JNCIE-ENT Certification Self-Study Bundle

Revision 1.b

Student Workbook and Lab Guide

Juniper Networks | Education Services

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

Course Number: EDU-JNCIE-ENT-BOOK

This document is produced by Juniper Networks, Inc.

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

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

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

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

Printed in USA.

Revision History:

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

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

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

YEAR 2000 NOTICE

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

SOFTWARE LICENSE

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

Contents

Contents 3
General Information 8
About this Publication 8
Education Services Offerings 8
Technical Publications 8
Juniper Networks Support 8
Target Audience 9
Exam strategy 9
Lab Preparation 9
Day Before the Exam 10
Exam Day 10
After the Exam 11
Lab Setup Information 11
Chapter 1: Junos System Features 16
Task 1: General System Services 16
Task 2: User Authentication and Authorization 17
Task 3: Syslog Configuration 18
Task 4: SNMP Configuration 18
Task 5: Junos Space 18
Task 6: Firewall Filters 19

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 2: Layer 2 Configuration and Features 20
Task 1: Layer 2 Switching 20
Task 2: Spanning Tree Configuration 22
Task 3: Layer 2 Features 22
Chapter 3: Protocol Independent Features 23
Task 1: GRE Tunnel 23
Task 2: Filter-Based Forwarding 24
Chapter 4: IGP 24
Task 1: OSPFv3 24
Task 2: OSPFv2 27
Task 3: IS-IS 28
Chapter 5: BGP 31
Task 1: IBGP Data Center 32
Task 2: IBGP Site 1 and Site 2 33
Task 3: EBGP Data Center 33
Chapter 6: EVPN/VXLAN 35
Task 1: EVPN VXLAN 35
Chapter 7: Class of Service 36
Task 1: vMX Class of Service 36
Task 2: EX Class of Service 39
Super Lab 1 42
Task 1: General System Features 46
Task 2: Layer 2 Configuration and Features 48

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: Protocol Independent Features 50
Task 4: IGP 51
Task 5: BGP 54
Task 6: EVPN/VXLAN (Spine and Leaf) 59
Task 7: Class of Service 60
Super Lab 2 64
Task 1: General System Features 68
Task 2: Layer 2 Configuration and Features 70
Task 3: Protocol Independent Features 72
Task 4: IGP and MPLS 74
Task 5: MPLS/LDP 78
Task 6: BGP 78
Task 7: VPN 82
Task 8: EVPN/VXLAN 83
Task 9: Class of Service 84
Appendices 87
Appendix – Additional Theory 88
MPLS Overview 88
Layer 3 VPN Overview 91
Appendix – Chapter 1: Junos System Features 95
Task 1: General System Services 95
Task 2: User Authentication and Authorization 99
Task 3: Syslog Configuration 106

© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: SNMP Configuration 109
Task 5: Junos Space 111
Task 6: Firewall Filters 119
Appendix – Chapter 2: Layer 2 Configuration and Features 128
Task 1: Layer 2 Switching 134
Task 2: Spanning Tree Configuration 148
Task 3: Layer 2 Features 156
Appendix – Chapter 3: Protocol Independent Features 162
Task 1: GRE Tunnel 162
Task 2: Filter-Based Forwarding 165
Appendix – Chapter 4: IGP 168
Task 1: OSPFv3 179
Task 2: OSPFv2 194
Task 3: IS-IS 198
Appendix – Chapter 5: BGP 214
Task 1: IBGP Data Center 225
Task 2: IBGP Site 1 and Site 2 236
Task 3: EBGP Data Center 240
Appendix – Chapter 6: EVPN/VXLAN 276
Task 1: EVPN VXLAN 280
Appendix – Chapter 7: Class of Service 292
Task 1: vMX Class of Service 297
Task 2: EX Class of Service 307

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Super Lab 1 330
Task 1: General System Features 330
Task 2: Layer 2 Configuration and Features 365
Task 3: Protocol Independent Features 380
Task 4: IGP 397
Task 5: BGP 417
Task 6: EVPN/VXLAN (Spine and Leaf) 463
Task 7: Class of Service 481
Appendix – Super Lab 2 494
Task 1: General System Features 494
Task 2: Layer 2 Configuration and Features 517
Task 3: Protocol Independent Features 538
Task 4: IGP 554
Task 5: MPLS/LDP 579
Task 6: BGP 588
Task 7: VPN 630
Task 8: EVPN/VXLAN 640
Task 9: Class of Service 659

© 2020 Juniper Networks, Inc. All rights reserved.

General Information

About this Publication

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

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

The JNCIE-ENT lab exam preparation workbook consists of 8 virtual MX Series devices running Junos version 18.3; 2 virtual SRX Series devices running Junos version 18.4; 2 virtual QFX Series devices running Junos version 18.4; 2 EX4300 devices running Junos version 18.4; 1 Junos Space with Network Director running release 18.4; and a Linux-based student desktop device.

Education Services Offerings

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

Technical Publications

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

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

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

Juniper Networks Support

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

© 2020 Juniper Networks, Inc. All rights reserved.

Target Audience

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

Exam strategy

Lab Preparation

To earn the JNCIE-ENT credential you must first pass the following Juniper Networks written exams: JNCIA-Junos, JNCIS-ENT, and JNCIP-ENT. Finally, you must also pass the very challenging 6 hour JNCIE-ENT lab exam. You must have extensive hands-on experience with enterprise routing and switching on vSRX, vMX, vQFX and EX platforms. When preparing for the JNCIE-ENT lab exam ensure you fully understand the exam topics and objectives. The latest update for the JNCIE-ENT lab exam topics and objectives can be found on the Juniper Networks certification website.

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

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

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

© 2020 Juniper Networks, Inc. All rights reserved.

Day Before the Exam

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

Exam Day

When you wake up, ensure you dress comfortable and take a good breakfast. This will be a very long day! Ensure you arrive 30 minutes before the exam start time. Ensure you do not forget to take a valid ID with you, exam registration confirmation and if you prefer your own keyboard. If you forget your ID, the JNCIE lab exam proctor will not allow you to take the exam. When it's time to enter the exam room, the JNCIE lab exam proctor will direct you to your workstation. Using the SecureCRT software you can access the JNCIE terminal server, which provides console and out of band access (fxp) to the exam devices.

When the proctor tells you that the exams begins, you are allowed to open the exam documents. Before immediately starting with the first question, we would like to advise you to read the entire exam and study the topology first to fully understand all requirements.

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

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

Do not assume any other requirement then mentioned in the tasks. If you are unsure about a task requirement then ask the proctor. While the JNCIE lab exam proctor cannot help you in solving the tasks, he or she can certainly help you in clarifying the task requirements.

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

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

© 2020 Juniper Networks, Inc. All rights reserved.

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

After the Exam

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

Lab Setup Information

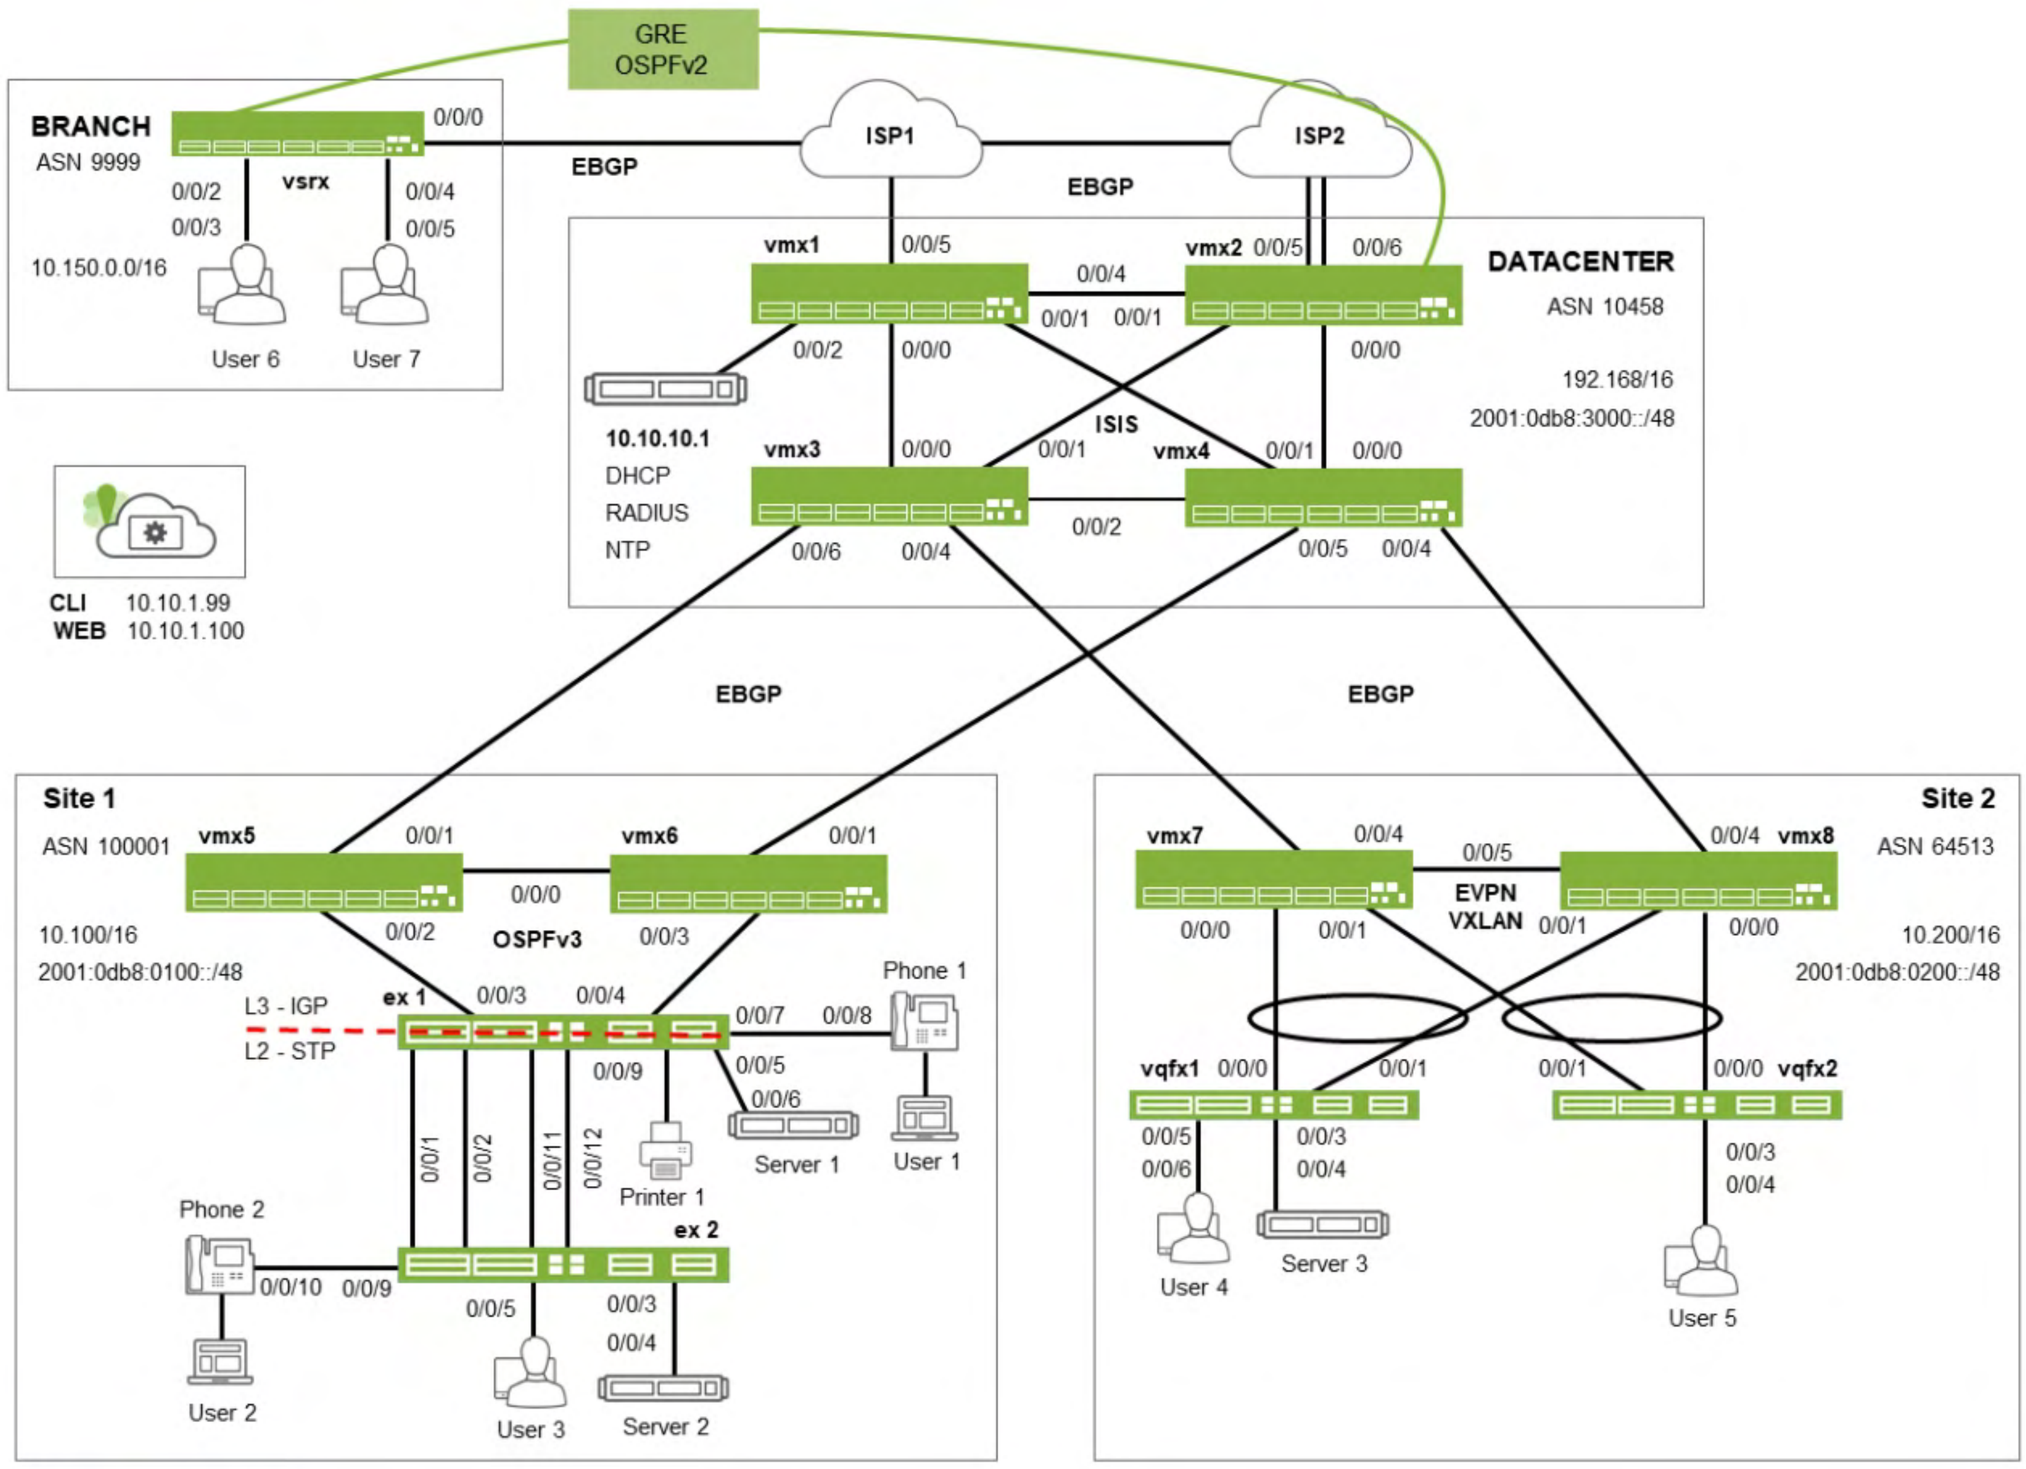

Figure 1 Topology

© 2020 Juniper Networks, Inc. All rights reserved.

In the figure above, the topology is shown that is used throughout this workbook. There are four vMX devices in the data center. vMX1 and vMX2 are both connected to an ISP. vMX3 and vMX4 are connected to vMX5 and vMX6 in Site1. Site 1 contains two EX4300 devices (EX1 and EX2) that have several devices attached, such as phones, printers, servers and users. Please note that some of these devices are simulated using routing-instances on the EX4300 switches. For example, server2 on EX2 has a layer3 port (ge-0/0/4) that participates in a routing-instance. There is a physical ethernet connection between port ge-0/0/4 and ge-0/0/3, where port ge-0/0/3 participates are a Layer 2 port in the inet.0 routing table. Printer1 and User2 are virtual machines that are accessible (see lab information). The data center vMX3 and vMX4 is also connected to Site2 vMX7 and vMX8. vMX7 and vMX8 have Layer 2 connections to vQFX1 and vQFX2. Like the EX4300 switches in Site1, the vQFX switches also have users and a server attached. There is also a branch office with two users that is connected to ISP1.

Device	OoB Interface	OoB IPv4 address
vMX1	fxp0	10.10.1.1/24
vMX2	fxp0	10.10.1.2/24
vMX3	fxp0	10.10.1.3/24
vMX4	fxp0	10.10.1.4/24
vMX5	fxp0	10.10.1.5/24
vMX6	fxp0	10.10.1.6/24
vMX7	fxp0	10.10.1.7/24
vMX8	fxp0	10.10.1.8/24
VR	fxp0	10.10.1.9/24
vSRX	fxp0	10.10.1.10/24
EX1	ge-0/0/0	10.10.1.11/24
EX2	ge-0/0/0	10.10.1.12/24
vQFX1	em0	10.10.1.13/24
vQFX2	em0	10.10.1.14/24
Junos Space	N/A	10.10.1.99/24 (CLI) 10.10.1.100/24 (WEB)

Table 2 Management Network (out of band)

The table above shows the devices used in the topology and their management IP addresses. These addresses are preconfigured and should not be changed and are listed here for informational purposes only.

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Interface	IPv4 address	IPv6 address
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128
	ge-0/0/0	192.168.13.1/24	2001:0db8:3000:0013::1/64
	ge-0/0/1	192.168.14.1/24	2001:0db8:3000:0014::1/64
	ge-0/0/2	10.10.10.100/24	N/A
	ge-0/0/4	192.168.12.1/24	2001:0db8:3000:0012::1/64
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128
	ge-0/0/0	192.168.24.2/24	2001:0db8:3000:0024::2/64
	ge-0/0/1	192.168.23.2/24	2001:0db8:3000:0023::2/64
	ge-0/0/4	192.168.12.2/24	2001:0db8:3000:0012::2/64
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64
	ge-0/0/6	172.16.3.2/24	2001:0dba:4000:0003::2/64
	GRE	173.17.99.2/24	
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128
	ge-0/0/0	192.168.13.3/24	2001:0db8:3000:0013::3/64
	ge-0/0/1	192.168.23.3/24	2001:0db8:3000:0023::3/64
	ge-0/0/2	192.168.34.3/24	2001:0db8:3000:0034::3/64
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128
	ge-0/0/0	192.168.24.4/24	2001:0db8:3000:0024::4/64
	ge-0/0/1	192.168.14.4/24	2001:0db8:3000:0014::4/64
	ge-0/0/2	192.168.34.4/24	2001:0db8:3000:0034::4/64
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128
	ge-0/0/0	10.100.56.5/24	2001:0db8:0100:0056::5/64
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64
	ge-0/0/2	10.100.15.5/24	2001:0db8:0100:015::5/64
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128
	ge-0/0/0	10.100.56.6/24	2001:0db8:0100:0056::6/64
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/3	10.100.116.6/24	2001:0db8:0100:0016::6/64
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128
	ge-0/0/0	Layer 2	Layer 2
	ge-0/0/1	Layer 2	Layer 2
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64
	ge-0/0/5	192.168.78.7/24	2001:0db8:3000:0078::7/64
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128
	ge-0/0/0	Layer 2	Layer 2
	ge-0/0/1	Layer 2	Layer 2
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64
	ge-0/0/5	192.168.78.8/24	2001:0db8:3000:0078::8/64
vSRX	lo0	10.150.0.99/32	
	ge-0/0/0	10.1.1.2/30	
	ge-0/0/2	Layer 2	
	ge-0/0/3	10.150.99.4/24	
	ge-0/0/4	Layer 2	
	ge-0/0/5	10.150.99.5/24	
	GRE	173.17.99.1/24	
EX1	lo0	10.100.0.1/32	2001:0db8:0100:0000::1/128
	ge-0/0/1	Layer 2	
	ge-0/0/2	Layer 2	
	ge-0/0/3	10.100.15.1/24	2001:0db8:0100:015::1/64
	ge-0/0/4	10.100.16.1/24	2001:0db8:0100:016::1/64
	ge-0/0/5	Layer 2	
	Ge0/0/6	dhcp	2001:0db8:0100:101::1/64
	ge-0/0/7	Layer 2	
	ge-0/0/8	10.100.201.1/24	2001:0db8:0100:201::1/64
	ge-0/0/11	Layer 2	
	ge-0/0/12	Layer 2	
	AE10	Layer 2	
	irb.100	10.100.100.254/24	2001:0db8:0100:100::254/64
	irb.101	10.100.101.254/24	2001:0db8:0100:101::254/64
	irb.200	10.100.200.254/24	2001:0db8:0100:200::254/64
	irb.201	10.100.201.254/24	2001:0db8:0100:201::254/64

© 2020 Juniper Networks, Inc. All rights reserved.

EX2	ge-0/0/1	Layer 2	
	ge-0/0/2	Layer 2	
	ge-0/0/3	Layer 2	
	ge-0/0/4	dhcp	
	ge-0/0/5	Layer 2	
	ge-0/0/6	Layer 2	
	ge-0/0/9	Layer 2	
	ge-0/0/10	10.100.201.2/24	2001:0db8:0100:201::2/64
	ge-0/0/11	Layer 2	
	ge-0/0/12	Layer 2	
	AE10	Layer 2	
vQFX1	xe-0/0/0	Layer 2	
	xe-0/0/1	Layer 2	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.101.1/24	2001:0db8:0200:101::1/64
	xe-0/0/5	Layer 2	
	xe-0/0/6	10.200.100.1/24	2001:0db8:0200:100::1/64
vQFX2	xe-0/0/0	Layer 2	
	xe-0/0/1	Layer 2	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.100.3/24	2001:0db8:0200:100::3/64

Table 3 Management Network (out of band)

The table above shows the IP addressing scheme that is used throughout the workbook chapters.

Important: To access the devices in this environment please use the lab/lab123 user credentials. For access to the Junos Space Web UI use the super/lab123 credentials.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 1: Junos System Features

This chapter focuses on initial system configuration and general system features. You will configure various tasks, such as configuring user authentication and authorization, NTP, SNMP, Syslog and RE protection.

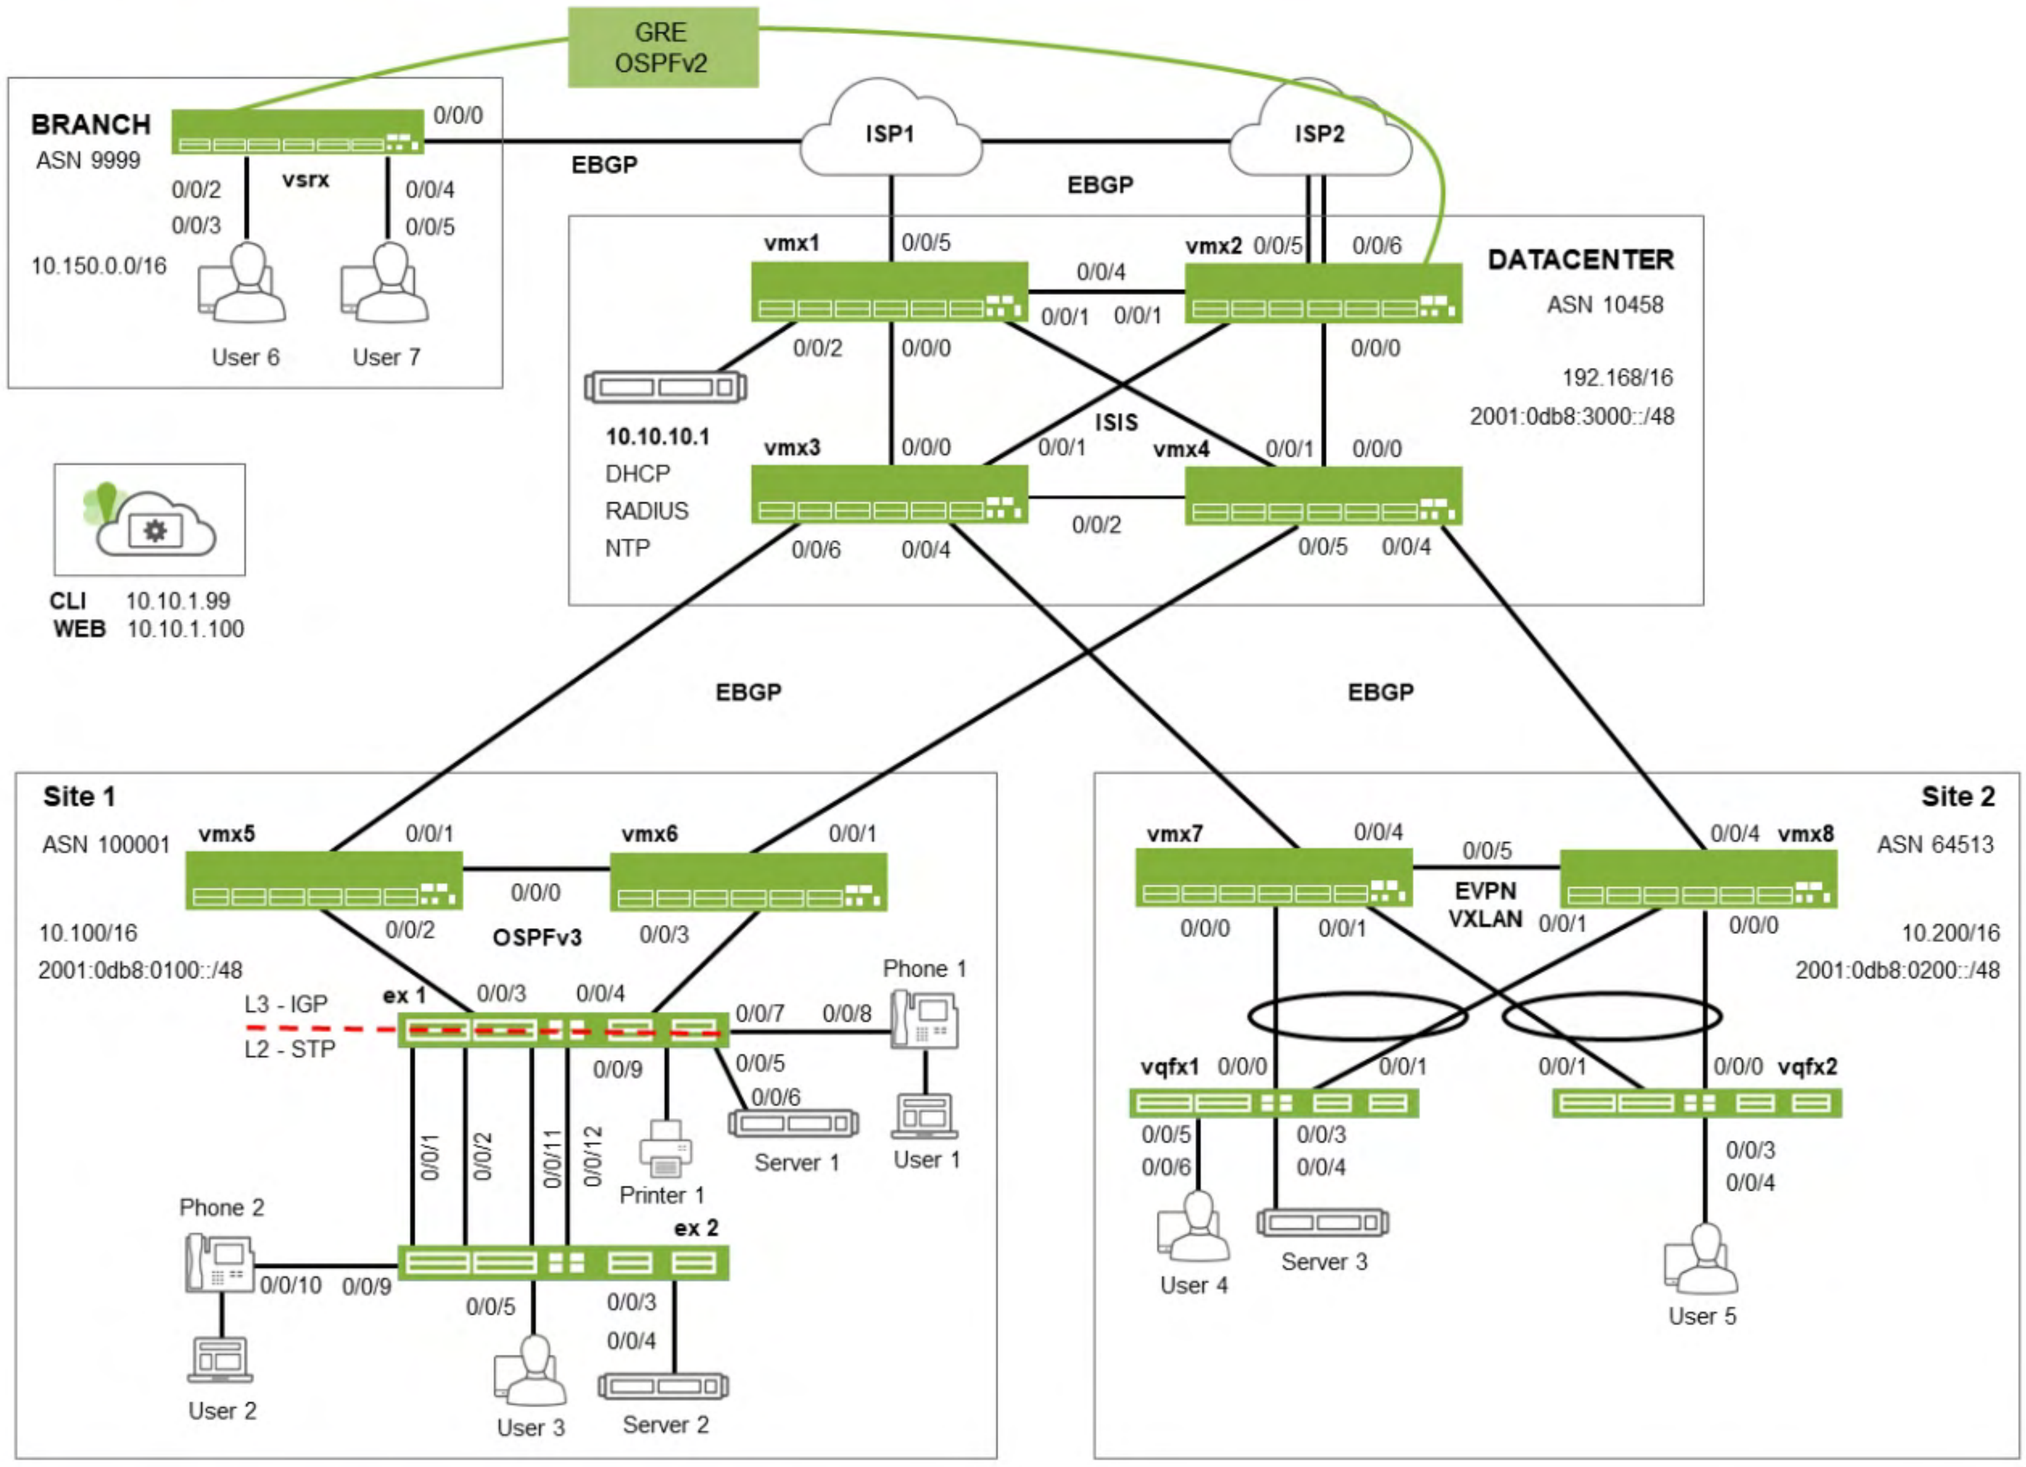

Task 1: General System Services

Note: For this task you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place.

1) Configure all vMX and EX devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute

© 2020 Juniper Networks, Inc. All rights reserved.

2) Configure server 10.10.10.1 as the DNS server.
3) Set the time zone to *America/Los_Angeles* on all vMX devices and EX1.
4) Make sure that all vMX devices and EX1 synchronize the time with NTP server 10.10.10.1.
5) Configure all vMX devices and EX1 transfer their configuration to the FTP server 10.10.10.1 each time the configuration is committed. You can use user name *lab* and password *lab123* for the FTP server access.

Note: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1*

Task 2: User Authentication and Authorization

In this part you will configure new users allowed to access the devices and define their privileges and permissions.

The following tasks should be configured on all vMX devices and EX1

1) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts 3 and timeout of 3 seconds. Use juniper123 as the RADIUS shared secret.
2) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices, EX1 as a super-user.
3) Configure additional users on all the devices as defined in Table 4.
4) Make sure that for all users in Table 4 sessions are automatically logout after 5 minutes of idle time-out. Ensure that new user classes will automatically inherit this setting.

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

Username	Password	Privileges
noc	noc123	Permissions “all”. Additionally, user cannot execute any of the “clear”, “configure”, “edit” or “start shell” commands.
ops	ops123	Can view “show interfaces” with all options and can execute “ping” and “traceroute” commands.

Table 4

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: Syslog Configuration

Ensure that all vMX devices and EX1 have the following Syslog configuration:

1) All “critical” messages regardless of the facility are displayed on terminals of all currently logged users.
2) All messages regardless of the facility with severity level of “info” and higher are sent to the default syslog file. Make sure that the repeated messages are individually logged.
3) A file named “interactive-commands” for command audit tracking receives records about the users and commands they execute.
4) All messages with severity level “info” and higher are sent to the syslog server 10.10.10.1. Log messages should be in a structured format
5) Ensure that the maximum archived log files are 10. Each logfile can contain up to 5 Mbyte of log information.

Task 4: SNMP Configuration

In this task, you will configure SNMPv2.

The following tasks should be configured on all vMX, vQFX and vSRX devices in the topology.

1) The SNMP read-only community string is jncie. Make sure that only the Junos Space server can poll the devices in the data center, Site1, Site2 and the branch device.
2) Configure all devices to send SNMP traps to Junos Space for link and/or chassis events. Make sure that traps are only send using the OoB management interface.

Task 5: Junos Space

In this task, you will configure the Junos Space Network Director

1) You have a Junos Space Network Director instance running for device management. You must discover all vMX devices using Junos Space Network Director with SNMPv2 and the community string jncie. Ensure that the devices are present under the Inventory window.

Note: Junos Space can be accessed by browsing to the management address using HTTPS.

2) Back up all vMX configurations onto Junos Space Network Director.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 6: Firewall Filters

In this task, you will configure RE (Routing Engine) protection using a firewall filter for all vMX devices in the data center.

1) Configure an IPv4 firewall filter which allows protocol messages from the BFD, VRRP, OSPF, BGP and GRE. Make sure that for the protocols VRRP and OSPF messages from directly connected subnets are accepted.
2) Configure the firewall filter to accept RADIUS, DNS, SNMP management protocols only from the 10.10.10/24 network.
3) Accept NTP messages only from the 10.10.10/24 and 10.10.1/24 networks.
4) Accept SSH and Telnet protocols only from the 10.10.1/24 management network.
5) Configure the devices to accept ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer, which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
6) Protect the devices from TCP syn/rst/fin attacks. Use a policer with the following parameters: bandwidth 500kbps and burst size 125kB. The excess traffic must be dropped.
7) Configure the firewall filter to discard any other traffic and count the packets.
8) Apply the firewall filter on vMX1 such as to ensure that it is used for the RE protection.
9) Configure an IPv6 firewall filter which allows protocol messages from OSPF and BGP.
10) Configure the devices to accept IPv6 ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
11) Configure the IPv6 firewall filter to discard any other traffic and count the packets.
12) Apply the IPv6 firewall filter on vMX1 such as to ensure that it is used for the RE protection.
13) Ensure that the configuration under firewall hierarchy is not displayed when printed out.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 2: Layer 2 Configuration and Features

Task 1: Layer 2 Switching

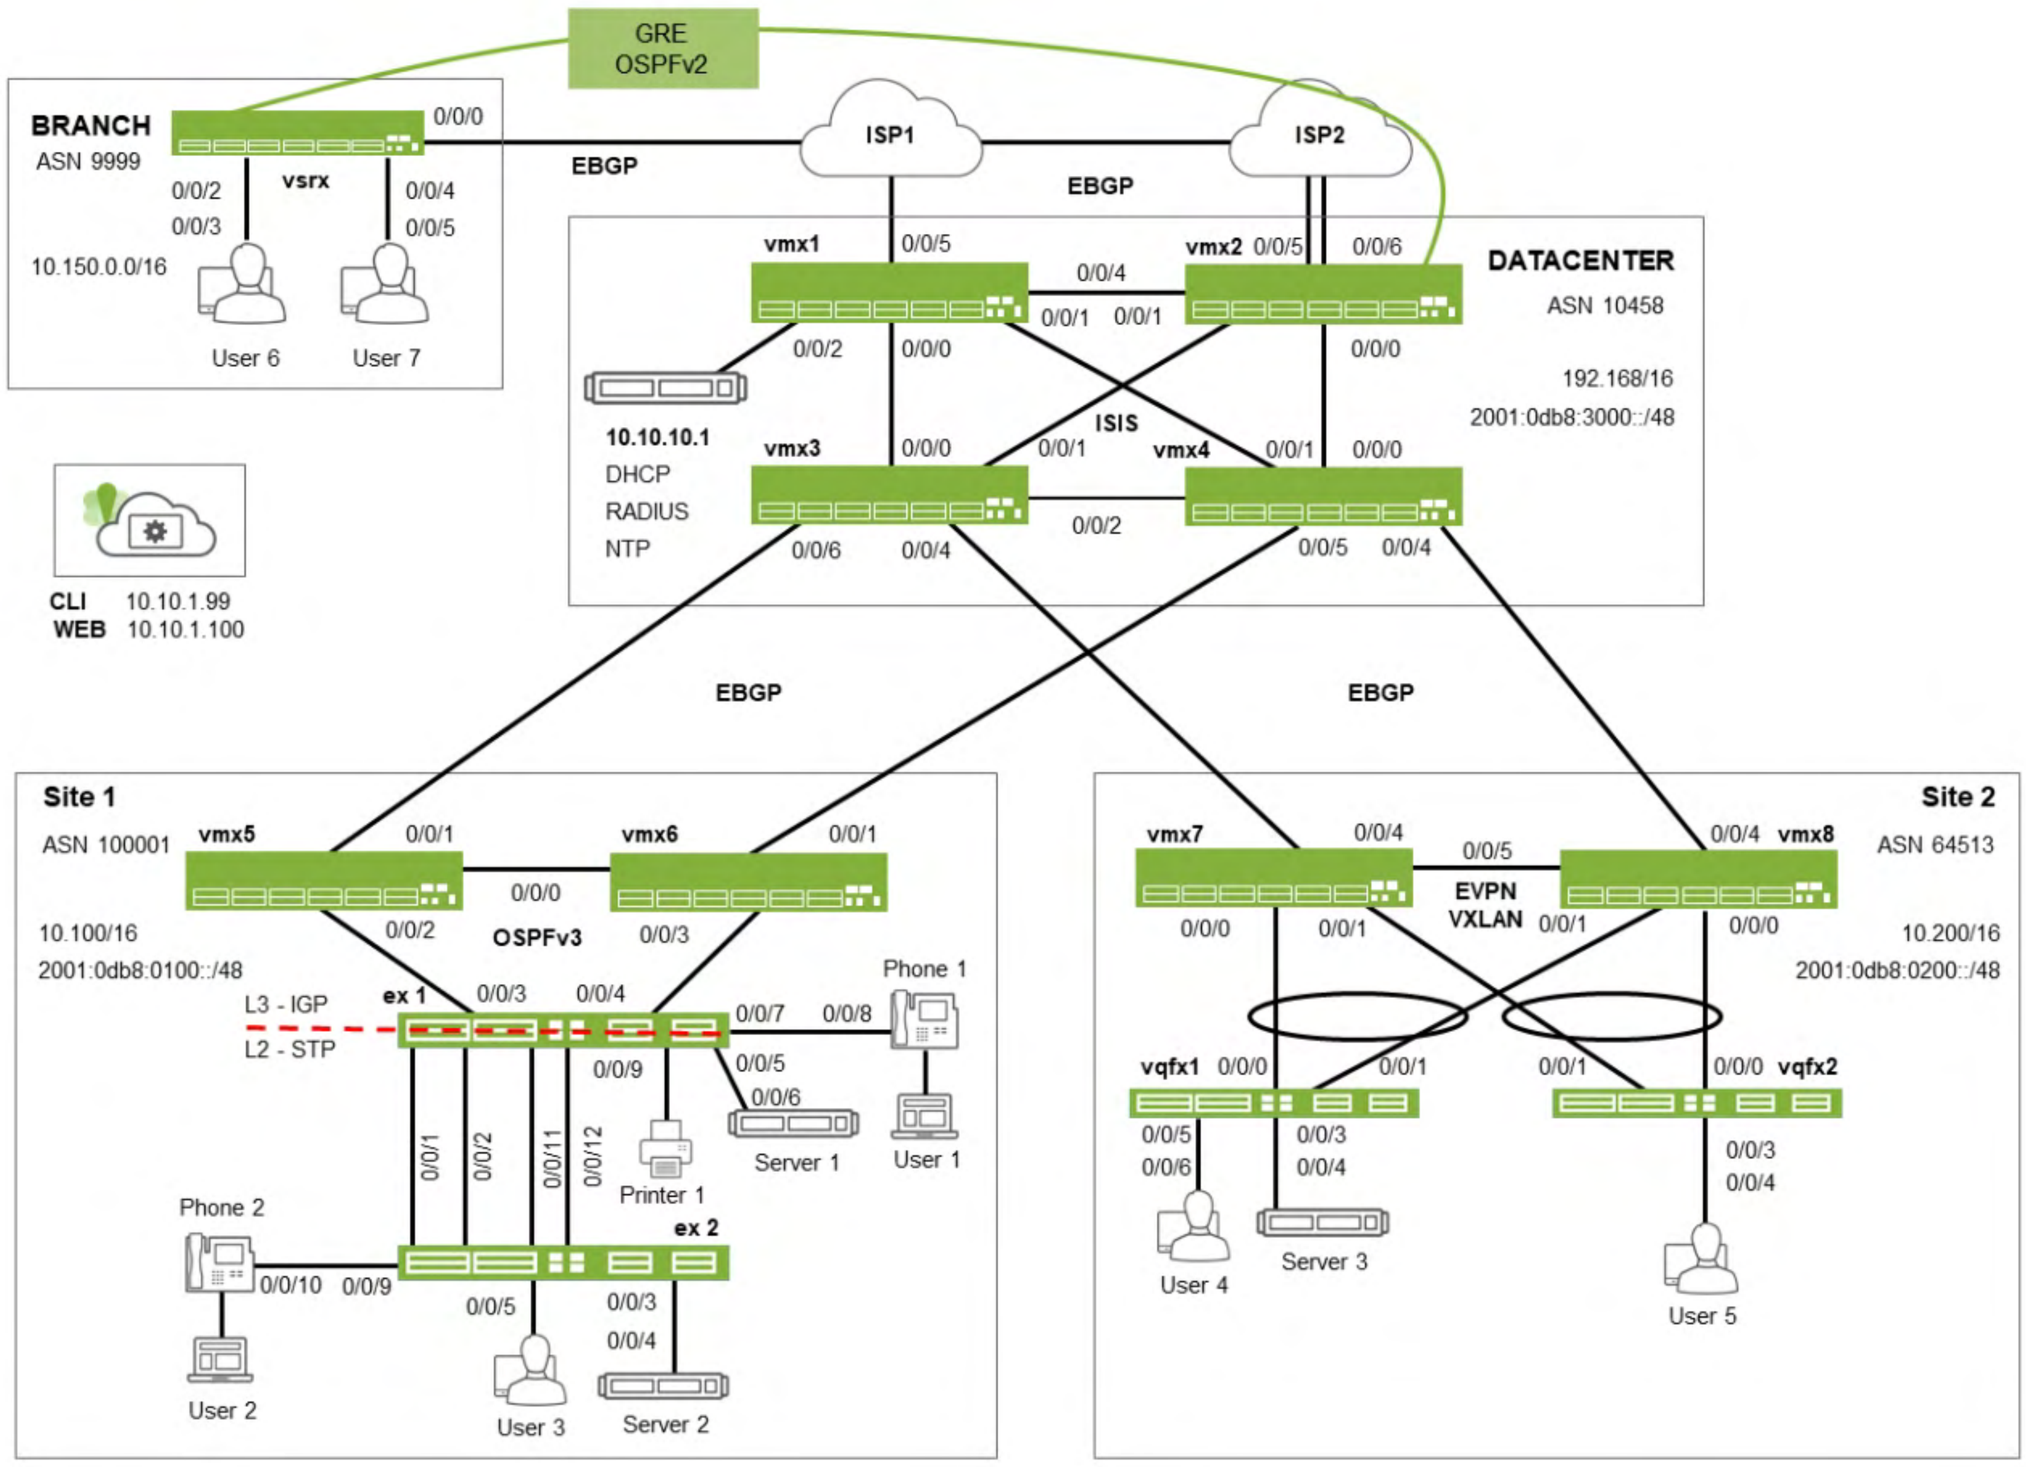

© 2020 Juniper Networks, Inc. All rights reserved.

VLAN Name	VLAN ID	Devices	Interface
user	100	EX1	ge-0/0/7 untagged
		EX2	ge-0/0/5 untagged
		EX2	ge-0/0/9 untagged
		vQFX1	ge-0/0/5 untagged
		vQFX2	ge-0/0/3 untagged
server	101	EX1	ge-0/0/5 untagged
		EX2	ge-0/0/3 untagged
		vQFX1	ge-0/0/3 untagged
printer	200	EX1	ge-0/0/9 untagged
phone	201	EX1	ge-0/0/7 tagged
		EX2	ge-0/0/9 tagged
guest	666	EX1, EX2	Dot1x guest vlan

Table 1 VLANs Site1 and Site 2

1) Configure the VLANs and corresponding interfaces in Site1 and Site 2 as shown in the table above
2) Configure aggregate ethernet interface (AE10) between EX1 and EX2. Ensure that the aggregate interface is shutdown if the bandwidth is less than 2Gbit.
3) The interfaces between EX1 and EX2 should be configured to transport all VLANs as shown in the table above.
4) Configure aggregate ethernet interface (AE0) on vQFX1 and vQFX2 towards the vMX devices. vQFX1 should only respond to LACP message.
5) Configure aggregate ethernet interface (AE0) on vMX7 and vMX8 towards vQFX1. Configure aggregate ethernet interface (AE1) on vMX7 and vMX8 towards vQFX2. Enable LACP on the AE interfaces. All Site 2 VLANs must be allowed on the AE0 interfaces.
6) Ensure that VLAN1 is transported untagged on all trunk interfaces in Site 1. This VLAN must be named ‘default’.
7) Configure interface ge-0/0/7 on EX1 and interface ge-0/0/9 on EX2 for Voice VLAN support. Make sure that tagged frames go to VLAN201 whereas untagged frames go to VLAN100.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Spanning Tree Configuration

In this task, you will configure the Spanning Tree protocol and various features.

1) Configure EX1 and EX2 in MSTP region jncie with two instances: Instance 1 must be bound to VLAN100-199. Instance 2 must be bound to VLAN200-299. Ensure that Instance 1 is rooted at EX1 and Instance 2 is rooted at EX2. The CIST root is at EX1.
2) Make sure that for the CIST interface AE10 is the root port on EX2. Only if interface AE10 and interface ge-0/0/2 are down the interface ge-0/0/1 should be the root port.
3) Make sure that for Instance 1 interface ge-0/0/2 is root port on EX2. Only if interface ge-0/0/1 and ge-0/0/2 are down the interface AE10 can be the root port.
4) Ensure that ports ge-0/0/5 , ge-0/0/7 and ge-0/0/9 on EX1 and ports ge-0/0/3, ge-0/0/5 and ge-0/0/9 on EX2 transition immediately to L2 forwarding state. Make sure that if BPDUs are received on any of those interfaces they will be automatically disabled. The edge port interfaces must be grouped under the MSTP configuration using the name “access-ports”.

Task 3: Layer 2 Features

1) Ensure Server1 and server 2 should receive an IPv4 address from the DHCP server.
2) Printer1 must be authenticated dynamically, but EX1 must not attempt 802.1x authentication. If authentication fails or the server is unreachable, the printer must be given guest access.
3) Configure interface ge-0/0/7 on EX1 and interface ge-0/0/9 on EX2 for POE and LLDP-MED support. Block location-id TLVs. No more than 20 WATT should be allocated the POE ports.
4) If more than 10 MAC addresses are learned on EX2 port ge-0/0/9 it should drop traffic.
5) Configure OAM LFM on the ge-0/0/1 and ge-0/0/2 interfaces between EX1 and EX2. Make sure EX1 and EX2 initiate the discovery process. Ensure OAM PDUs are sent every 600ms. If 5 PDUs are missed, the peer should be declared lost/down.
6) Configure EX1 to shutdown a port if Unknown Unicast or Broadcast traffic consumes more than 5% of the interface bandwidth. Multicast traffic should not be included in your solution.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 3: Protocol Independent Features

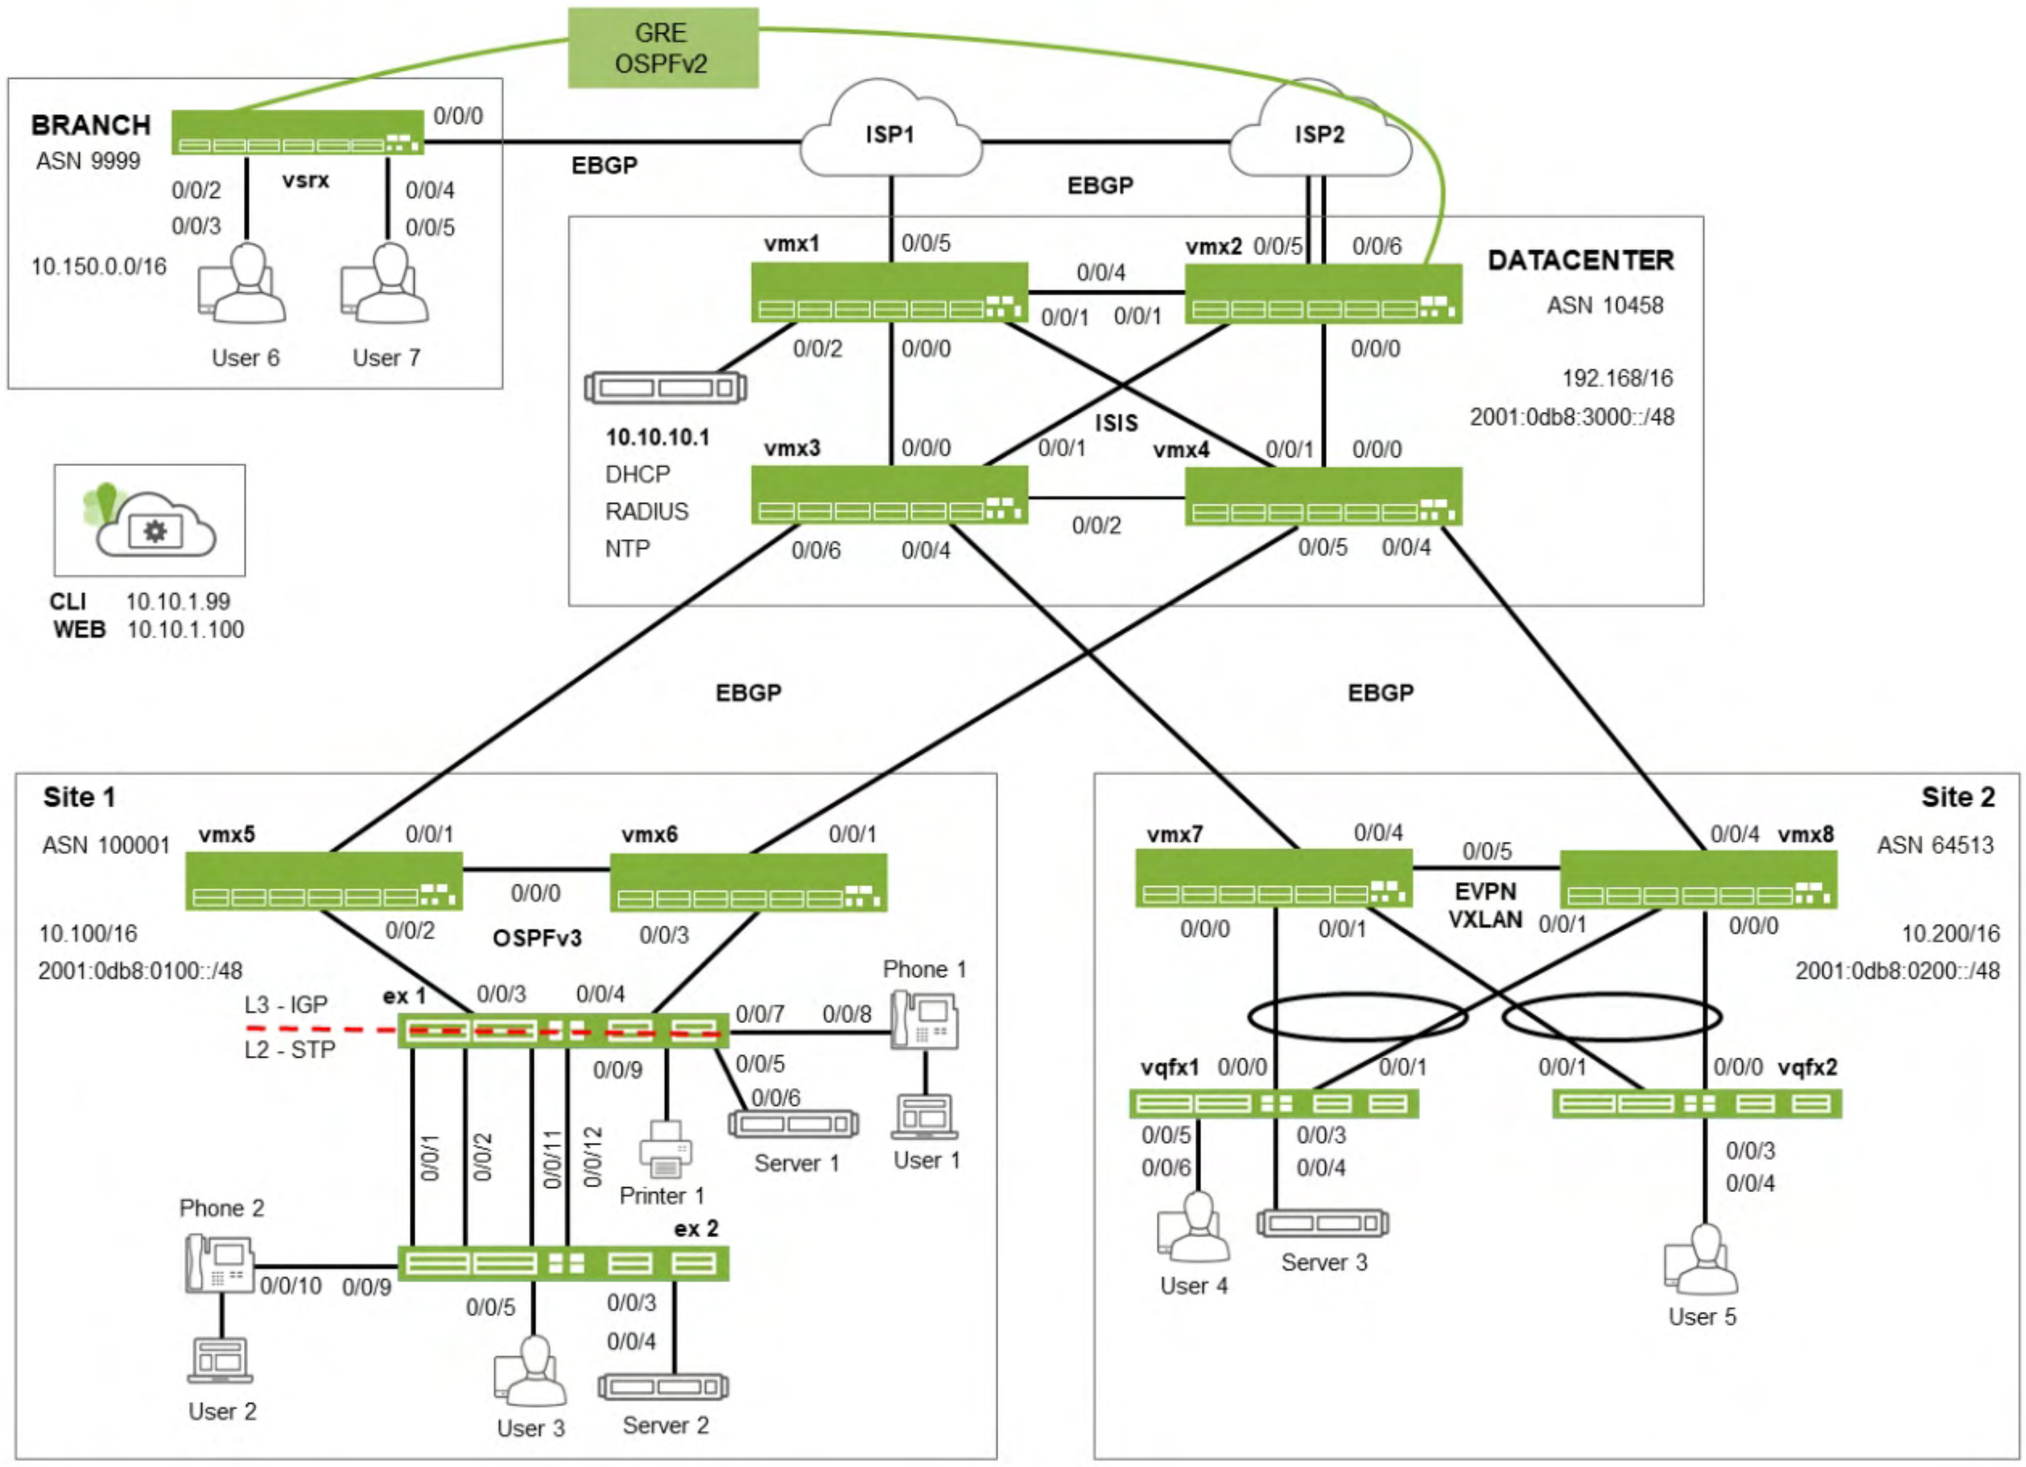

Task 1: GRE Tunnel

In this task, you will configure a GRE tunnel between the data center and the branch office.

1) Configure a GRE tunnel between the branch office and vMX2 in the data center. Use the IP addressing as shown in the Lab setup Information chapter. Use the loopback interface as GRE tunnel source address.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Filter-Based Forwarding

1) Ensure that vMX2 forwards IPv4 webtraffic received from Site1 destinated to 6.83.176.1 using the ge-0/0/6 interface. If interface ge-0/0/6 is unavailable, then interface ge-0/0/5 must be used. Other traffic should follow the routing requirements as mentioned in the other tasks

Chapter 4: IGP

Task 1: OSPFv3

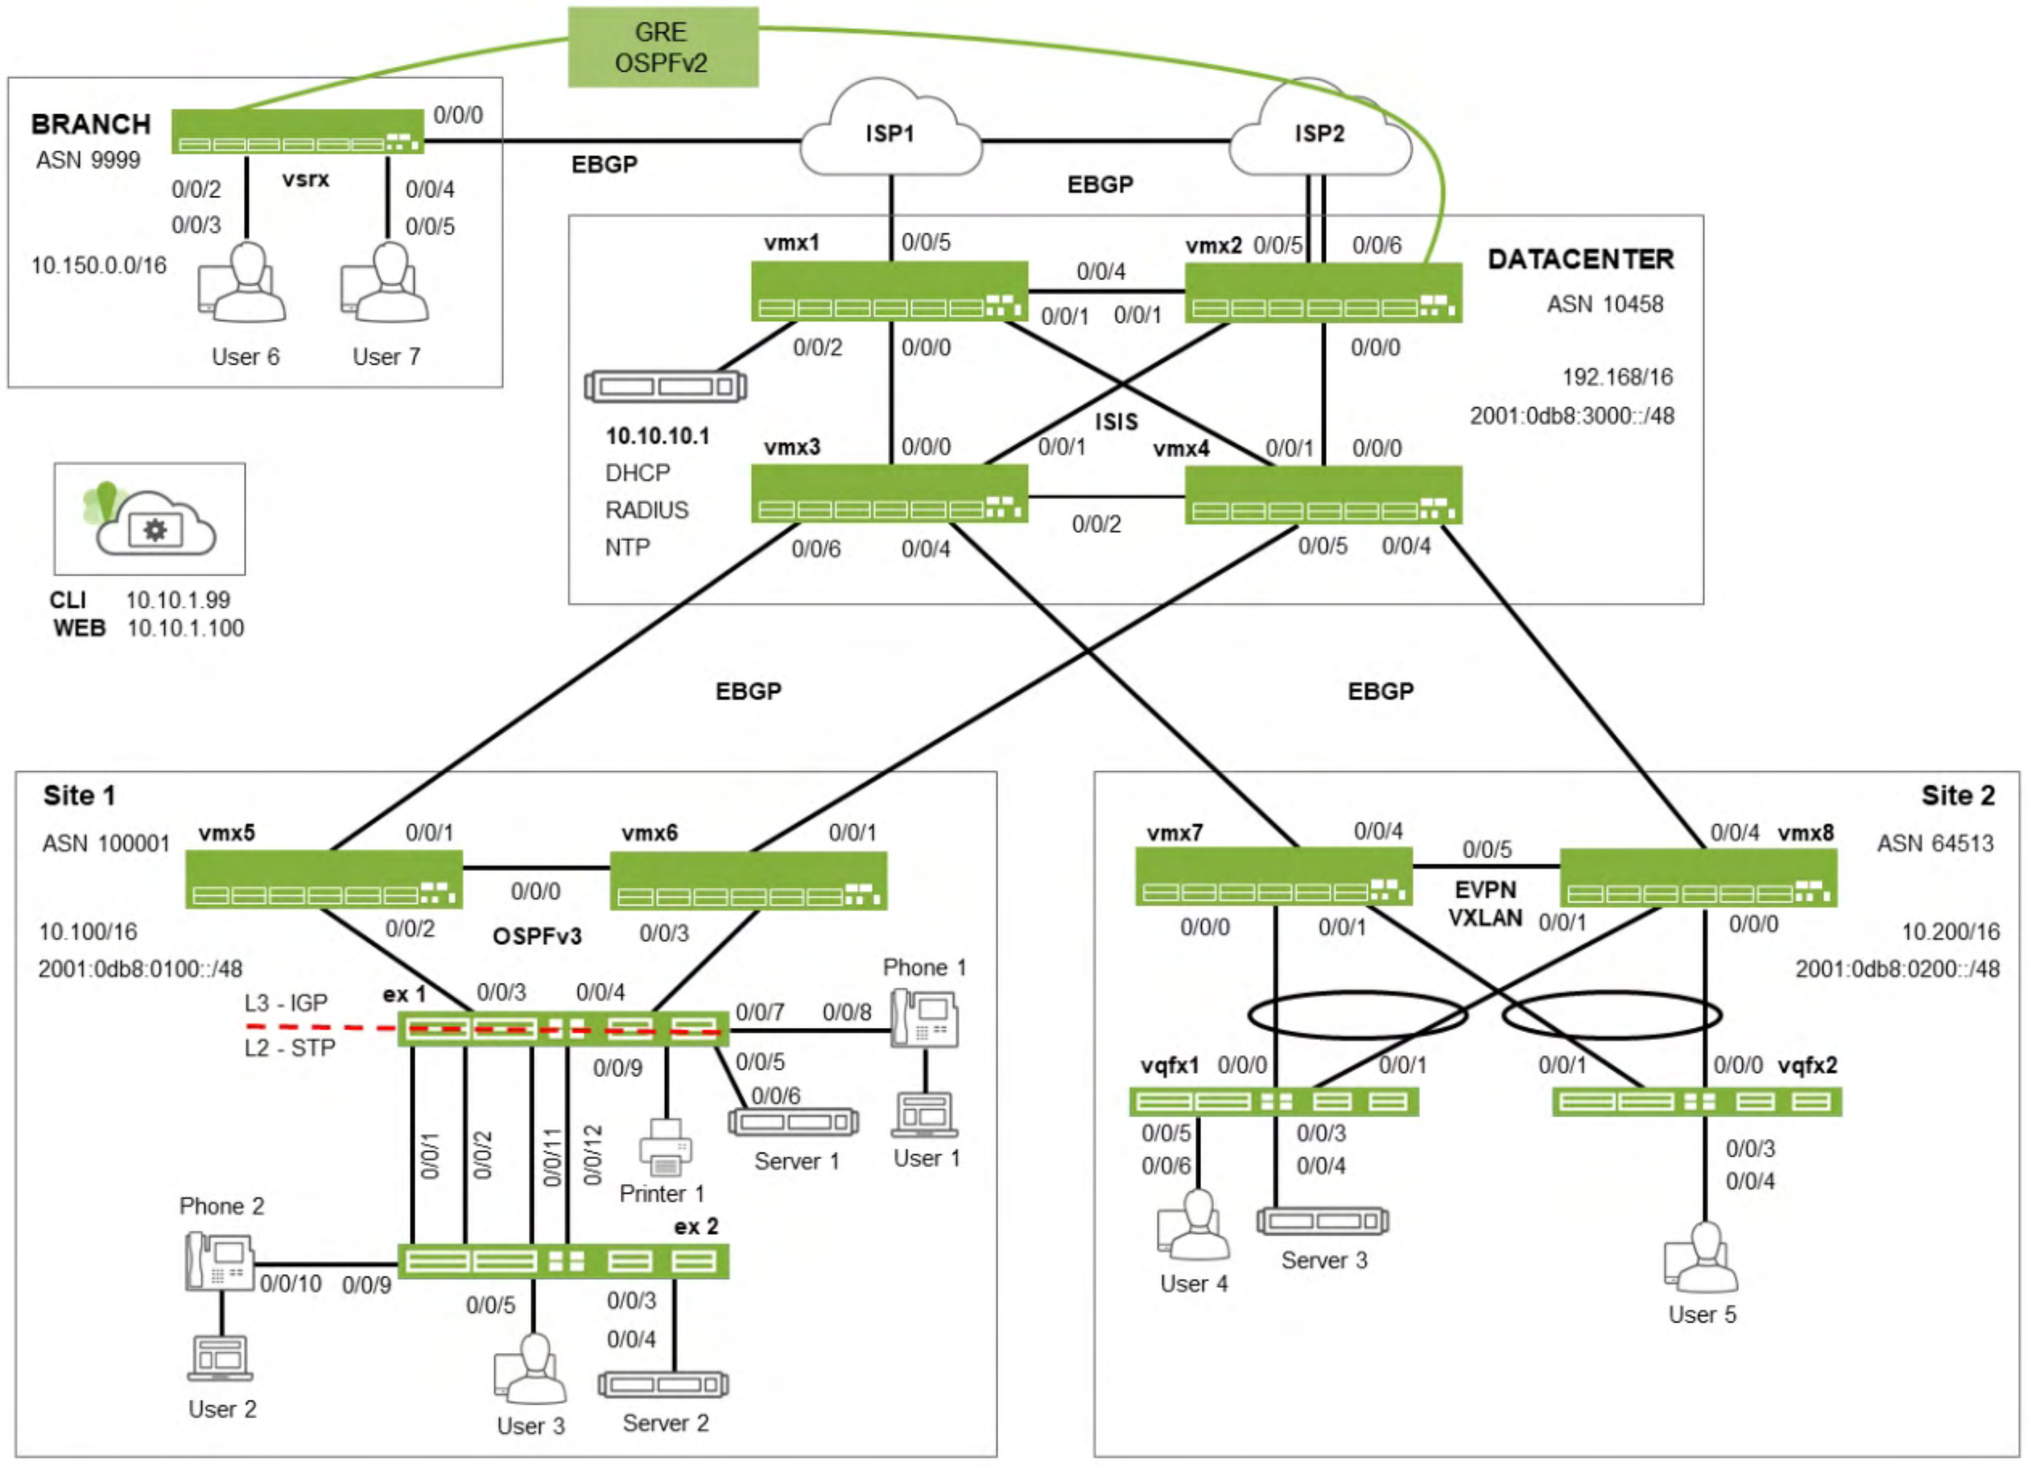

© 2020 Juniper Networks, Inc. All rights reserved.

In this task, you will configure OSPFv3 in Site 1 and Site 2.

Device	Interface	area
vMX5	lo0	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/2	0.0.0.1
vMX6	lo0	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/3	0.0.0.1
EX1	lo0	0.0.0.1
	ge-0/0/3	0.0.0.1
	ge-0/0/4	0.0.0.1

Table 2 OSPF Information Site 1

Site 1:

1) Configure OSPFv3 in Site 1 using the information as shown in the tables above.
2) Ensure type 0x2002 LSAs are not found in the OSPFv3 databases in Site 1.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four seconds
4) Area 1 should contain Intra Area Prefixes LSAs for the IPv4 IRB addresses of EX1
5) Area 1 should contain External LSAs for the IPv6 IRB addresses of EX1.
6) Area 0 should contain an InterArea Prefix LSA that encompasses the IPv4 IRB addresses of EX1. The individual IPv4 networks of EX1 IRB interface should not exist in Area 0.
7) No more than 1000 non self-generated LSAs are allowed for IPV4. Only generate a warning if this limit is exceeded
8) The time to wait before running an SFP calculation should be 50 ms for IPV4 and 200ms for IPV6.
9) Ensure that self-generated LSAs exchanged between the vMX devices in Site1 are only refreshed if the content of the LSA changes.

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Interface	area
vMX7	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0
vMX8	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0

Table 3 OSPF Information Site 2

Site 2:

1) Configure OSPFv3 in Site 2 using the information as shown in the tables above.
2) Ensure no LSA type 0x2002 are found in the OSPFv3 databases in Site 1.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four seconds.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: OSPFv2

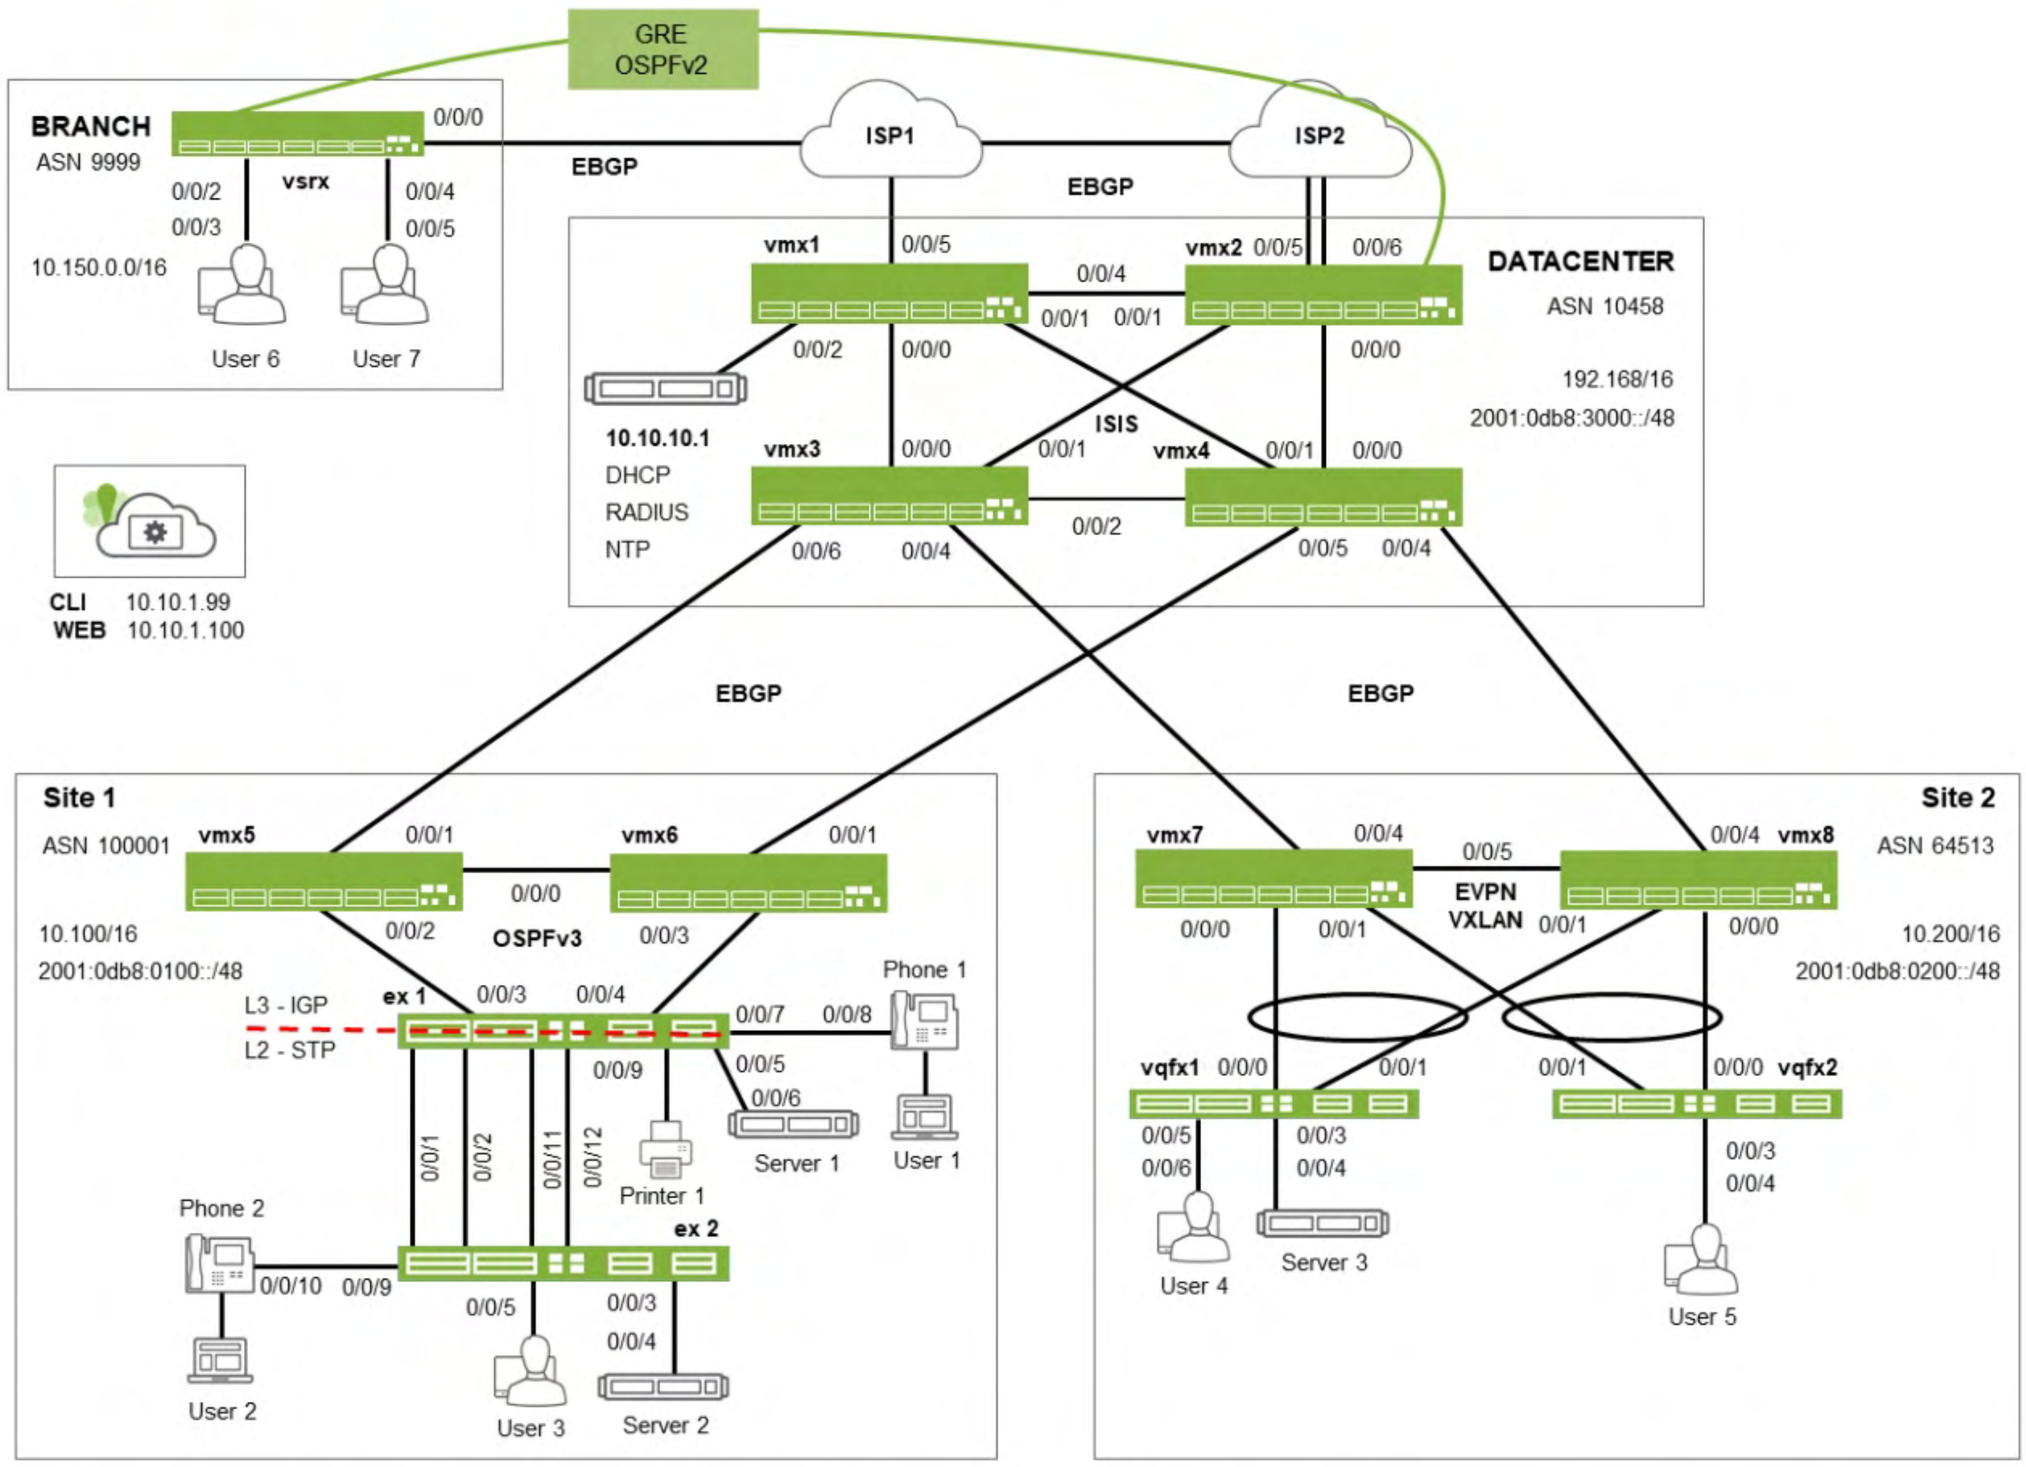

In this task you will configure a OSPFv2 over the GRE tunnel between the data center and the branch office.

Device	Interface	area
vMX2	GRE interface	0.0.0.0
vSRX	GRE interface	0.0.0.0

Table 4 OSPF Information for GRE Tunnel

1) Enable OSPFv2 over the GRE interface as shown in the table above.
2) Ensure the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.

© 2020 Juniper Networks, Inc. All rights reserved.

3) vSRX should receive a default route and the data center IPv4 aggregate route from vMX2. Both routes should exist in the OSPFv2 database.
4) Configure OSPF in such a way that vSRX only accepts the default route from vMX2.

Task 3: IS-IS

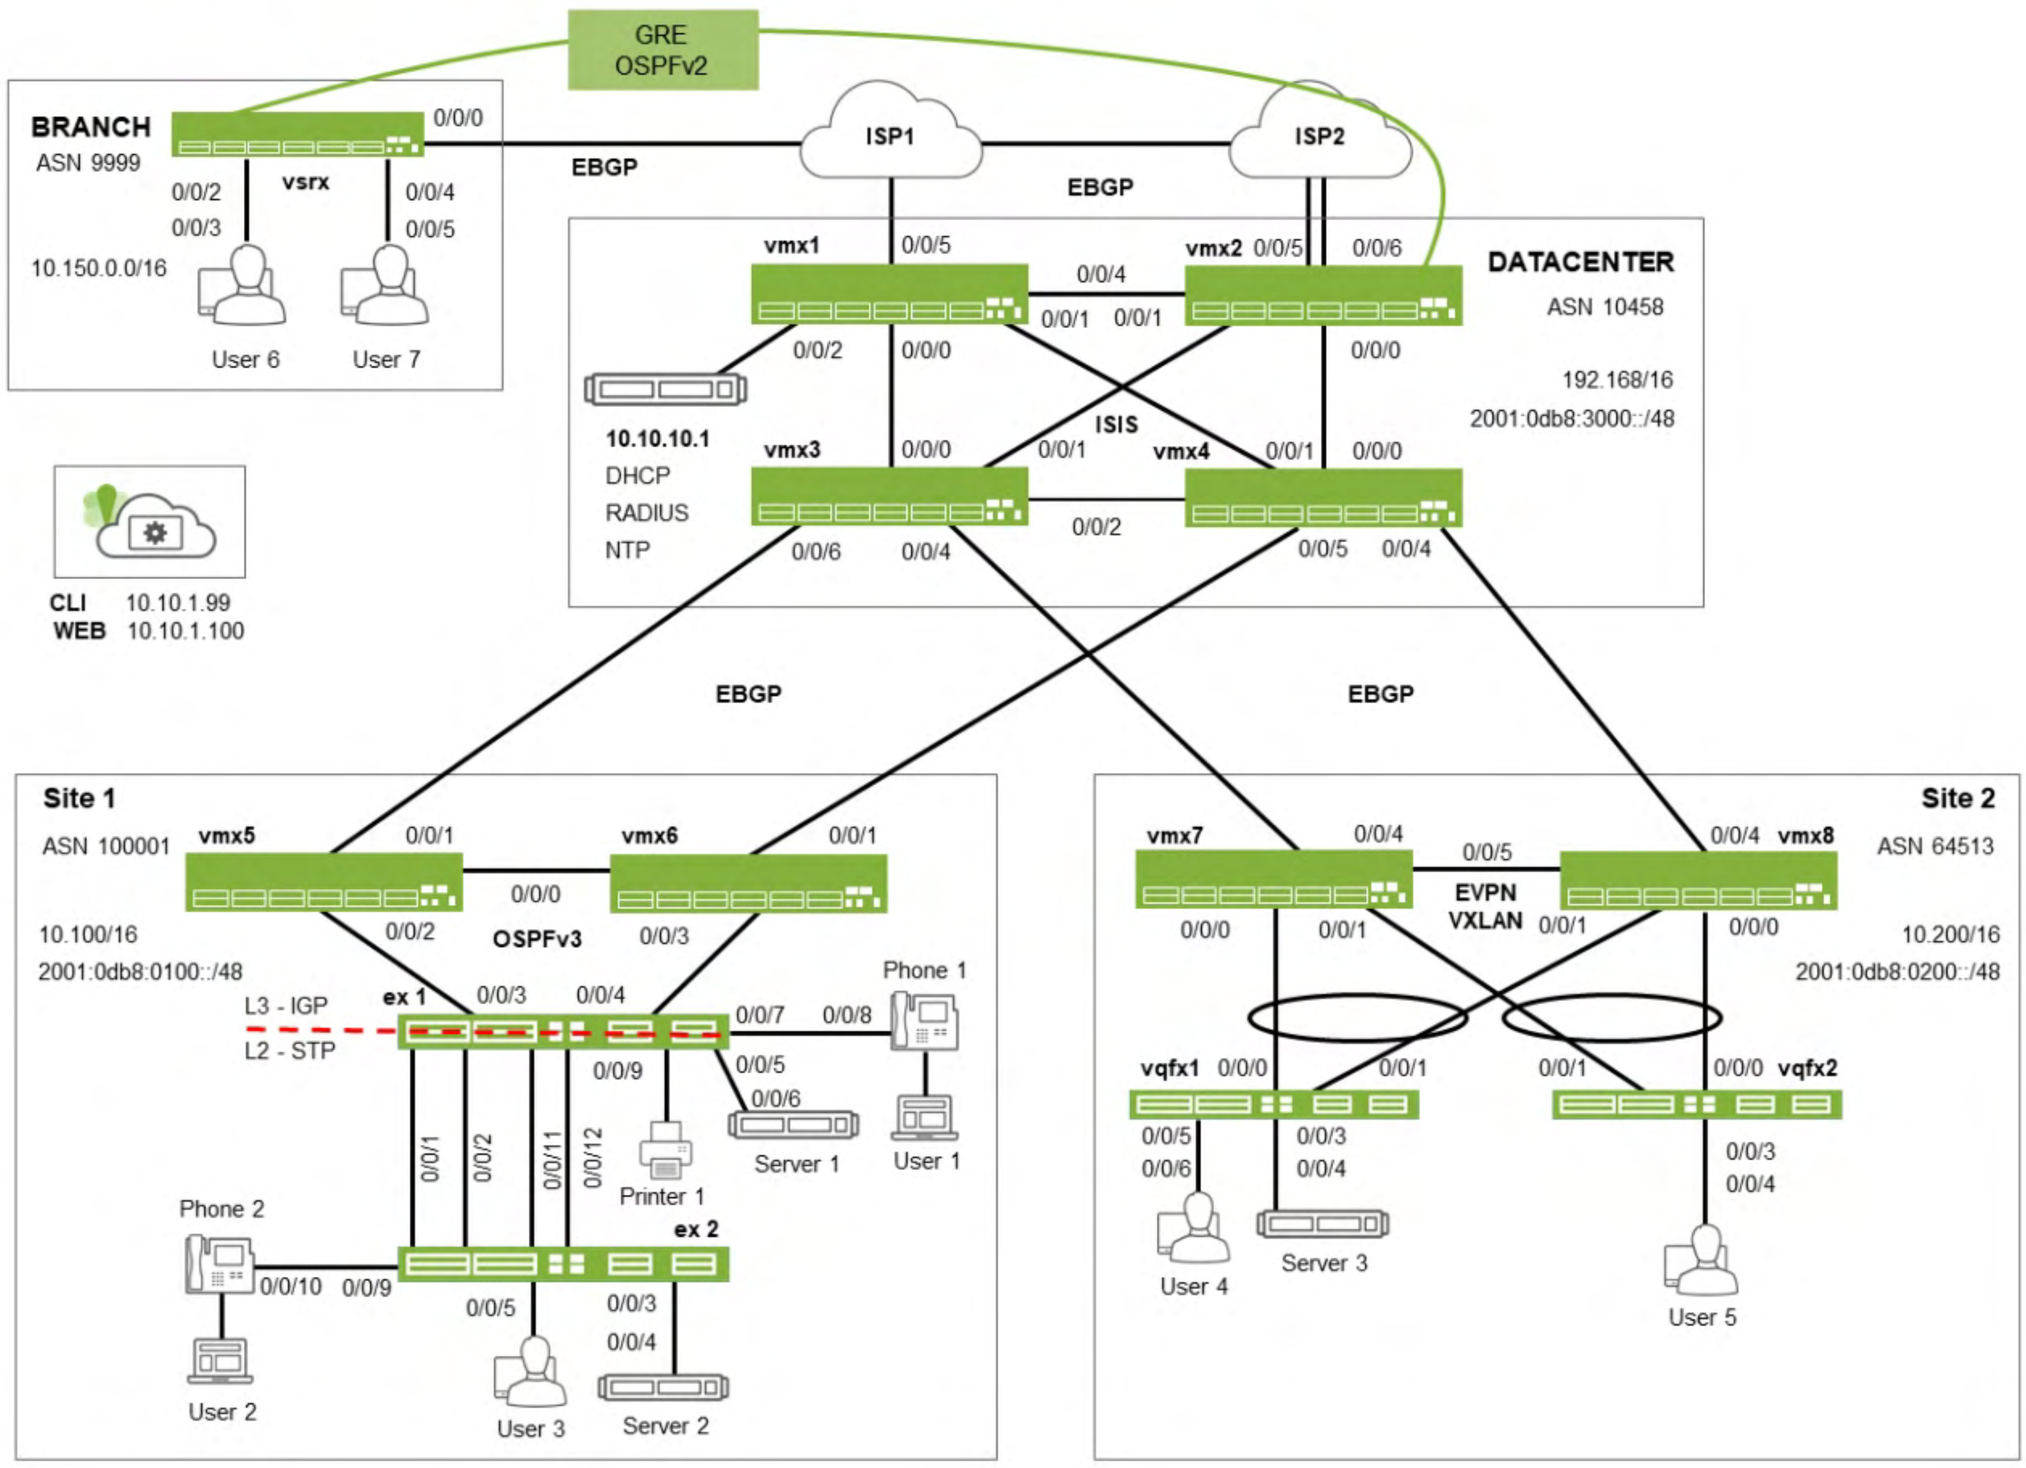

© 2020 Juniper Networks, Inc. All rights reserved.

In this task, you will configure IS-IS as the IGP in the data center.

1) Configure IS-IS in the data center using the information as shown in the tables below.

Device	Level	Interfaces
vMX1	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
		ge-0/0/2
vMX2	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
vMX3	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2
vMX4	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2

Table 5 IS-IS Information

Device	Required IS-I
vMX1	Sysid: 0168.0000.0001 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, Ipv6 is enabled
vMX2	Sysid: 0168.0000.0002

© 2020 Juniper Networks, Inc. All rights reserved.

	Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled
vMX3	Sysid: 0168.0000.0003 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled
vMX4	Sysid: 0168.0000.0004 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled

2) Do not enable Level 1 routing.
3) Ensure that no DIS is elected.
4) Ensure interface ge-0/0/2 on vMX1 participates in IS-IS, but no adjacency can be formed.
5) The maximum number of routes that can be redistributed in IS-IS should be 100.
6) Enable MD5 authentication for level 2. Use “juniper123” as the key.
7) Configure the hello and hold interval to 1 and 4 for the IS-IS level 2 interface between vMX1 and vMX3.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 5: BGP

This chapter focuses on the implementation of BGP for routing within the data center and within Site 2. BGP is also used for routing between the data center, ISP1, ISP2 and between the data center, Site 1 and Site 2. At various points in the topology you need to perform route redistribution between and IGP and BGP for end to end connectivity.

Device	Interface	Ipv4 address	Ipv6 address	Description
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128	IBGP source interface v4 EBGP to AS 2222
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64	v6 EBGP to AS 2222
	ge-0/0/6	172.16.3.2/24	Link local	v6 EBGP to AS 2222
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128	IBGP source interface
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64	EBGP to AS 64513
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64	EBGP to AS 100001
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128	IBGP source interface
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64	EBGP to AS 64513
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64	EBGP to AS 100001
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table 6 BGP peering information

© 2020 Juniper Networks, Inc. All rights reserved.

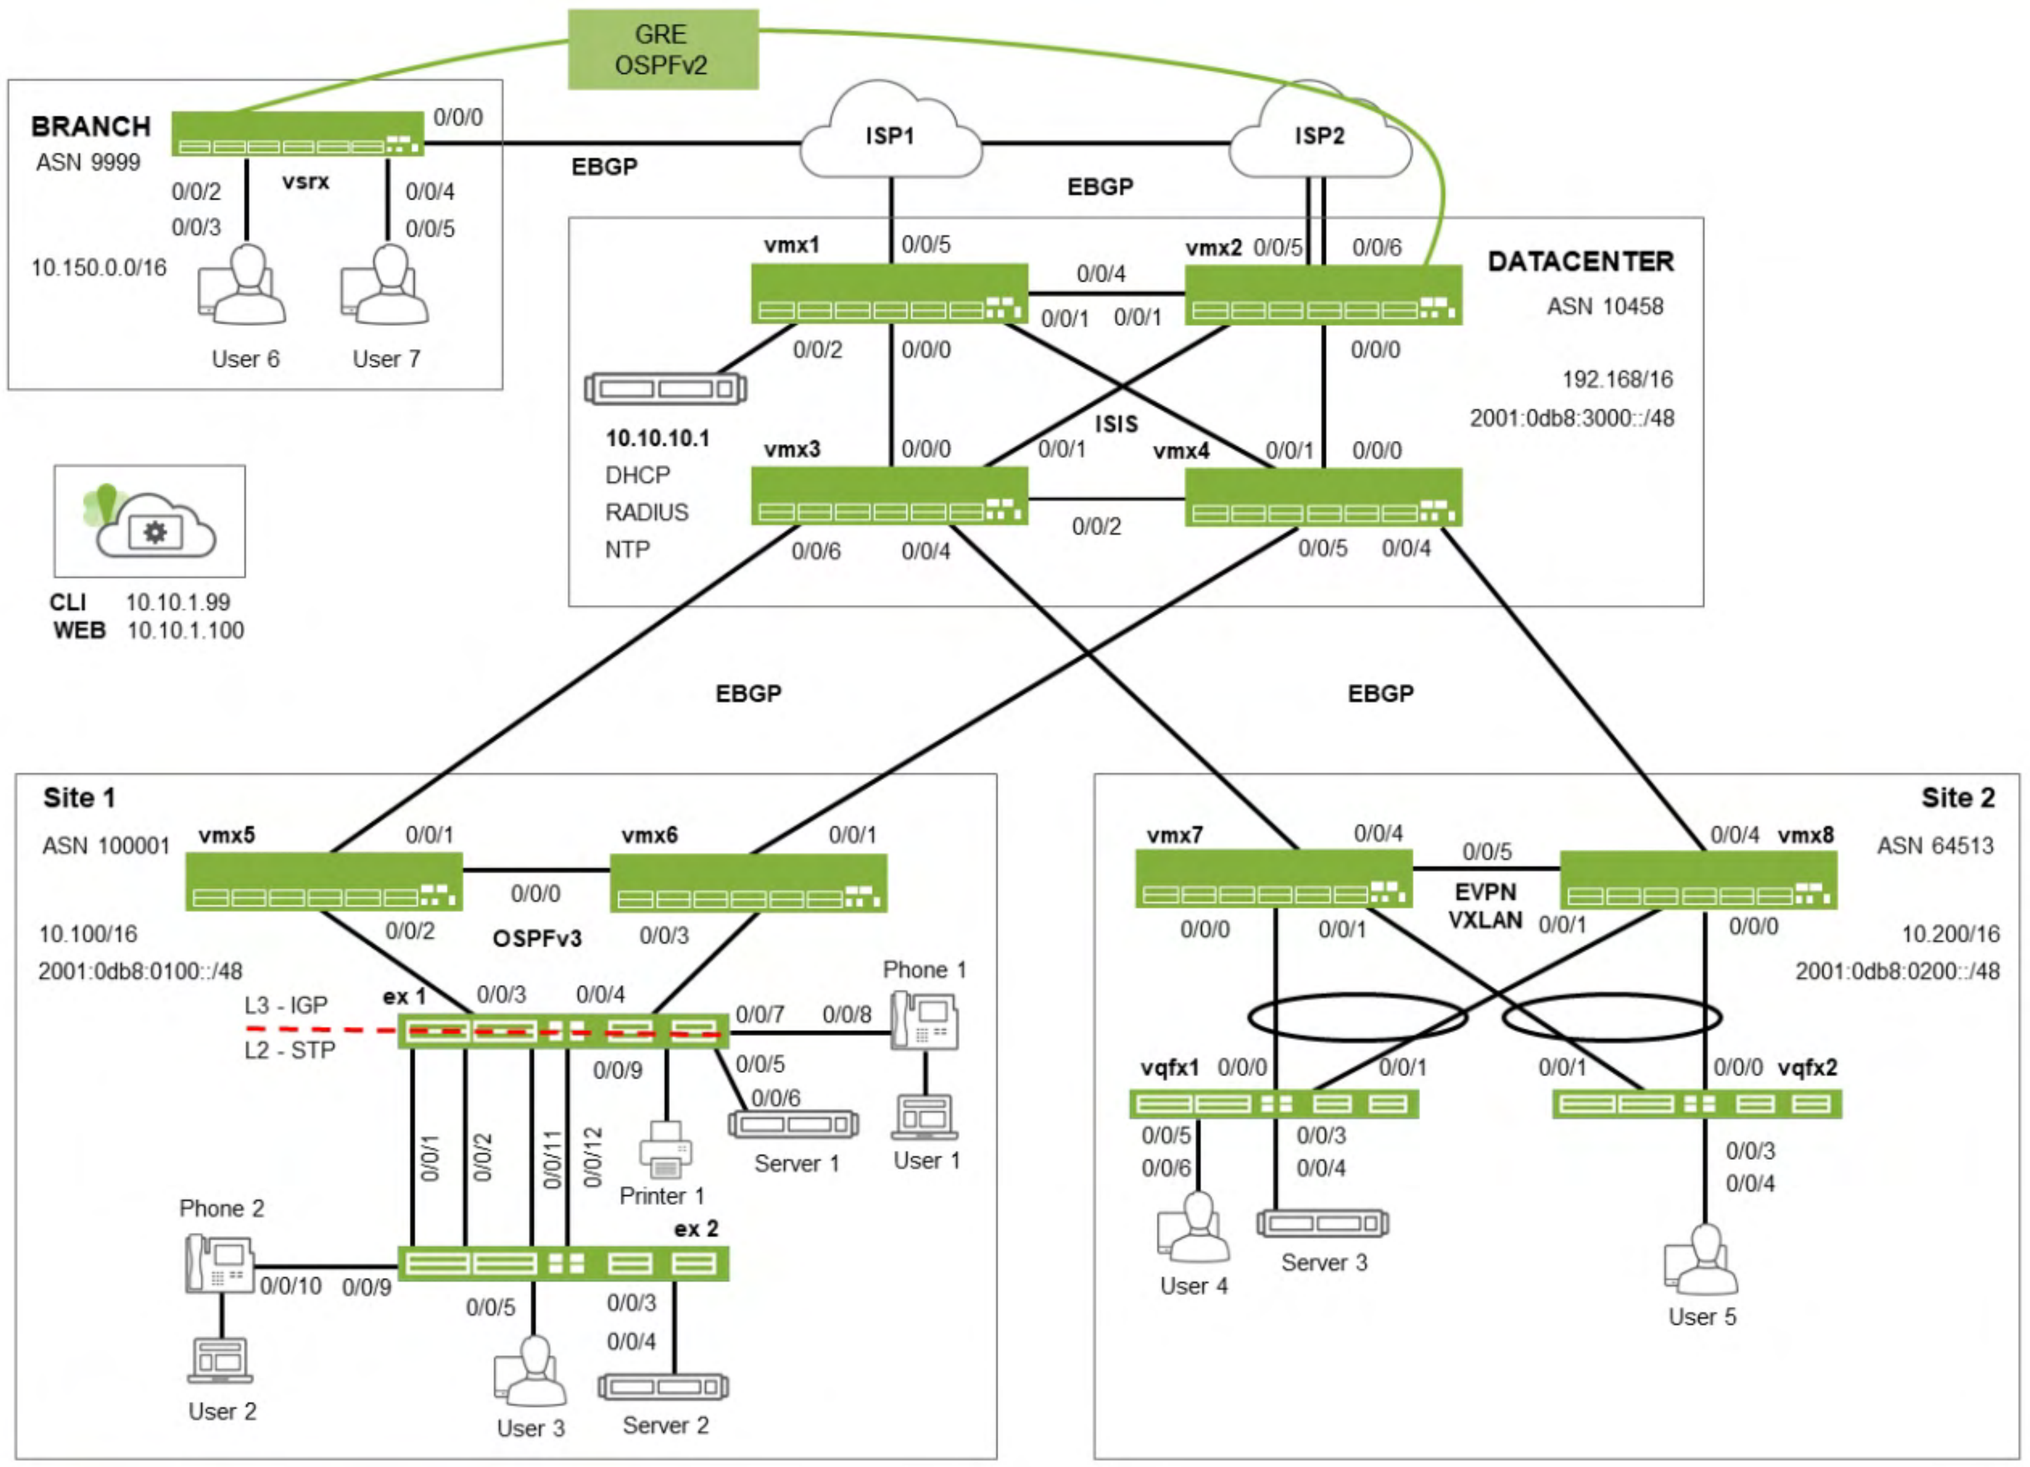

Task 1: IBGP Data Center

1) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have an IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and also accept routes received from each other.
2) Do not configure neighbor statements for IBGP sessions on vMX1 and vMX2.
3) Enable BFD with authentication for IBGP sessions in the data center. Allow multihop sessions. Do not configure BFD between the Route Reflectors.
4) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS and the Branch Office.
5) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down. vMX2 should not advertise the default route if the BGP peering with ISP2 goes down.

© 2020 Juniper Networks, Inc. All rights reserved.

6) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the AS.

Task 2: IBGP Site 1 and Site 2

1) Configure IBGP in Site1 between vMX5 and vMX6 and in Site 2 between vMX7 and vMX8. Use the information as shown in the table above.

Task 3: EBGP Data Center

1) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
2) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.
3) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.
4) Advertise the IPv4 summary routes received from Site 1, Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs
5) Advertise a single /32 route to both ISPs that encompasses the IPv6 network of the data center, Site 1 and Site 2.
6) Ensure that ISP2 is the primary path to all IPv4 address blocks announced by your AS. Both ISPs apply route policy based on received communities. Use the information in the table below to meet the task requirement

ISP	COMMUNITY	ISP ACTION
ISP1	1111:1	Prepend 1x
	1111:2	Prepend 2x
	1111:3	Prepend 3x
ISP2	2222:1	Prepend 1x
	2222:2	Prepend 2x
	2222:3	Prepend 1x

7) Ensure the ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.

© 2020 Juniper Networks, Inc. All rights reserved.

8) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2. Prefixes from Site1 that are advertised to both ISPs should contain the ASN of Site 1.

9) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Site 1 and Site 2.

10) Site 1 and Site 2 should learn the summarized prefixes of the other site, data center, and branch office. Site 1 and Site 2 should also learn the IPv4 and IPv6 default route and the management network 10.10.10/24 that originated in the data center.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 6: EVPN/VXLAN

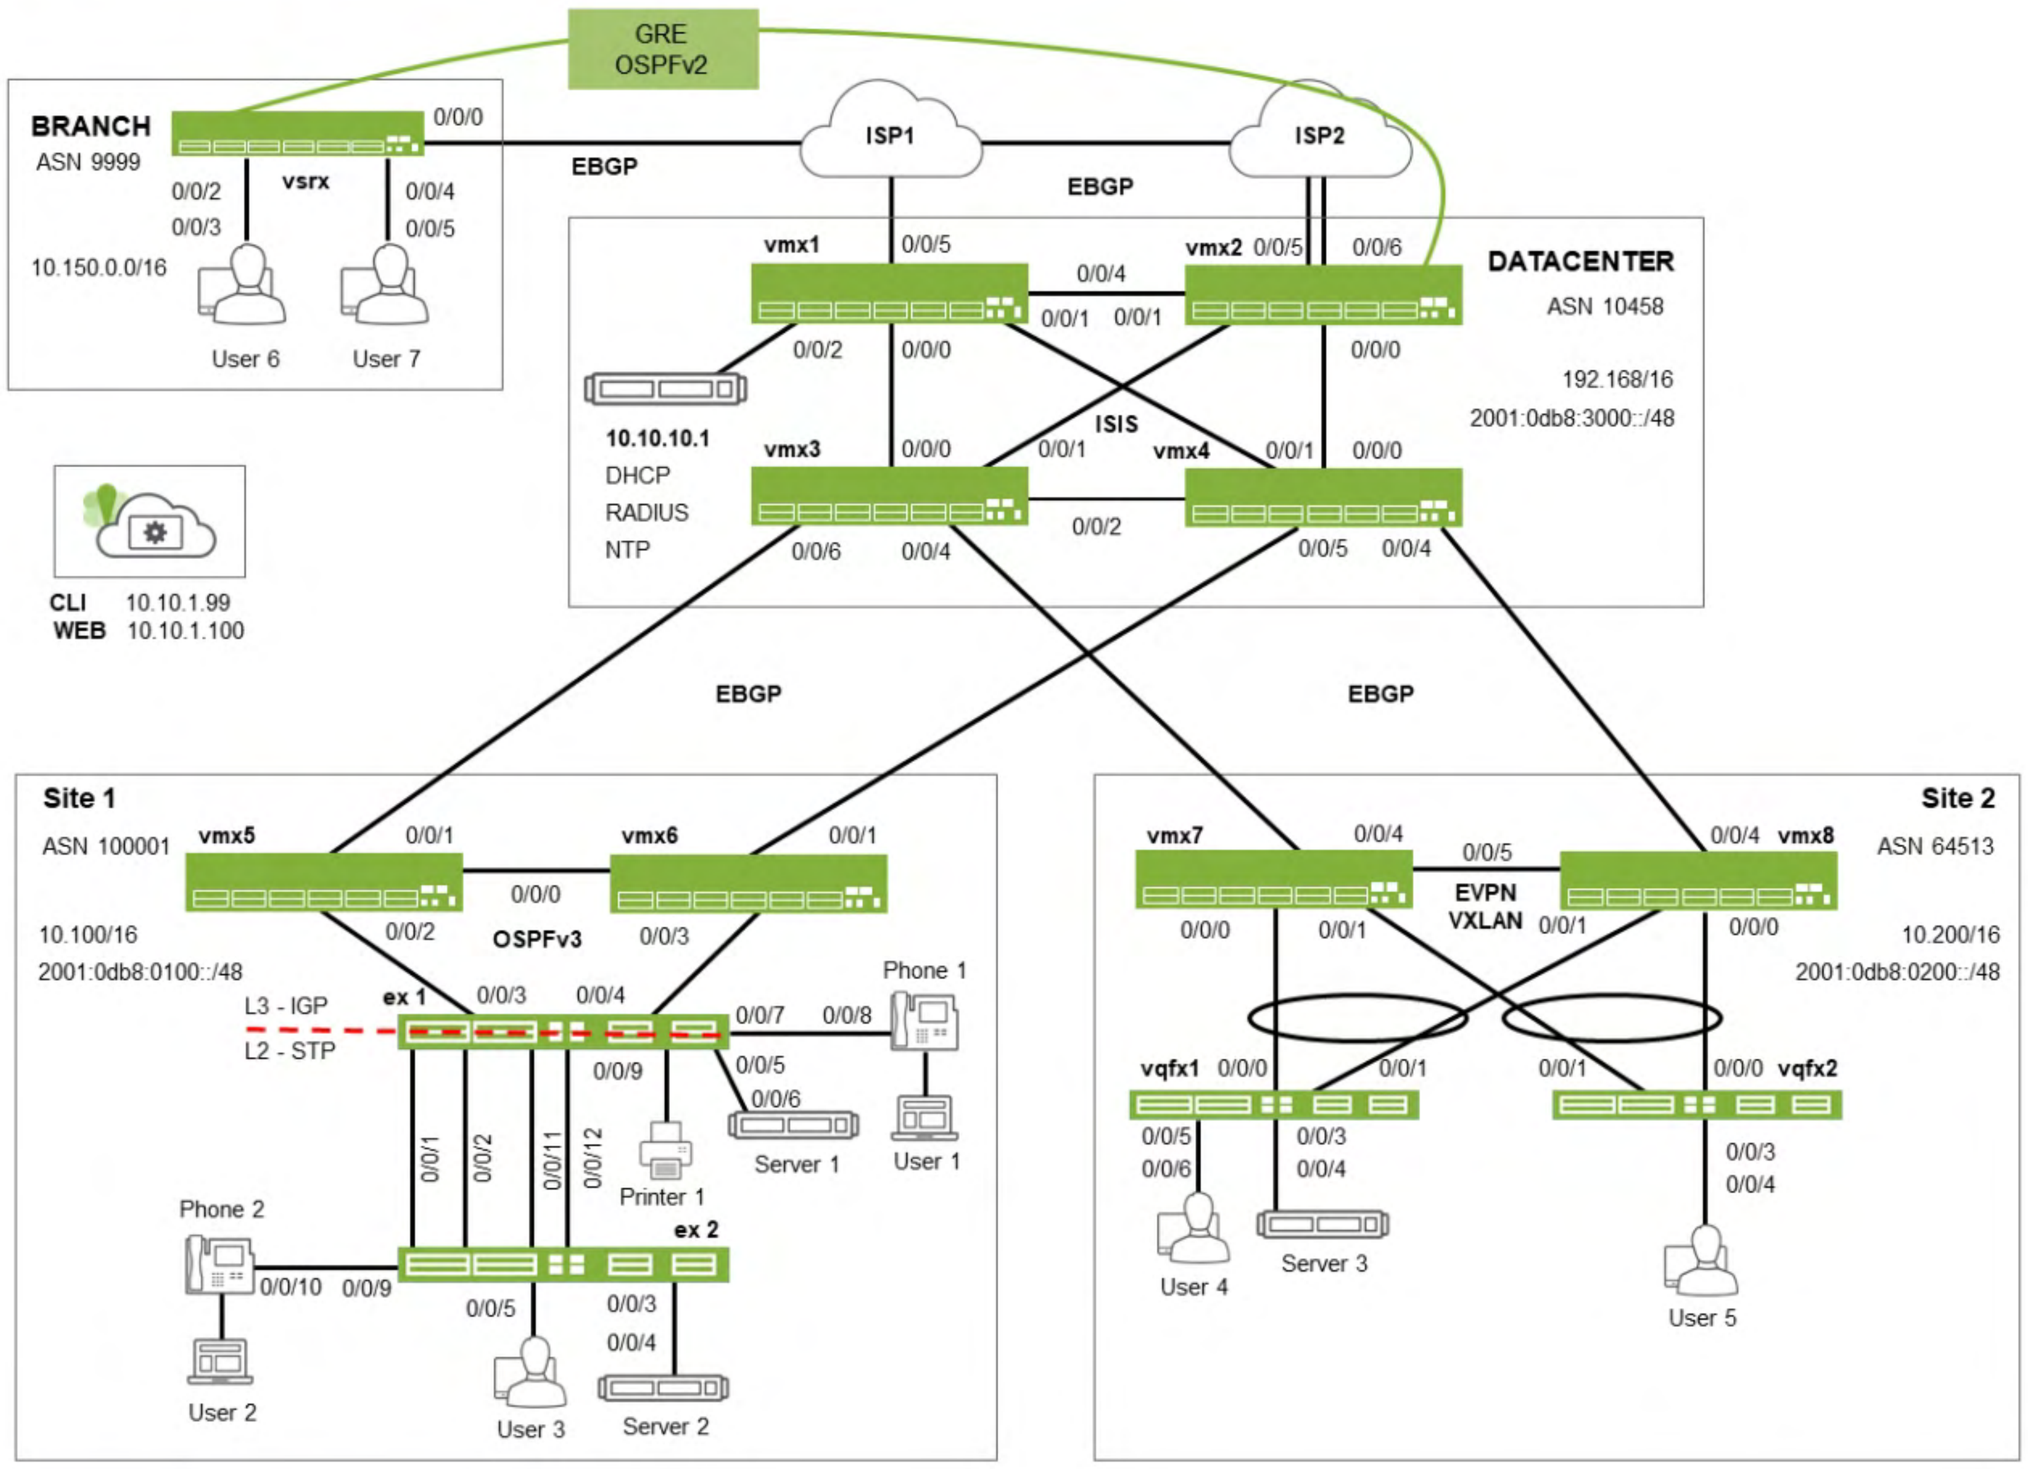

Task 1: EVPN VXLAN

In this task, you will configure the Layer 2 connectivity in Site 2 using the EVPN/VXLAN technology.

1) Configure one EVPN instance on vMX7 and VMX8 that uses VXLAN transport. The EVPN should provide layer 2 connectivity for the VLANs in Site 2. Ensure that active multi-homing is supported for unicast traffic.
2) Ensure that route targets are automatically generated.
3) Ensure that all virtual gateway addresses reply to ICMP packets.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 7: Class of Service

This chapter is focused on Class of Services. You will configure and monitor MF and BA Classifiers, Two-color Policers, different Forwarding classes with respected Queues and Schedulers, DSCP Rewrite markers, RED Drop profiles and Class based forwarding.

Task 1: vMX Class of Service

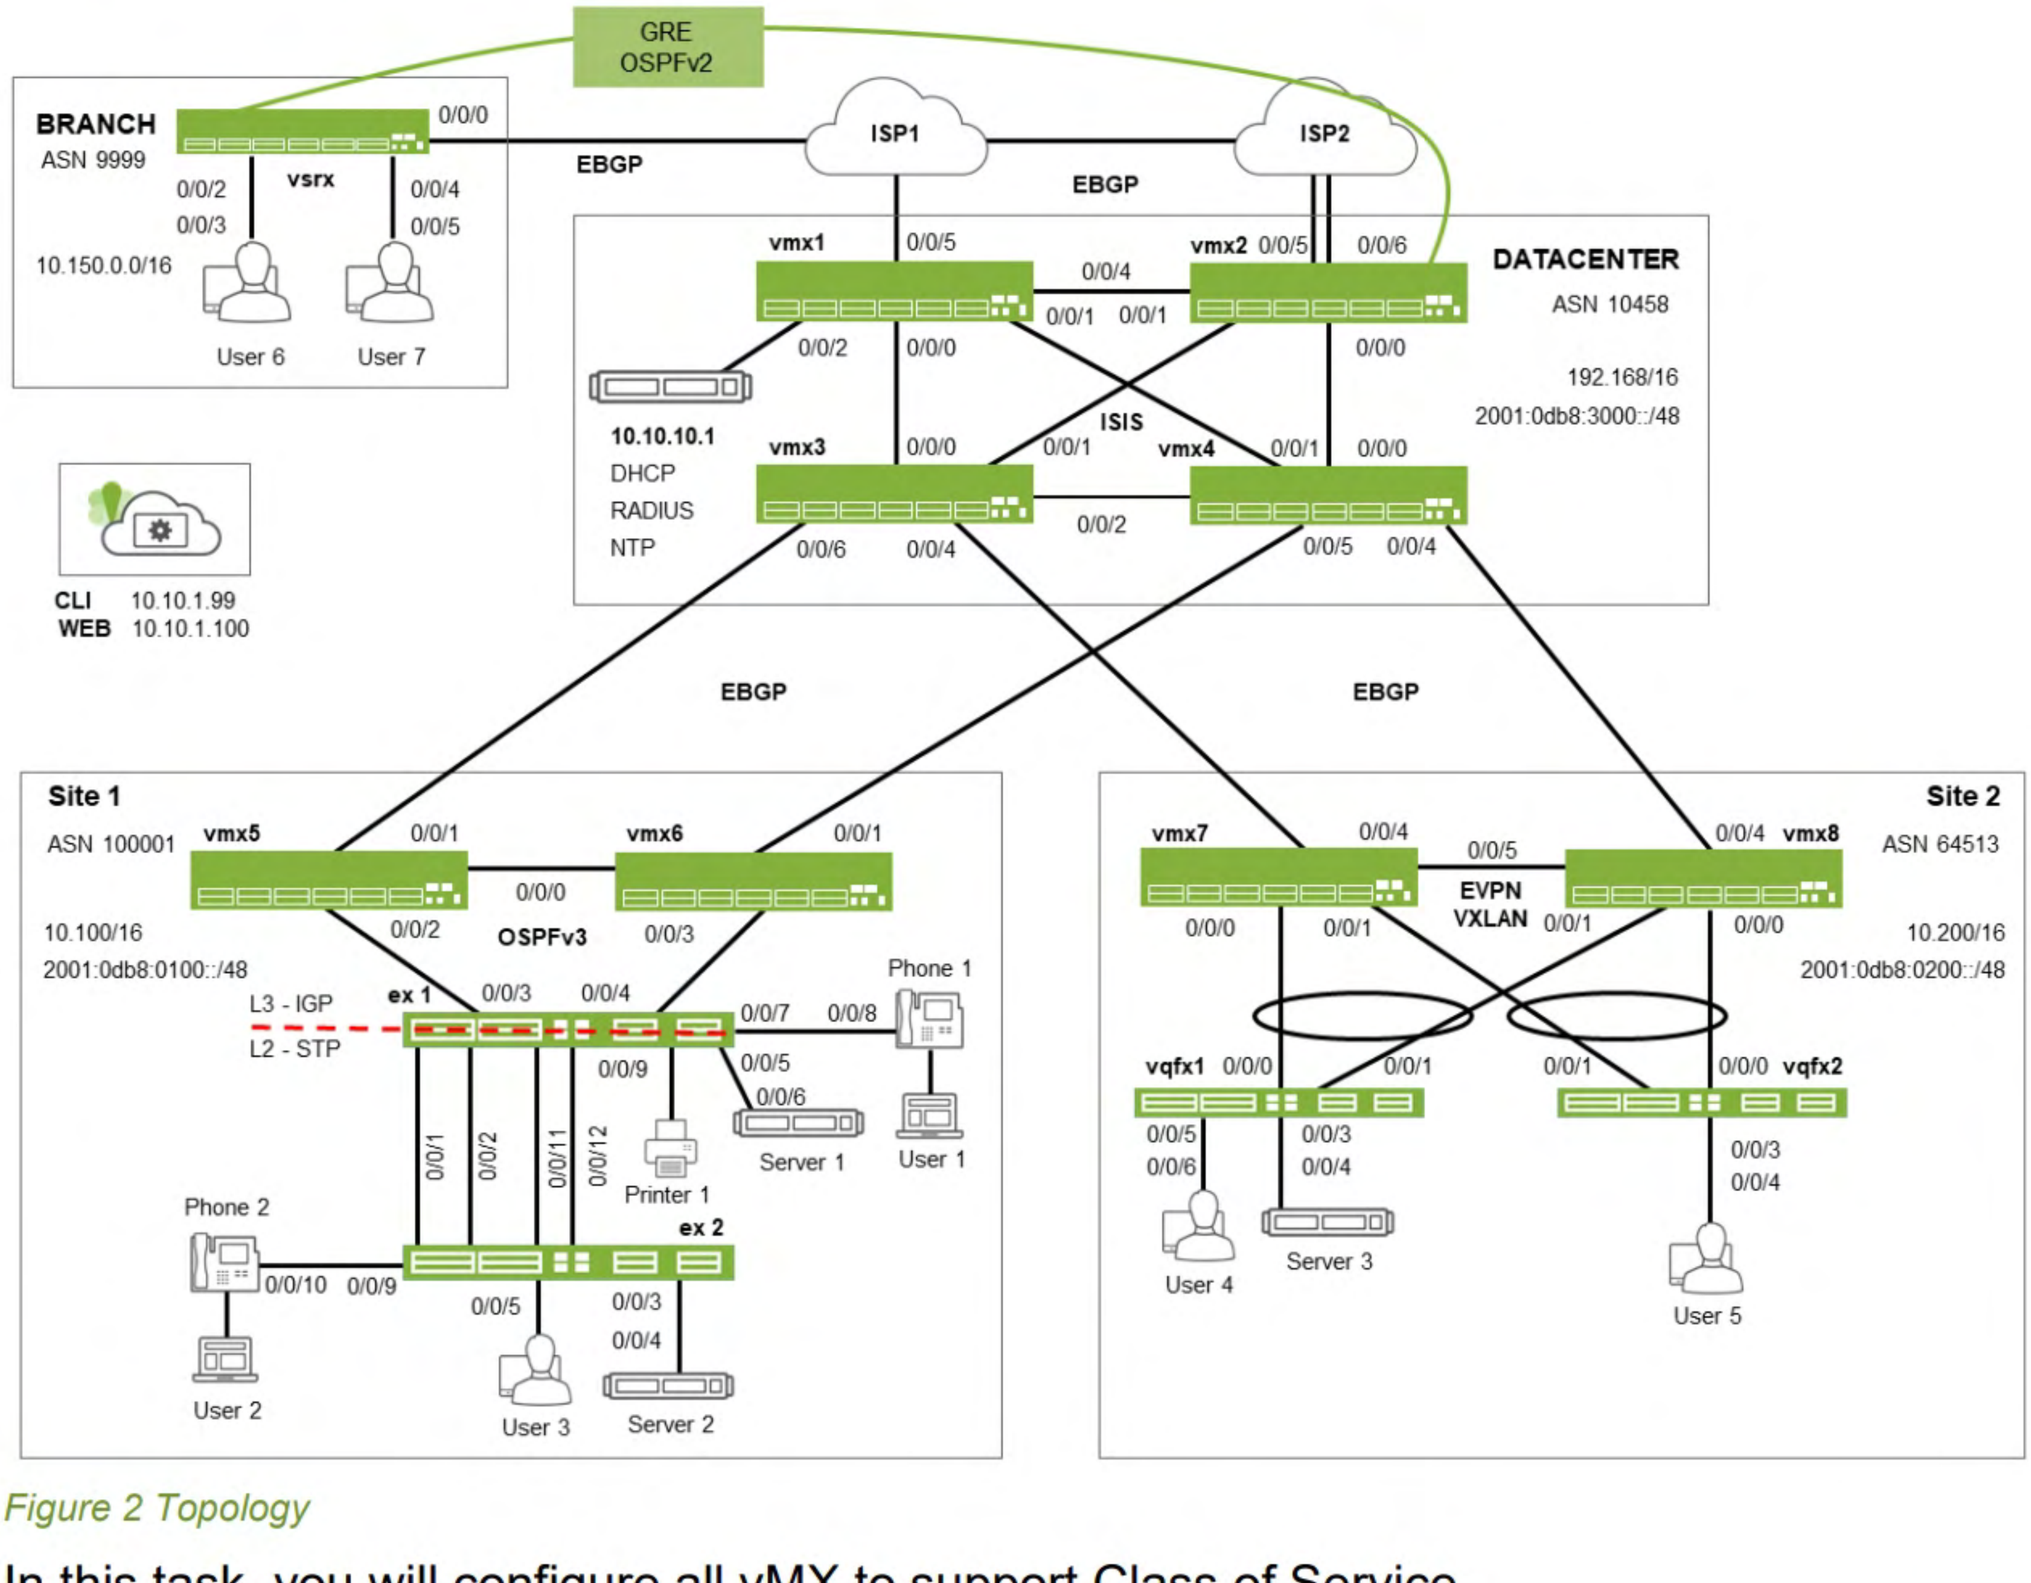

Figure 2 Topology

In this task, you will configure all vMX to support Class of Service.

1) Configure Forwarding Classes and map them to the output Queues as indicated in the table below.

© 2020 Juniper Networks, Inc. All rights reserved.

Forwarding Class	Queue	Scheduler
control	7	control-sched
voice	5	voice-sched
server-gold	3	server-gold-sched
server-silver	2	server-silver-sched
server-bronze	1	server-bronze-sched
best-effort	0	best-effort-sched

2) Configure schedulers with the parameters in Table 7.

Scheduler	Parameter	Value
control-sched	priority	medium-high
	transmit rate	5%
	buffer-size	10%
voice-sched	priority	strict-high
	transmit rate	10%
	buffer-size	20 (milliseconds)
server-gold-sched	priority	high
	transmit rate	25%
	buffer-size	10% exact
server-silver-sched	priority	medium-high
	transmit rate	35%
	buffer-size	50% exact
	drop profile high	relaxed
server-bronze-sched	priority	medium-low
	transmit rate	remainder
	buffer-size	remainder
	drop profile low	relaxed
	drop profile high	aggressive
best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder
	drop profile any	aggressive

Table 7

3) Configure two interpolated Drop profiles listed in Table 8.

© 2020 Juniper Networks, Inc. All rights reserved.

Drop Profile	Fill Level (%)	Drop Probability (%)
relaxed	75	0
	90	75
	100	100
aggressive	50	0
	75	50
	90	75
	100	100

Table 8

4) Apply the schedulers to all physical gigabit ethernet interfaces.
5) Configure IPv4 and IPv6 Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

6) Apply the classifiers to all (logical) interfaces.
7) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	high
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low

© 2020 Juniper Networks, Inc. All rights reserved.

best-effort	001110	high
	000000	high

8) Apply the rewrite rules on all logical interfaces.

Task 2: EX Class of Service

In this task, you will configure EX1 and EX2 in Site1 to support Class of Service.

1) Configure the following Forwarding Classes and map them to the output Queues as indicated in Table 9.

Forwarding Class	Queue	Schedulers L2 and L3 core Interfaces	Schedulers L2 access ports
control	7	control-sched-core	control-sched-user
voice	5	voice-sched-core	voice-sched-user
server-gold	3	server-gold-sched	server-gold-sched
server-silver	2	server-silver-sched	server-silver-sched
server-bronze	1	server-bronze-sched	server-bronze-sched
best-effort	0	best-effort-sched	best-effort-sched

Table 9

2) Configure schedulers with the parameters as described in Table 10 below.

Scheduler	Parameter	Value
control-sched-core	priority	strict-high
	shaping rate	5%
	buffer-size	10%
control-sched-user	priority	strict-high
	shaping rate	2%
	buffer-size	10%
voice-sched-core	priority	strict-high
	shaping rate	10%
	buffer-size	5% exact
voice-sched-user	priority	strict-high
	shaping rate	5%
	buffer-size	5% exact
server-gold-sched	priority	strict-high
	shaping rate	25%

© 2020 Juniper Networks, Inc. All rights reserved.

server-silver-sched	buffer-size	10% exact
	priority	low
	transmit rate	35%
server-bronze-sched	buffer-size	50%
	priority	low
	transmit rate	remainder
best-effort-sched	buffer-size	remainder
	priority	low
	transmit rate	remainder
	buffer-size	remainder

Table 10

3) Apply the schedulers to the physical interfaces on EX1 and EX2. Make sure that OoB interfaces are not affected.
4) Configure IPv4, IPv6 and dot1p Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	Dot1p	Dot1p Loss-priority	DSCP value	DSCP Loss-priority
control	111	low	110000	low
			111000	low
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 11

5) Apply the classifiers to all interfaces as described in Table 124. Make sure that the OoB interfaces are not affected.

Device	Classifiers	Interfaces
EX1	BA IPv4, IPv6 and dot1p classifiers	All interfaces
EX2	Dot1p	All interfaces

Table 12

© 2020 Juniper Networks, Inc. All rights reserved.

6) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	Dot1p	Dot1p loss priority	DSCP value	Loss-priority
control	111	low	110000	low
			111000	high
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 13

7) Apply the rewrite-rules to all interfaces as described in the table below. Make sure that the OoB interfaces are not affected.

Device	Rewrite-rules	Interfaces
EX1	IPv4, IPv6 and dot1p rewrite-rules	All interfaces
EX2	Dot1p rewrite rules	All interfaces

Table 14

8) Configure all voice enabled interfaces on EX1 and EX2 to advertise the Forwarding-Class ‘voice’ to the VoIP phones.

© 2020 Juniper Networks, Inc. All rights reserved.

Super Lab 1

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

Lab Setup Information

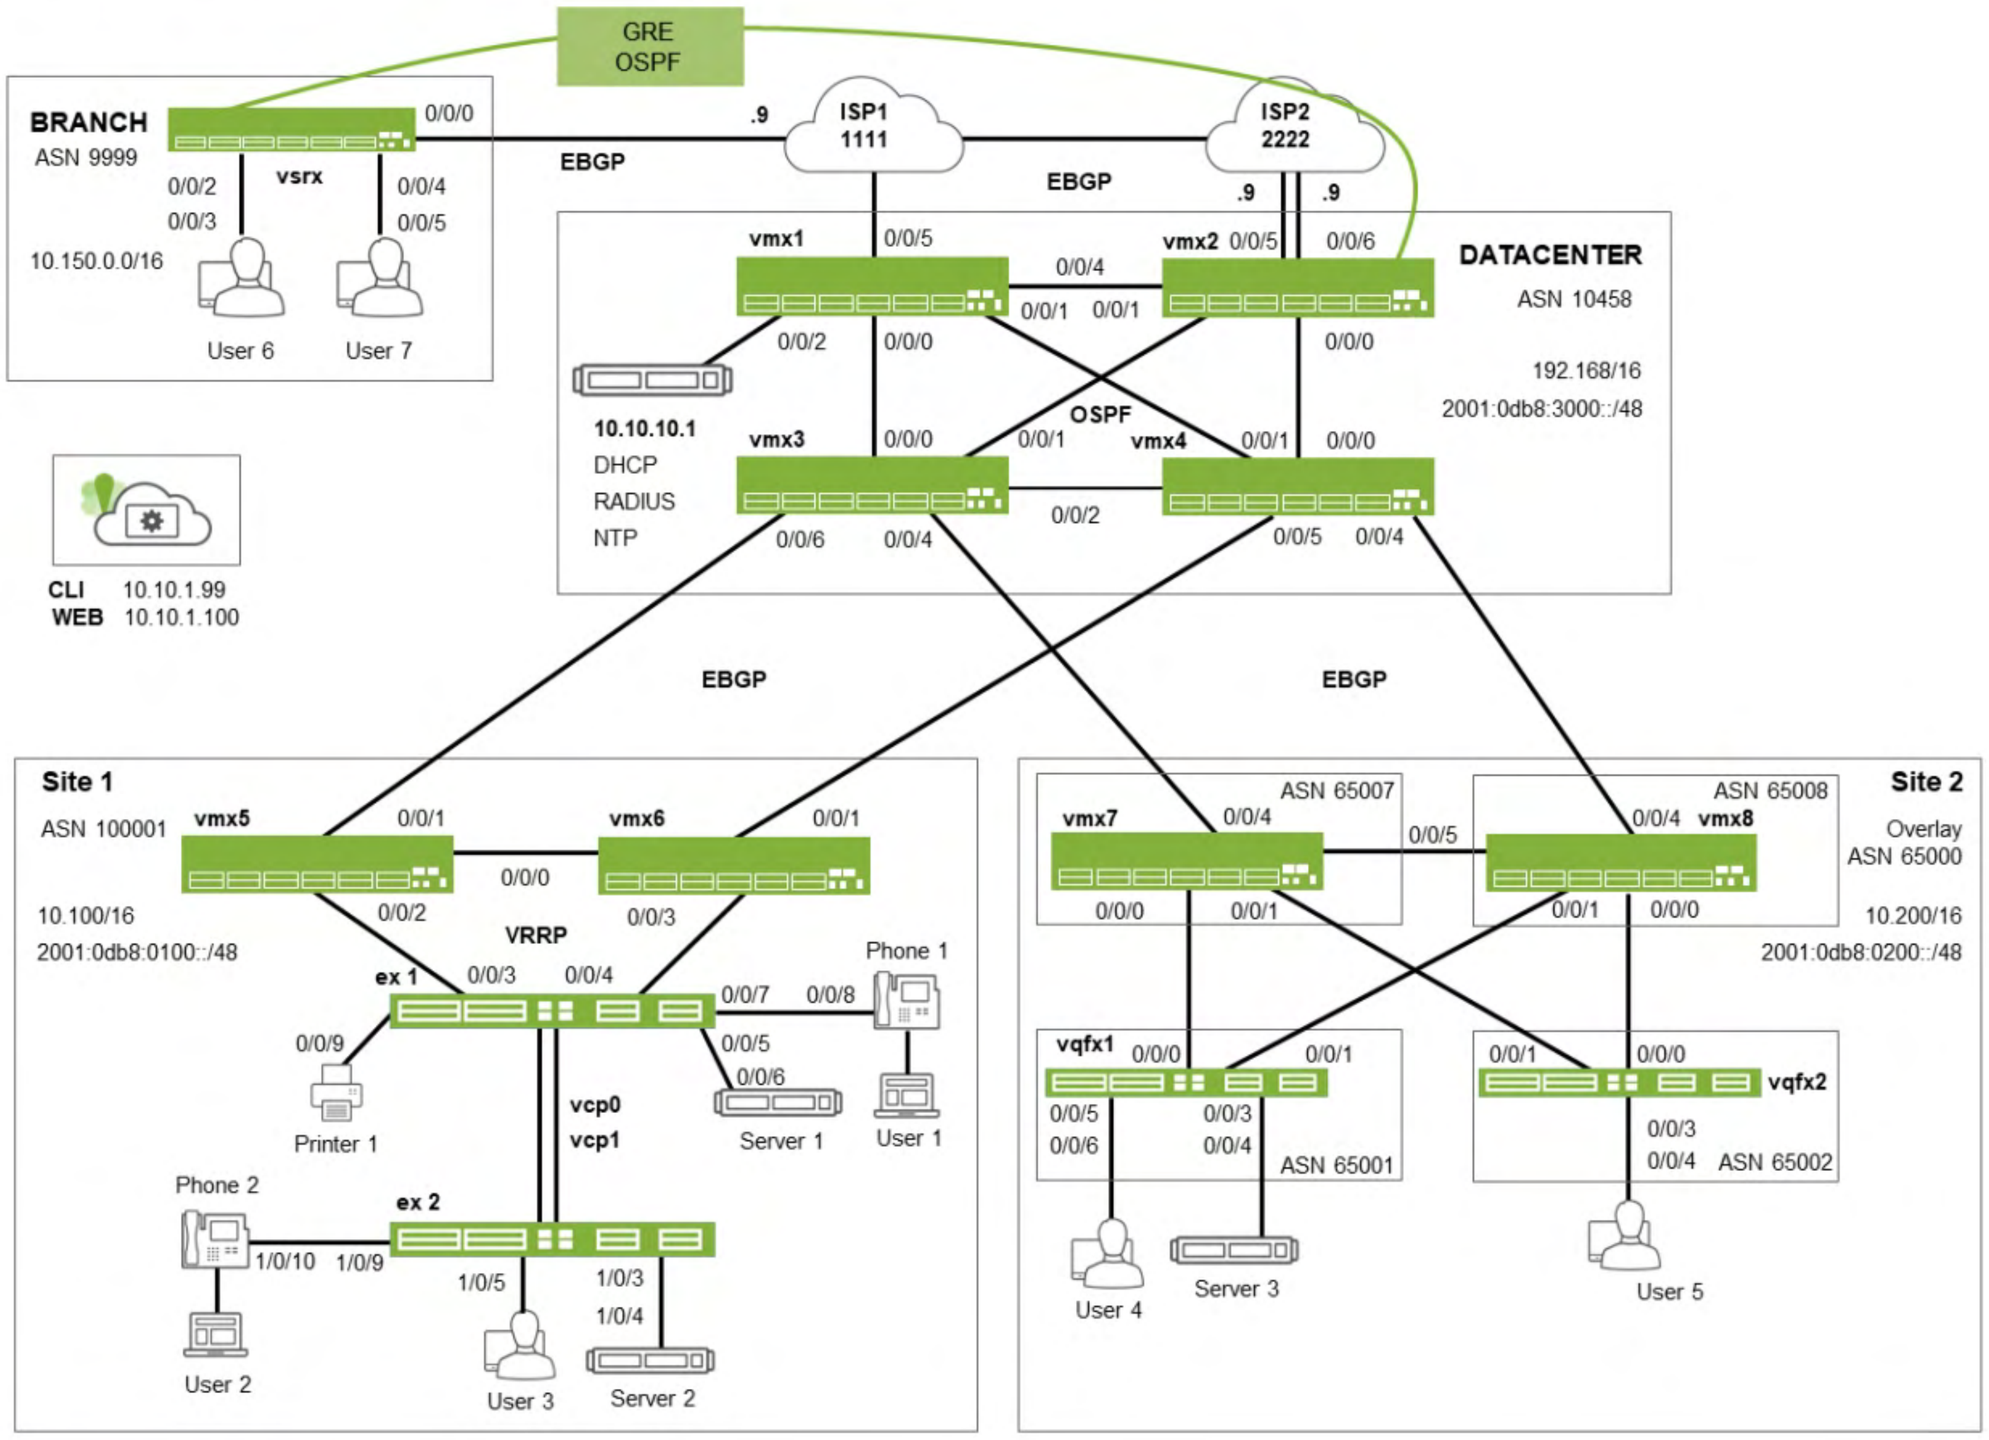

Figure 3 Topology

In the figure above, the topology is shown that is used for this super lab. There are four vMX devices in the data center. vMX1 and vMX2 are both connected to an ISP. vMX3 and vMX4 are connected to vMX5 and vMX6 in Site1. Site 1 contains two EX4300 devices (EX1

© 2020 Juniper Networks, Inc. All rights reserved.

and EX2) that need to be configured as a virtual chassis. The virtual chassis has several devices attached, such as phones, printers, servers and users. Please note that some of these devices are simulated using routing-instances on the EX4300 switches. For example, server2 on EX2 has a layer3 port (ge-1/0/4) that participates in a routing-instance. There is a physical ethernet connection between port ge-1/0/4 and ge-1/0/3, where port ge-1/0/3 participates are a Layer 2 port in the inet.0 routing table. Printer1 and User2 are virtual machines that are accessible (see lab information). The data center vMX3 and vMX4 is also connected to Site2 vMX7 and vMX8. vMX7 and vMX8 have Layer 2 connections to vQFX1 and vQFX2. Like the EX4300 switches in Site1, the vQFX switches also have users and a server attached. There is also a branch office with two users that is connected to ISP1.

Device	*OoB Interface*	*OoB IPv4 address*
vMX1	fxp0	10.10.1.1/24
vMX2	fxp0	10.10.1.2/24
vMX3	fxp0	10.10.1.3/24
vMX4	fxp0	10.10.1.4/24
vMX5	fxp0	10.10.1.5/24
vMX6	fxp0	10.10.1.6/24
vMX7	fxp0	10.10.1.7/24
vMX8	fxp0	10.10.1.8/24
VR	fxp0	10.10.1.9/24
vSRX	fxp0	10.10.1.10/24
EX1	ge-0/0/0	10.10.1.11/24
EX2	ge-0/0/0	10.10.1.12/24
vQFX1	em0	10.10.1.13/24
vQFX2	em0	10.10.1.14/24

© 2020 Juniper Networks, Inc. All rights reserved.

Junos Space	N/A	10.10.1.99/24 (CLI) 10.10.1.100/24 (WEB)

Table 2 Management Network (out of band)

The table above shows the devices used in the topology and their management IP addresses. These addresses are preconfigured and should not be changed and are listed here for informational purposes only.

Device	Interface	Ipv4 address	Ipv6 address
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128
	ge-0/0/0	192.168.13.1/24	2001:0db8:3000:0013::1/64
	ge-0/0/1	192.168.14.1/24	2001:0db8:3000:0014::1/64
	ge-0/0/2	10.10.10.100/24	N/A
	ge-0/0/4	192.168.12.1/24	2001:0db8:3000:0012::1/64
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128
	ge-0/0/0	192.168.24.2/24	2001:0db8:3000:0024::2/64
	ge-0/0/1	192.168.23.2/24	2001:0db8:3000:0023::2/64
	ge-0/0/4	192.168.12.2/24	2001:0db8:3000:0012::2/64
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64
	ge-0/0/6	172.16.3.2/24	2001:0dba:4000:0003::2/64
	GRE	173.17.99.2/24	
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128
	ge-0/0/0	192.168.13.3/24	2001:0db8:3000:0013::3/64
	ge-0/0/1	192.168.23.3/24	2001:0db8:3000:0023::3/64
	ge-0/0/2	192.168.34.3/24	2001:0db8:3000:0034::3/64
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128
	ge-0/0/0	192.168.24.4/24	2001:0db8:3000:0024::4/64
	ge-0/0/1	192.168.14.4/24	2001:0db8:3000:0014::4/64
	ge-0/0/2	192.168.34.4/24	2001:0db8:3000:0034::4/64
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128
	ge-0/0/0	10.100.56.5/24	2001:0db8:0100:0056::5/64
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64
	ge-0/0/2	10.100.15.5/24	2001:0db8:0100:015::5/64
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/0	10.100.56.6/24	2001:0db8:0100:0056::6/64
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64
	ge-0/0/3	10.100.116.6/24	2001:0db8:0100:0016::6/64
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128
	ge-0/0/0	10.200.17.7/24	N/A
	ge-0/0/1	10.200.27.7/24	N/A
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64
	ge-0/0/5	192.168.78.7/24	2001:0db8:3000:0078::7/64
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128
	ge-0/0/0	10.200.28.8/24	N/A
	ge-0/0/1	10.200.18.8/24	N/A
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64
	ge-0/0/5	192.168.78.8/24	2001:0db8:3000:0078::8/64
vSRX	lo0	10.150.0.99/32	
	ge-0/0/0	10.1.1.2/30	
	ge-0/0/2	Layer 2	
	ge-0/0/3	10.150.99.4/24	
	ge-0/0/4	Layer 2	
	ge-0/0/5	10.150.99.5/24	
	GRE	173.17.99.1/24	
EX1	ge-0/0/3	10.100.15.1/24	2001:0db8:0100:015::1/64
	ge-0/0/4	10.100.16.1/24	2001:0db8:0100:016::1/64
	ge-0/0/5	Layer 2	
	ge-0/0/6	dhcp	2001:0db8:0100:101::1/64
	ge-0/0/7	Layer 2	
	ge-0/0/8	10.100.201.1/24	2001:0db8:0100:201::1/64
	irb.100	10.100.100.254/24	2001:0db8:0100:100::254/64
	irb.101	10.100.101.254/24	2001:0db8:0100:101::254/64
	irb.156	10.100.156.1/24	2001:0db8:0100:0156::1
	irb.200	10.100.200.254/24	2001:0db8:0100:200::254/64
	irb.201	10.100.201.254/24	2001:0db8:0100:201::254/64
EX2	ge-1/0/3	Layer 2	
	ge-1/0/4	dhcp	
	ge-1/0/5	Layer 2	
	ge-1/0/6	Layer 2	
	ge-1/0/9	Layer 2	
	ge-1/0/10	10.100.201.2/24	2001:0db8:0100:201::2/64
vQFX1	xe-0/0/0	10.200.17.1/24	

© 2020 Juniper Networks, Inc. All rights reserved.

	xe-0/0/1	10.200.18.1/24	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.101.1/24	2001:0db8:0200:101::1/64
	xe-0/0/5	Layer 2	
	xe-0/0/6	10.200.100.1/24	2001:0db8:0200:100::1/64
vQFX2	xe-0/0/0	10.200.28.2/24	
	xe-0/0/1	10.200.27.2/24	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.100.3/24	2001:0db8:0200:100::3/64

The table above shows the IP addressing scheme that is used throughout this super lab.

Task 1: General System Features

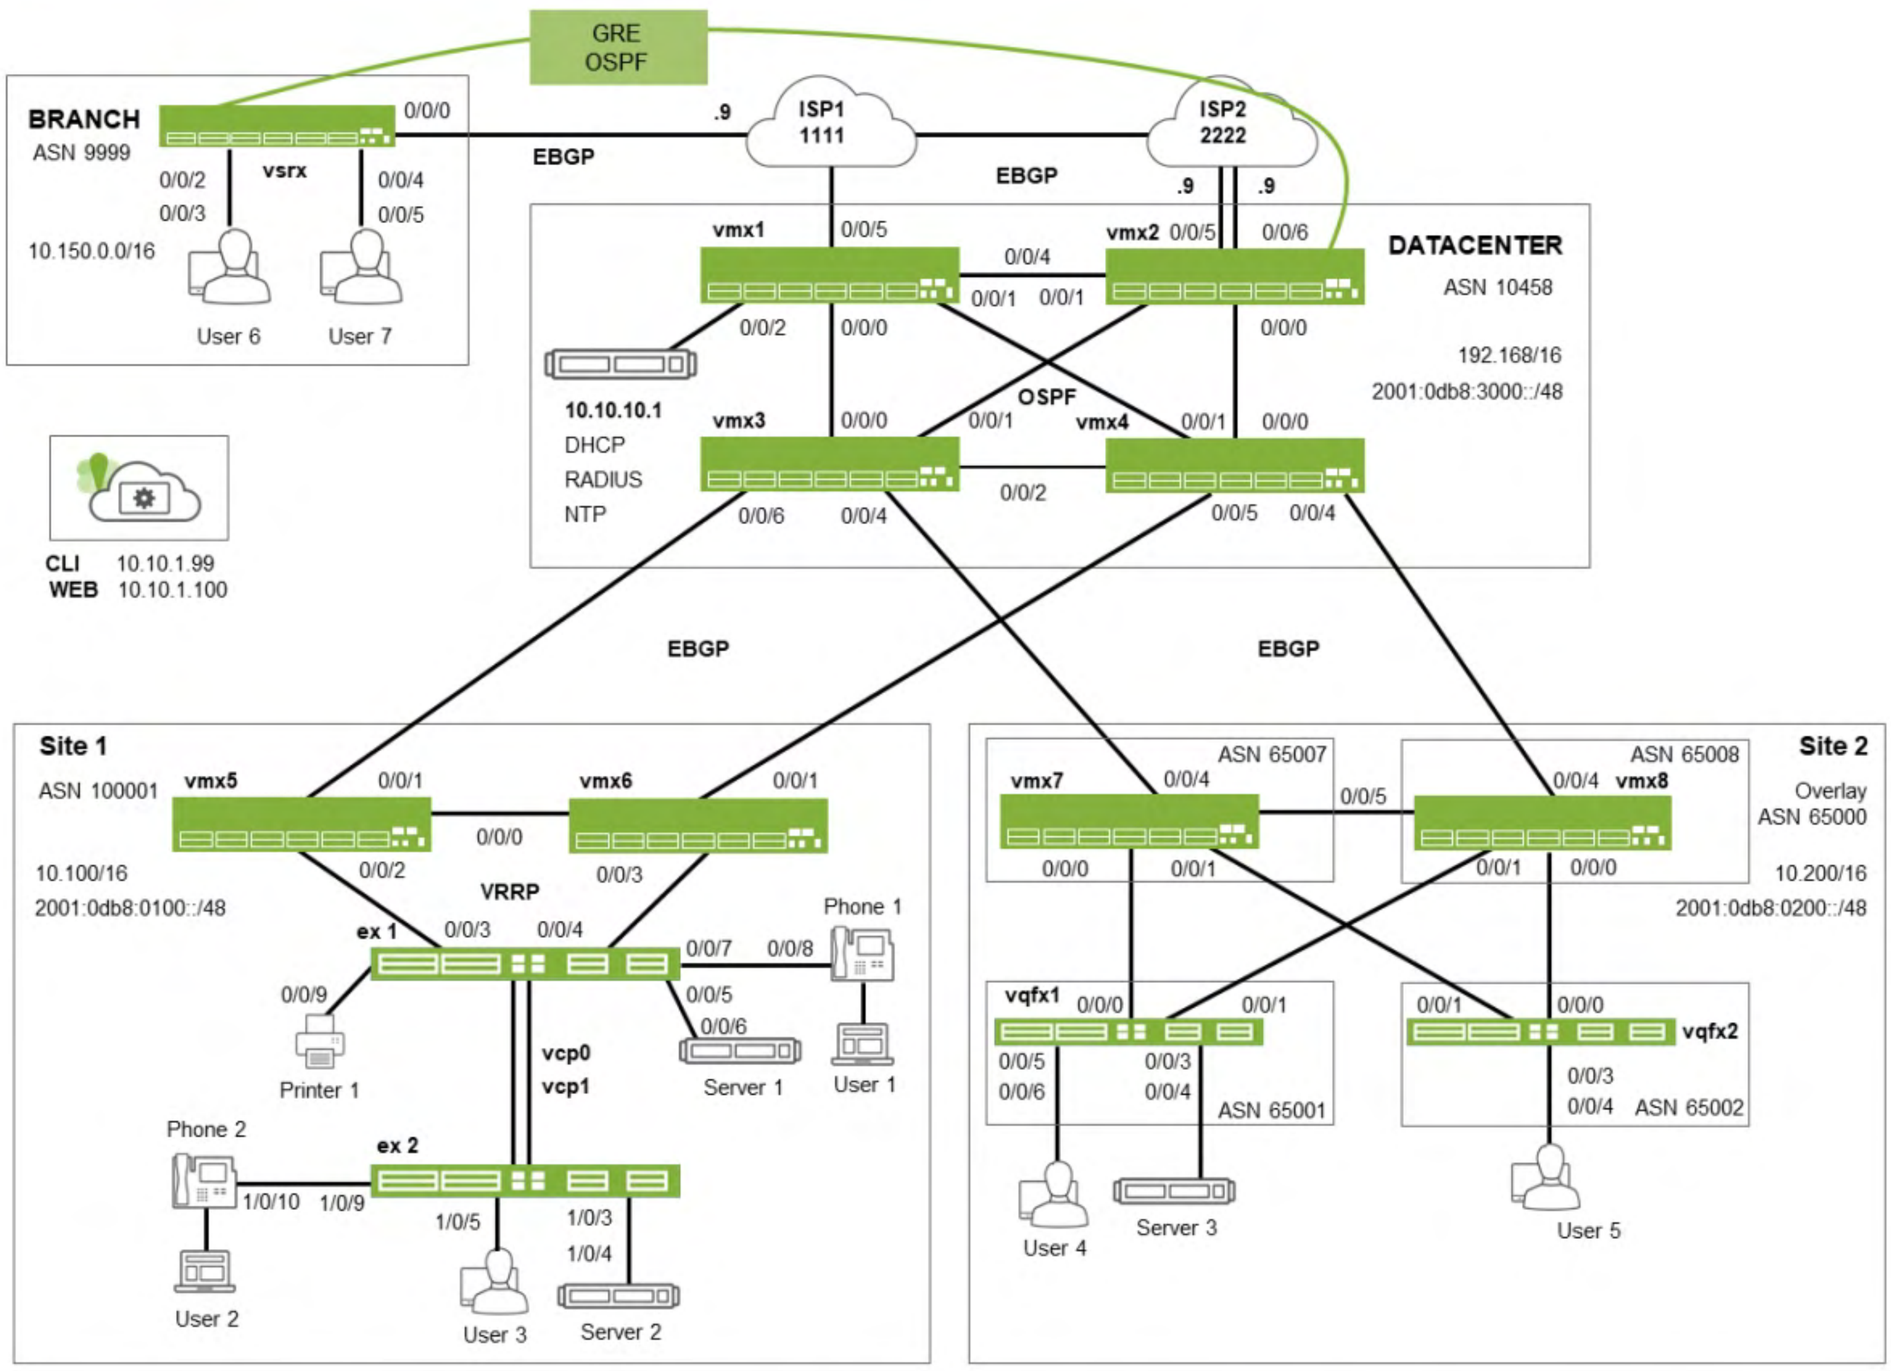

Note: For this task you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure all vMX and EX devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute.
2) Configure server 10.10.10.1 as the DNS server.
3) Set the time zone to America/Los_Angeles on all vMX devices and EX1.
4) Make sure that all vMX devices and EX1 synchronize the time with NTP server 10.10.10.1.

NOTE: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1*

5) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts 3 and timeout of 3 seconds. Use juniper123 as the RADIUS shared secret.
6) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices, EX1 as a super-user.
7) The SNMP read-only community string is jncie. Make sure that only the Junos Space server can poll the devices in the data center, Site1, Site2 and the branch device.
8) Configure all devices to send SNMP traps to Junos Space. in case of link and/or chassis events. Make sure that traps are only send using the OoB management interface.
9) You have a Junos Space--Network Director instance running for device management. You must discover all vMX devices using Junos Space Network Director with SNMPv2 and the community string jncie. Ensure that the devices are present under the Inventory window.

Note: Junos Space can be accessed by browsing to the management address using HTTPS.

10) Back up all vMX configurations onto Junos Space Network Director.
11) Ensure that RFC 4741 compliant Python scripts can connect to vMX1 over secured shell session using Netconf.
12) Allow 6 concurrent connections and a maximum of 2 sessions per minute.
13) Prevent unsupported Junos configuration hierarchies from being exported in get-config RPC outputs.
14) Configure an IPv4 firewall filter which allows protocol messages from the BFD, VRRP, OSPF, BGP and GRE. Make sure that for the protocols VRRP and OSPF messages from directly connected subnets are accepted.

© 2020 Juniper Networks, Inc. All rights reserved.

15) Future added BGP peers should be automatically allowed in the firewall filter.
16) Configure the firewall filter to accept RADIUS, DNS, SNMP management protocols only from the 10.10.10/24 network.
17) Accept NTP messages only from the 10.10.10/24 and 10.10.1/24 networks.
18) Accept SSH and Telnet protocols only from the 10.10.1/24 management network.
19) Configure the devices to accept ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer, which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
20) Protect the devices from TCP syn/rst/fin attacks. Use a policer with the following parameters: bandwidth 500kbps and burst size 125kB. The excess traffic must be dropped.
21) Allow traffic to the control plane that is required for other tasks.
22) Configure the firewall filter to discard any other traffic and count the packets.
23) Apply the firewall filter such as to ensure that it is used for the RE protection.

Task 2: Layer 2 Configuration and Features

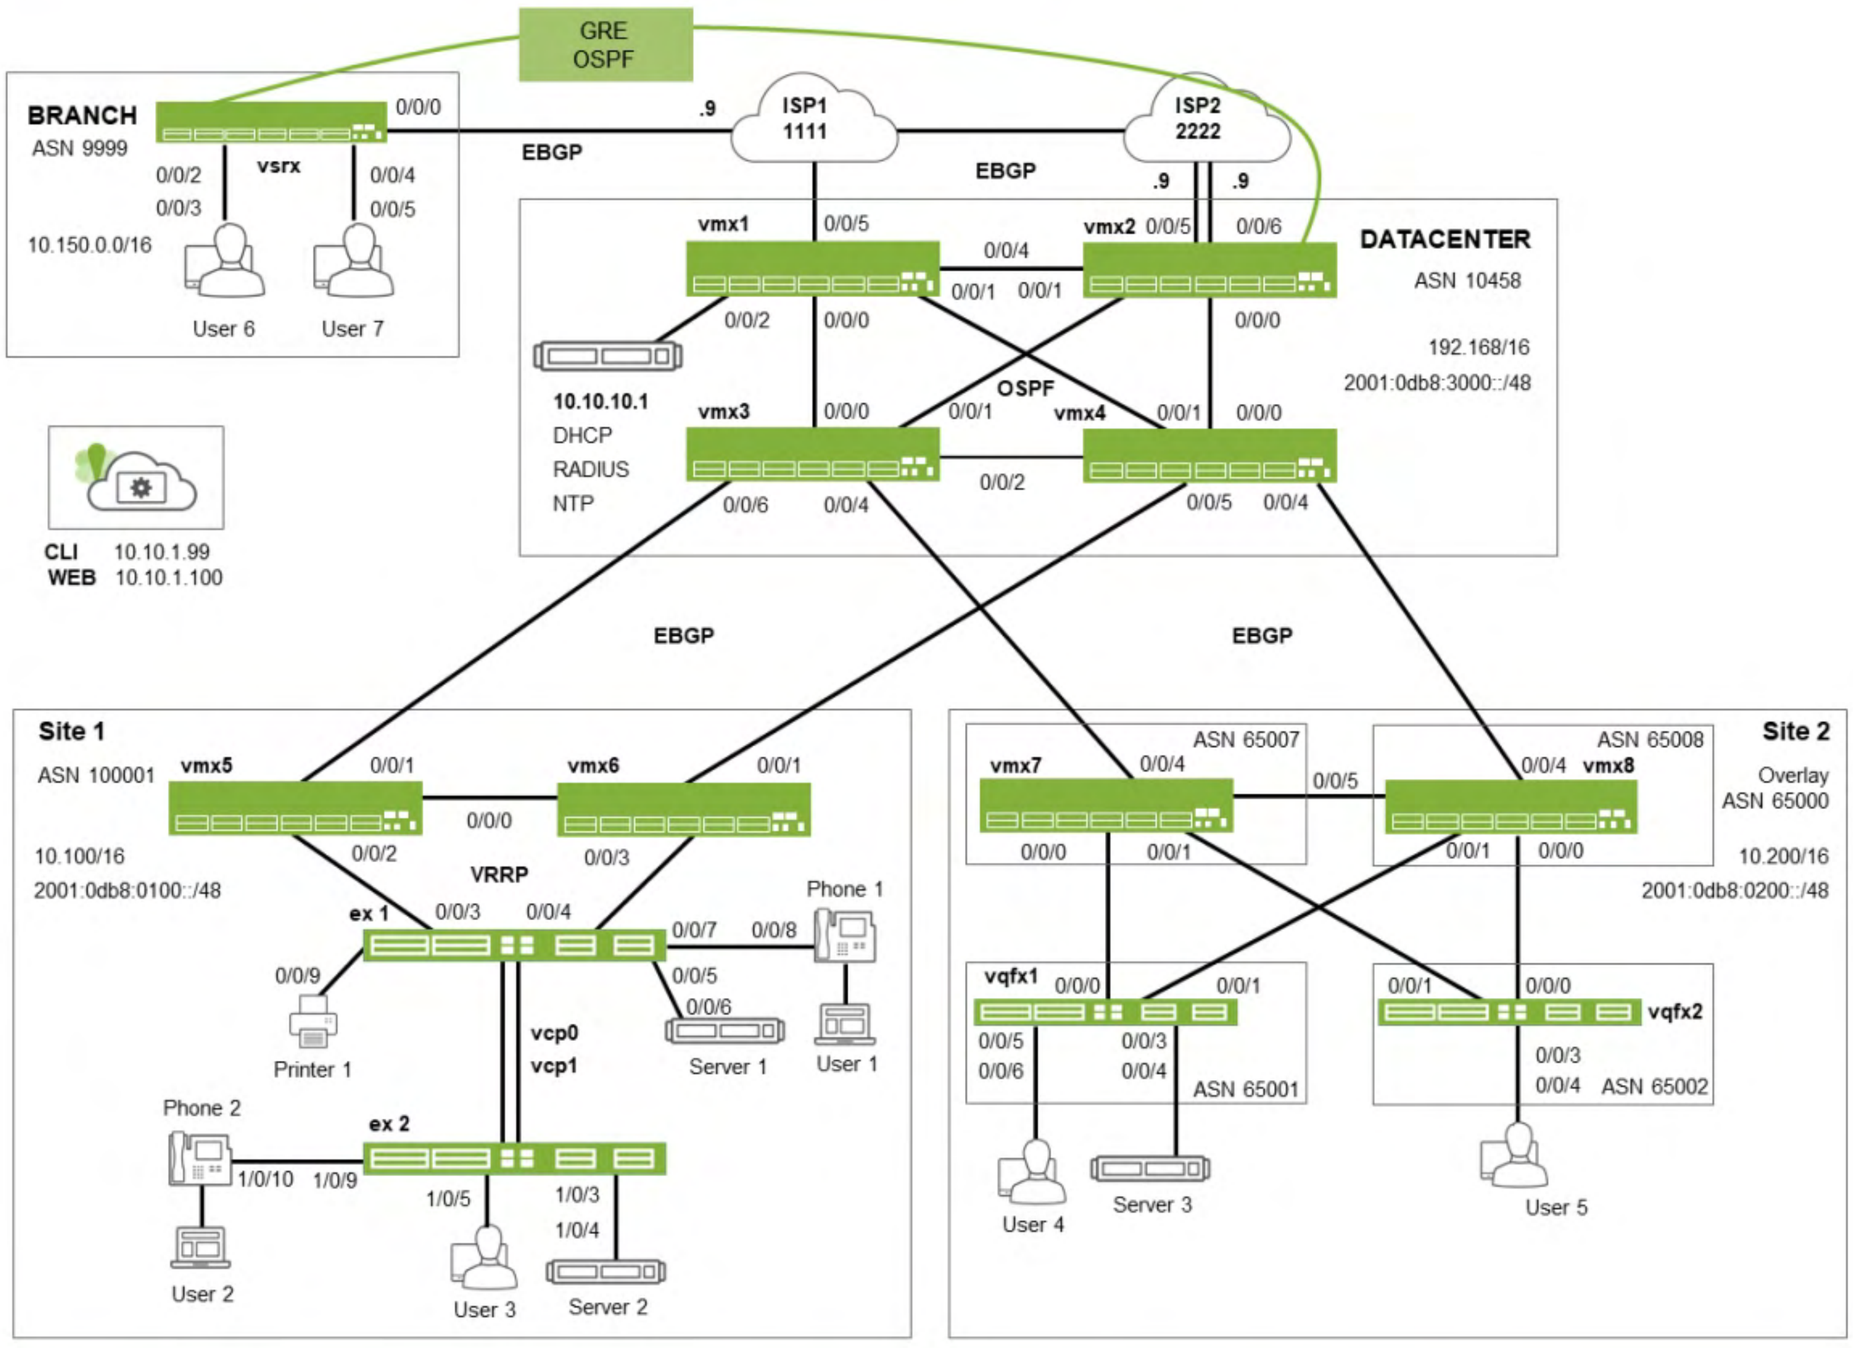

© 2020 Juniper Networks, Inc. All rights reserved.

In this task, you will configure a Virtual chassis consisting of EX1 and EX2. You will configure (VLAN) interfaces on the switches.

1) Configure EX1 and EX2 as a virtual chassis. Ensure EX1 is FPC0 and EX2 is FPC1. Ensure the hostname for the virtual-chassis is EX1.
2) Configure Juniper's recommended priority settings for a two-member Virtual Chassis.
3) Ensure that the master device retains mastership by not taking the linecard role in case the backup device fails.
4) Virtual Chassis master election events should be logged in detail to a file called "me".

VLAN Name	VLAN ID	Devices	Interface
user	100	EX1	ge-0/0/7 untagged
		EX2	ge-0/0/5 untagged
		EX2	ge-0/0/9 untagged
		vQFX1	ge-0/0/5 untagged
		vQFX2	ge-0/0/3 untagged
server	101	EX1	ge-0/0/5 untagged
		EX2	ge-0/0/3 untagged
		vQFX1	ge-0/0/3 untagged
printer	200	EX1	ge-0/0/9 untagged
phone	201	EX1	ge-0/0/7 tagged
		EX2	ge-0/0/9 tagged
guest	666	EX1, EX2	Dot1x guest vlan
vmx5vmx6ex1	156	EX1	ge-0/0/3
			ge-0/0/4

Table 15 VLANs Site1 and Site 2

5) Configure the VLANs and corresponding interfaces in Site1 and Site 2 as shown in the table above.
6) Configure interface ge-0/0/7 and interface ge-1/0/9 on virtual chassis EX1 for Voice VLAN support. Make sure that tagged frames go to VLAN201 whereas untagged frames go to VLAN100.
7) Ensure Server1 and server 2 should receive an IPv4 address from the DHCP server.
8) Printer1 must be authenticated dynamically, but EX1 must not attempt 802.1x authentication. If authentication fails or the server is unreachable, the printer must be given guest access.

© 2020 Juniper Networks, Inc. All rights reserved.

9) Configure interface ge-0/0/7 on EX1 and interface ge-1/0/9 on EX2 for POE and LLDP-MED support. Send location-id TLVs. No more than 20 WATT should be allocated the POE ports.
10) The maximum IP/MAC bindings should be 5 on EX2 port ge-1/0/9.
11) Configure EX1 to shutdown a port if Unknown Unicast or Broadcast traffic consumes more than 5% of the interface bandwidth. Multicast traffic should not be included in your solution.

Task 3: Protocol Independent Features

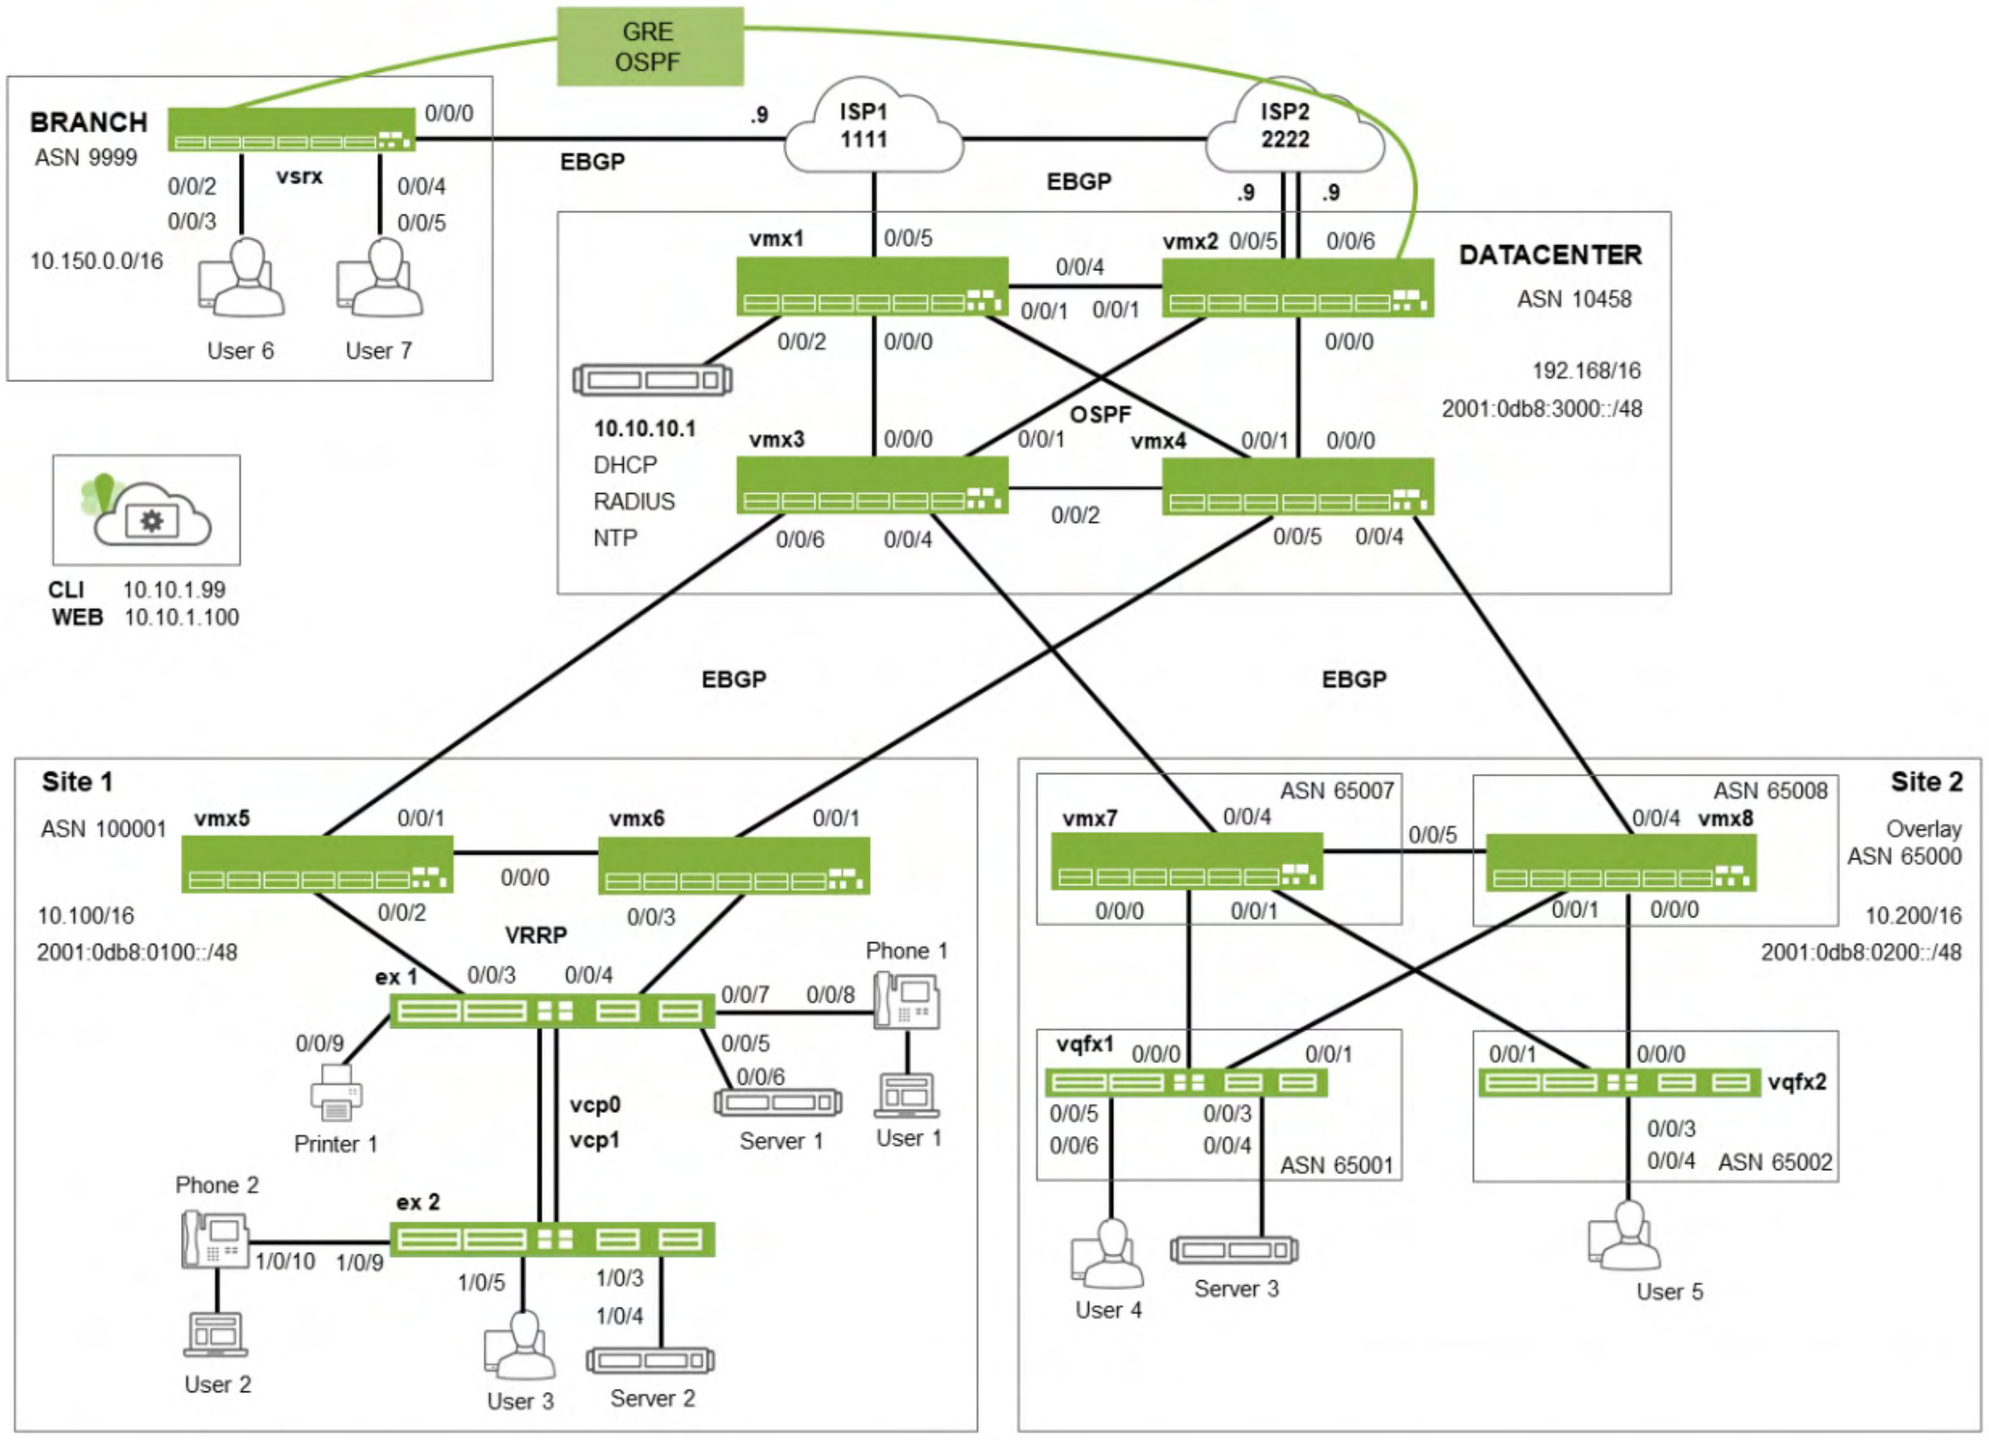

1) Configure a GRE tunnel between the branch office and vMX2 in the data center. Use the IP addressing as shown in the Lab setup Information chapter. Use the loopback interface as GRE tunnel source address.
2) Ensure that vMX2 forwards ssh and unencrypted webtraffic received from Site1 and Site2 destinated to the 18.42.206.0/24 network using the ge-0/0/6 interface. If

© 2020 Juniper Networks, Inc. All rights reserved.

interface ge-0/0/6 is unavailable, then interface ge-0/0/5 must be used. Other traffic should follow the routing requirements as mentioned in the other tasks.

3) vMX1 has a commit script stored that needs to be enabled. The commit script provides a warning message if vMX1 has IPv4 control plane filtering disabled. The commit script is currently not working and needs to be fixed.
4) Ensure that interface ge-0/0/0 on vMX1 is disabled by the Junos OS when it goes down 2 times within 1 minutes. Do not use an external script.
5) Ensure that event-policies are not allowed to be executed simultaneously.
6) Ensure that on-box Python scripts are allowed on vMX1.
7) Ensure that Python commit scripts are allowed to be stored off-box. Other commit scripts should only be allowed to be stored on-box.

Task 4: IGP

This chapter focuses on the implementation of various IGPs like OSPFv2, OSPv3 and also static routing in combination with VRRP

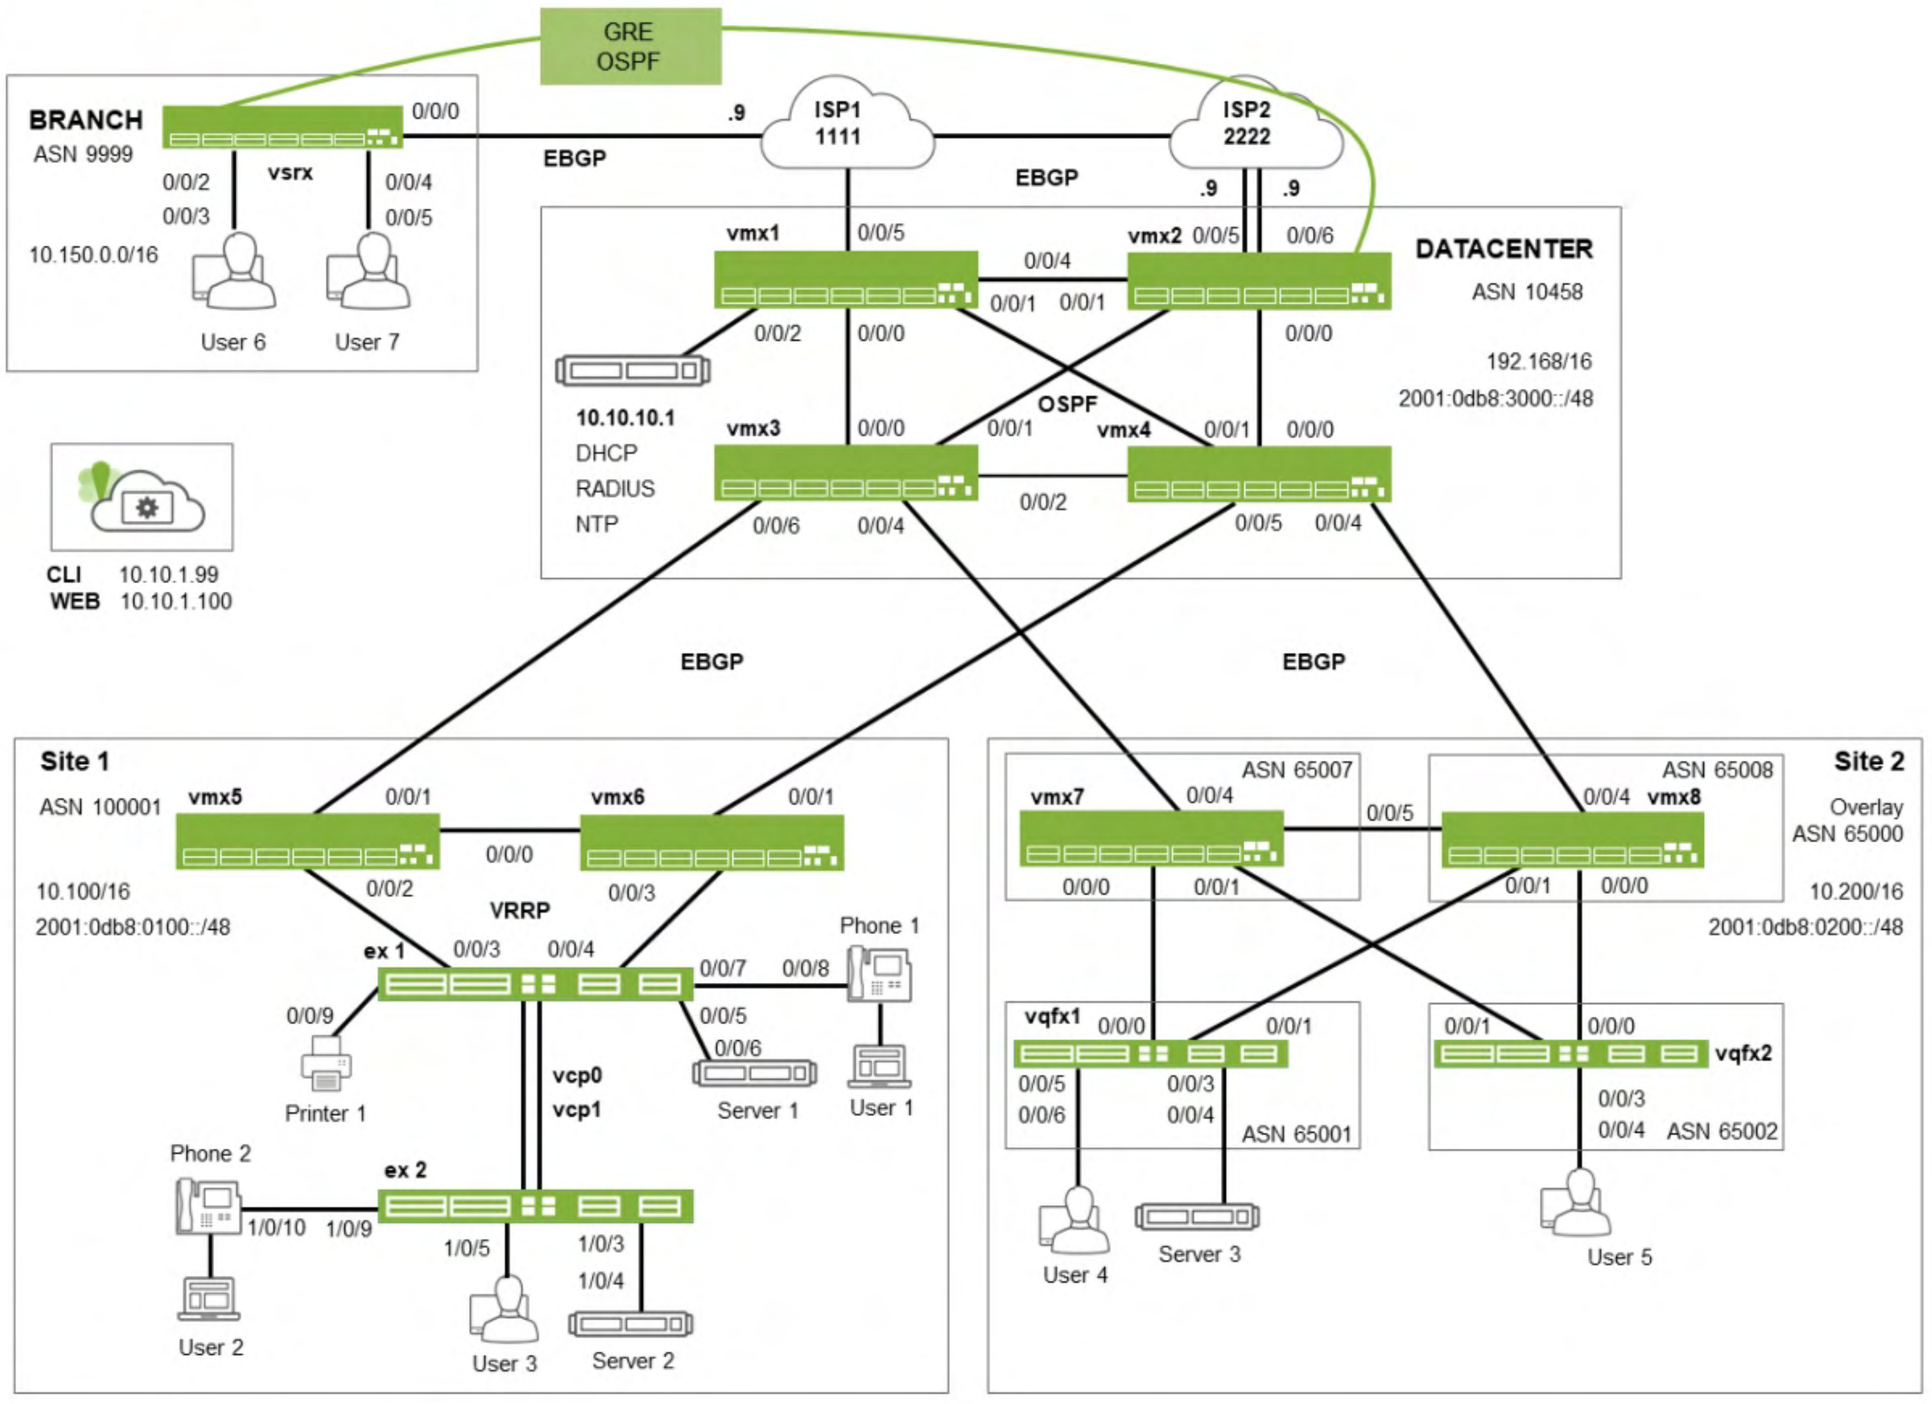

© 2020 Juniper Networks, Inc. All rights reserved.

Site 1:

1) Configure static routing over the ge-0/0/0 interface to ensure vMX5 and vMX6 can reach each other loopback address.
2) Create two static summary routes on vMX5 and vMX6 that encompass the EX1 IRB interfaces.
3) Configure VRRP on vMX5 and vMX6 to provide a redundant default gateway for EX1. Make sure that the virtual address responds to icmp packets and enable MD5 authentication with password “juniper”.
4) vMX5 should be the VRRP master for IPv4 traffic. vMX6 should be the master for IPv6 traffic. After recovery of a failure the VRRP master should revert back to its role.
5) vMX5 should become VRRP master for IPv6 in case vMX6 interface ge-0/0/1 fails, then
6) Create two static default routes on EX1 towards the VRRP virtual address (.254) of vMX5 and vMX6.

Data Center:

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 2	lo0	0.0.0.0
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 3	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 4	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1

Table 16 OSPF Information Data Center

7) Configure OSPFv3 in the data center using the information as shown in the tables above.
8) Ensure no LSA type 0x2002 are found in the OSPFv3 databases in Site 1.

© 2020 Juniper Networks, Inc. All rights reserved.

9) Configure the OSPF interface timers so that a neighbor is declared dead after four missed hello packets.
10) Ensure that each vMX devices in the data center learn each other's IPv6 loopback address.

Site 2:

Device	Interface	area
vMX7	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0
vMX8	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0

Table 17 OSPF information site 2

11) Configure OSPFv3 in Site 2 using the information as shown in the tables above.
12) Ensure no LSA type 0x2002 are found in the OSPFv3 databases in Site 1.
13) Configure the OSPF interface timers so that a neighbor is declared dead after four seconds.

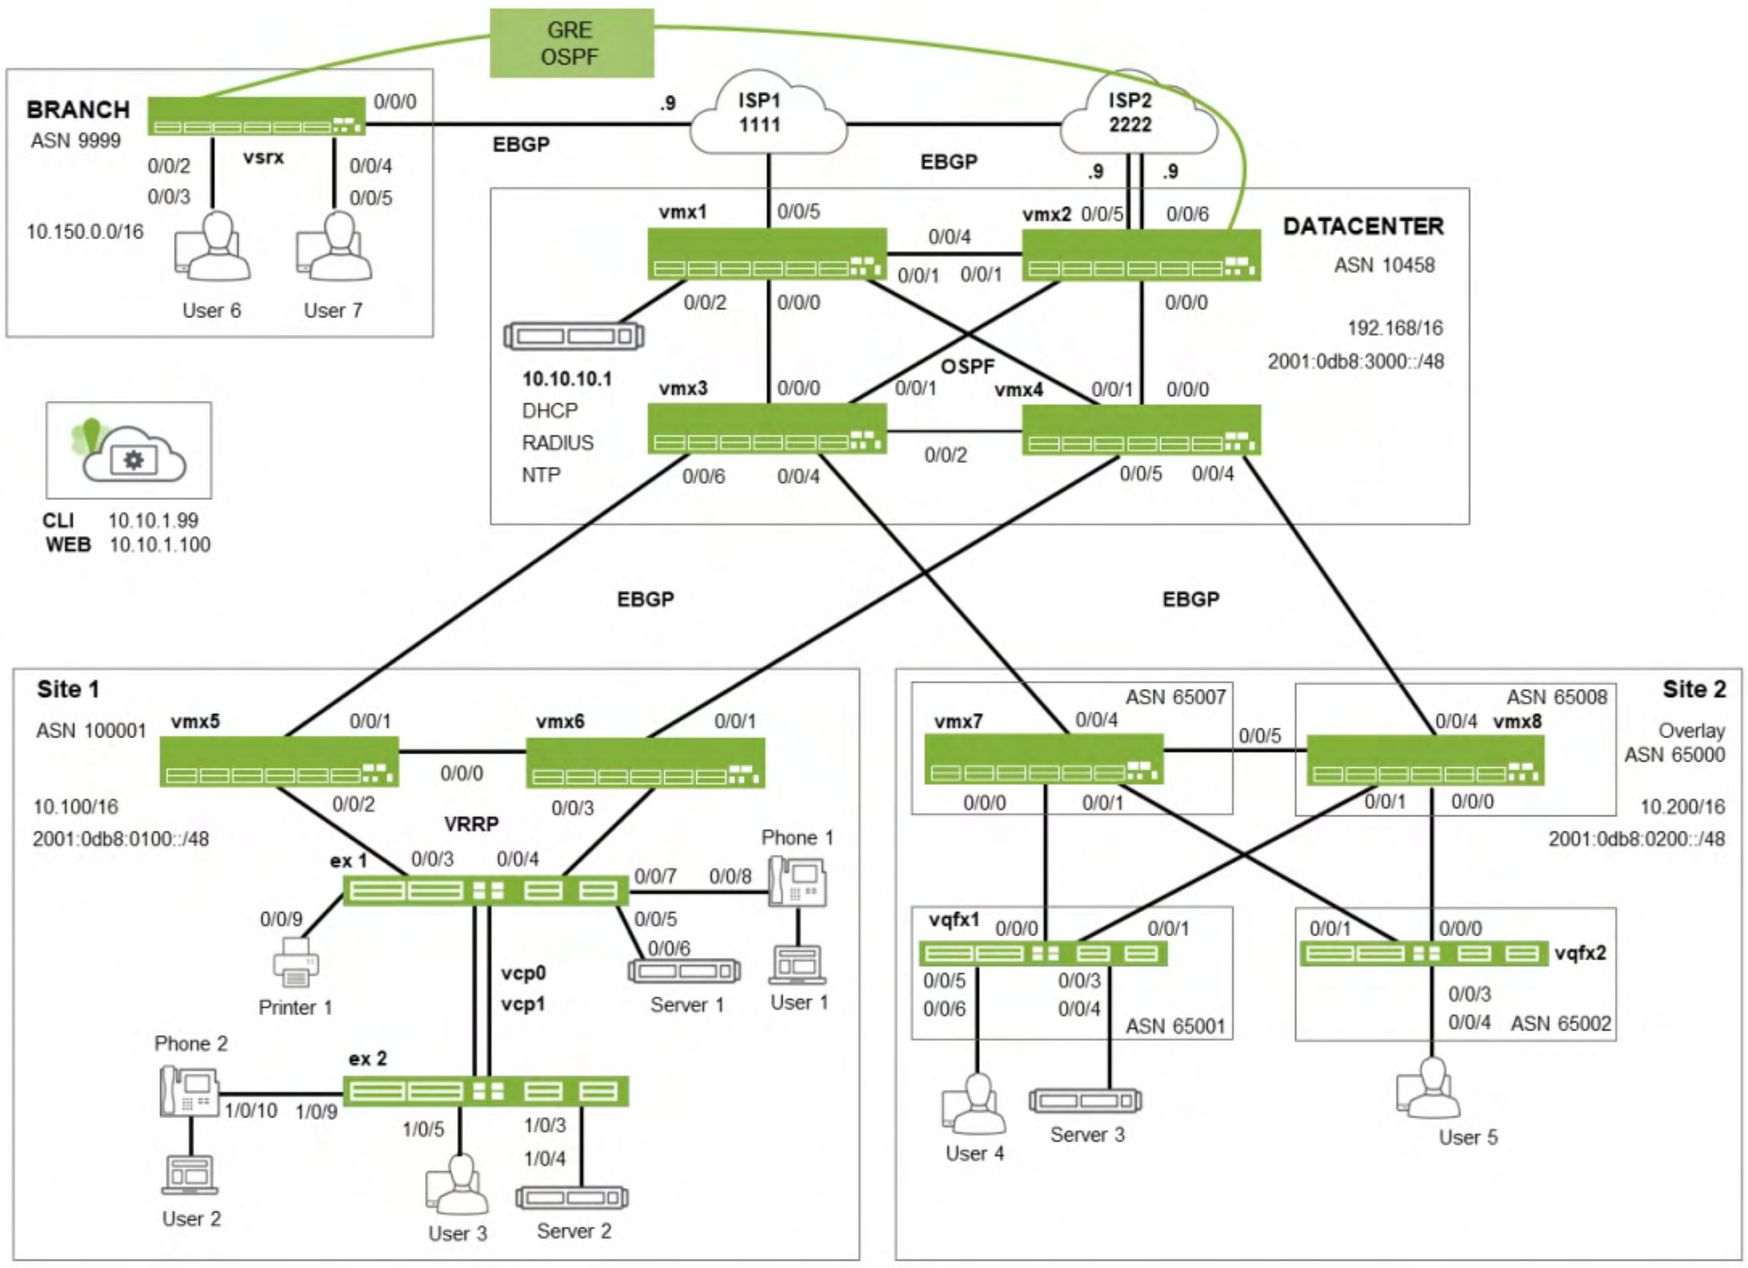

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Interface	Routing instance	area
vMX2	GRE interface	ospf-branch	0.0.0.0
vSRX	GRE interface	default	0.0.0.0

Table OSPF Information for GRE Tunnel

14) Enable OSPFv2 over the GRE interface as shown in the table above.
15) Ensure that the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.
16) Ensure that vSRX OSPF database contains a default route and the other OSPFv2 destinations received from vMX2. Configure OSPF on vSRX that the default route is the only external LSA that is installed in its routing table.

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 2	lo0	0.0.0.0
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 3	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX 4	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1

17) Enable OSPFv2 for IPv4 routing on the vMX devices in the data center. Use the information as shown in the table above.
18) Ensure that OSPFv2 calculates backup paths for (most) of the LSAs found in the OSPF database.
19) Ensure that Gigabit Ethernet interfaces have an OSPF cost of 100. Do not make any changes under the OSPF interface hierarchy.

Task 5: BGP

This task focuses on the implementation of BGP for routing within the data center and within Site 2. BGP is also used for routing between the data center, ISP1, ISP2, and

© 2020 Juniper Networks, Inc. All rights reserved.

between the data center, Site 1, and Site 2. At various points in the topology you need to perform route redistribution between and IGP and BGP for end-to-end connectivity.

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Interface	IPv4 address	IPv6 address	Description
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128	IBGP source interface
				v4 EBGP to AS 2222
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64	v6 EBGP to AS 2222
	ge-0/0/6	172.16.3.2/24	Link local	v6 EBGP to AS 2222
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128	IBGP source interface
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64	EBGP to AS 64513
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64	EBGP to AS 100001
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128	IBGP source interface
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64	EBGP to AS 64513
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64	EBGP to AS 100001
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table BGP Peering Information

© 2020 Juniper Networks, Inc. All rights reserved.

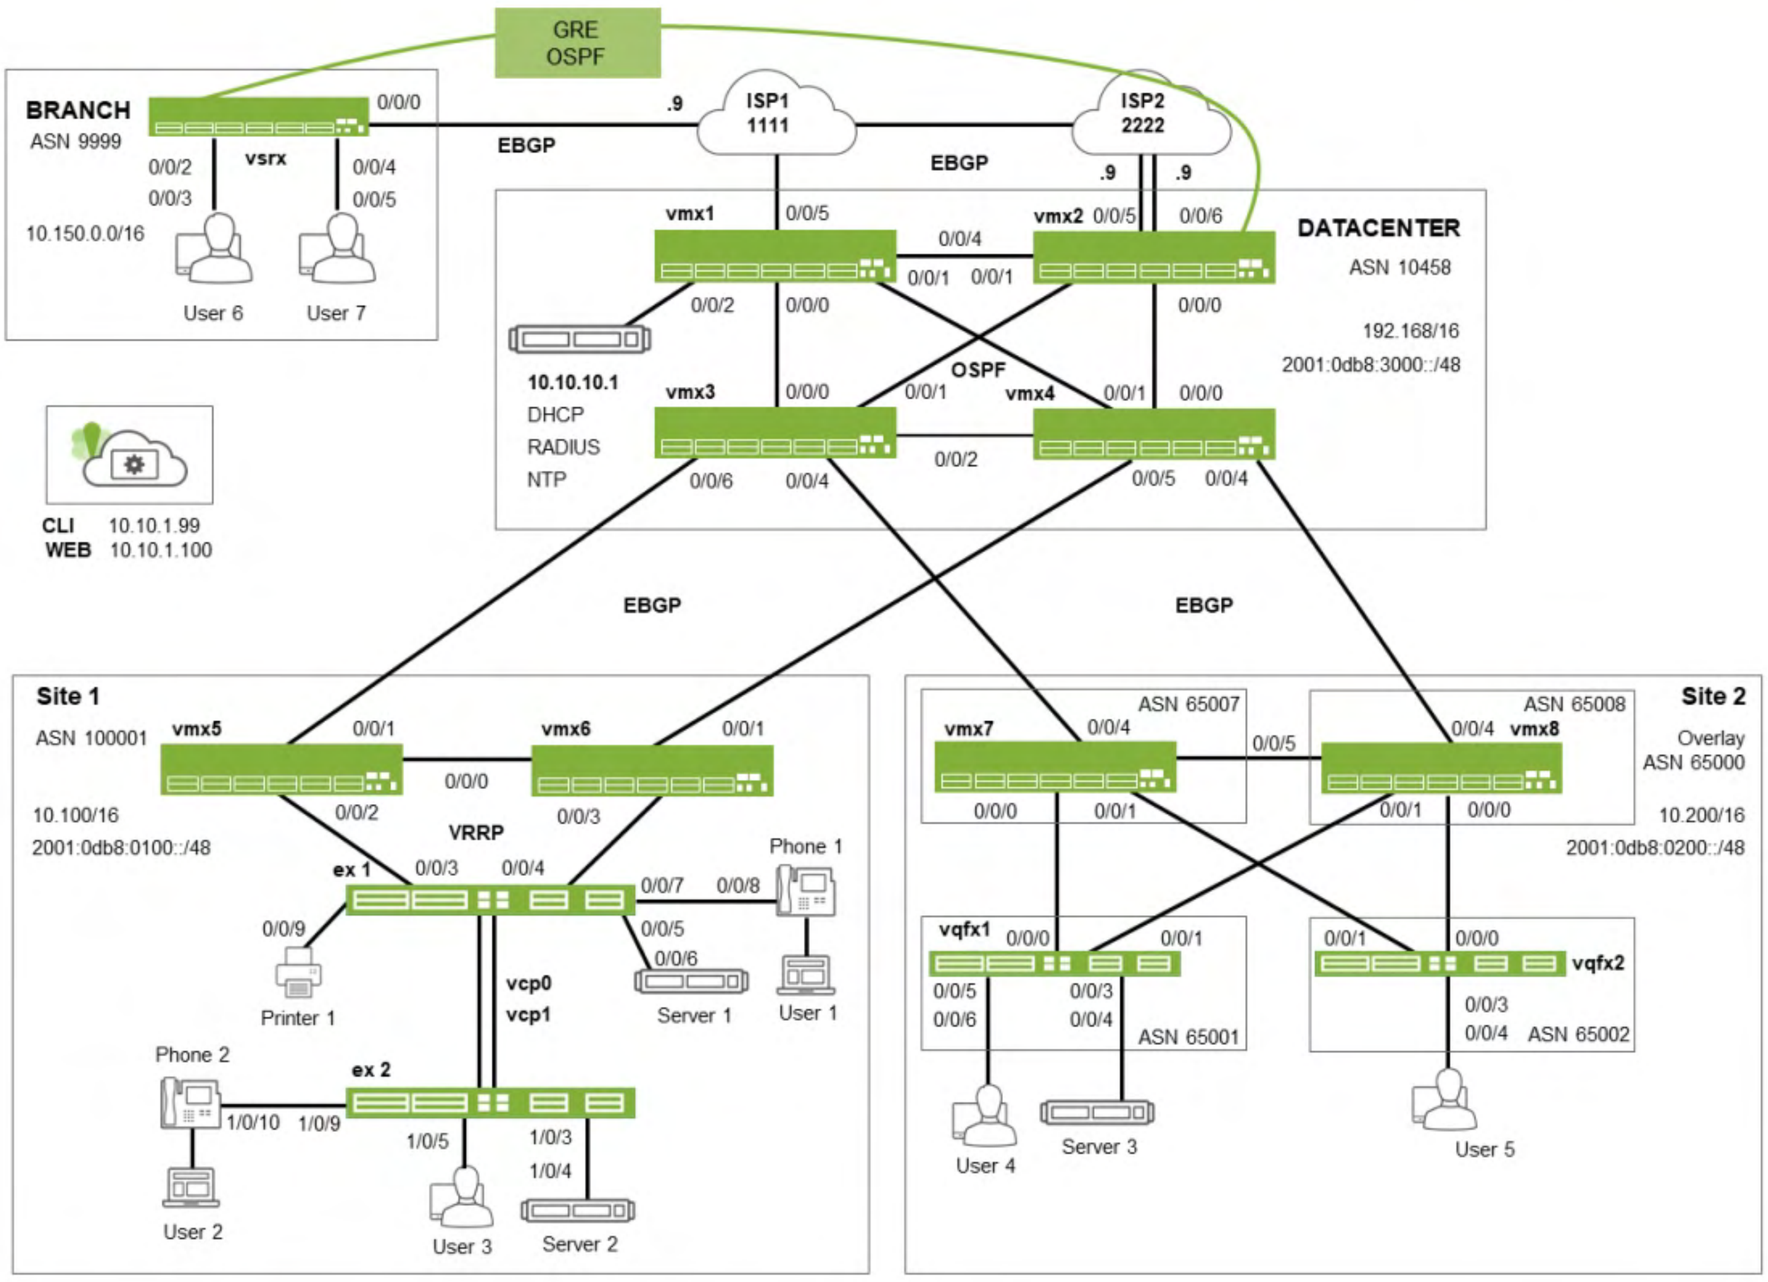

IBGP Data Center

1) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have an IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and accept routes received from each other.
2) All IBGP sessions on vMX1 and vMX2 should be established on port 179 locally. Do not use the passive command or a firewall filter to accomplish this.
3) Enable BFD with authentication for IBGP sessions in the data center. Allow multihop sessions. Do not configure BFD between the Route Reflectors.
4) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS and the Branch Office.
5) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down. vMX2 should not advertise the default route if the BGP peering with ISP2 goes down.
6) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the AS.

© 2020 Juniper Networks, Inc. All rights reserved.

IBGP Site 1 and Site 2

7) Configure IBGP in Site1 between vMX5 and vMX6. Use the information as shown in the table above.

EBGP Data Center

8) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
9) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.
10) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.
11) Advertise the IPv4 summary routes received from Site 1, Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs.
12) Advertise a single /32 route to both ISPs that encompasses the IPv6 networks of the data center, Site 1, and Site 2.
13) Ensure that ISP2 is the primary path to all IPv4 address blocks announced by your AS. Both ISPs apply route policy based on received communities. Use the information in the table below to meet the task requirement.

ISP	COMMUNITY	ISP ACTION
ISP1	1111:1	Prepend 1x
	1111:2	Prepend 2x
	1111:3	Prepend 3x
ISP2	2222:1	Prepend 1x
	2222:2	Prepend 2x
	2222:3	Prepend 1x

14) Ensure the ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.
15) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2. Prefixes from Site1 that are advertised to both ISPs should contain the ASN of Site 1.
16) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Site 1 and Site 2.

© 2020 Juniper Networks, Inc. All rights reserved.

17) Site 1 and Site 2 should learn the summarized prefixes of the other Site, data center and Branch office. Site 1 and Site 2 should also learn the IPv4 and IPv6 default route and the management network 10.10.10/24 that originated in the data center.
18) Ensure the device in Site1, Site2, and Branch Office have full reachability to each other and the Internet.

Task 6: EVPN/VXLAN (Spine and Leaf)

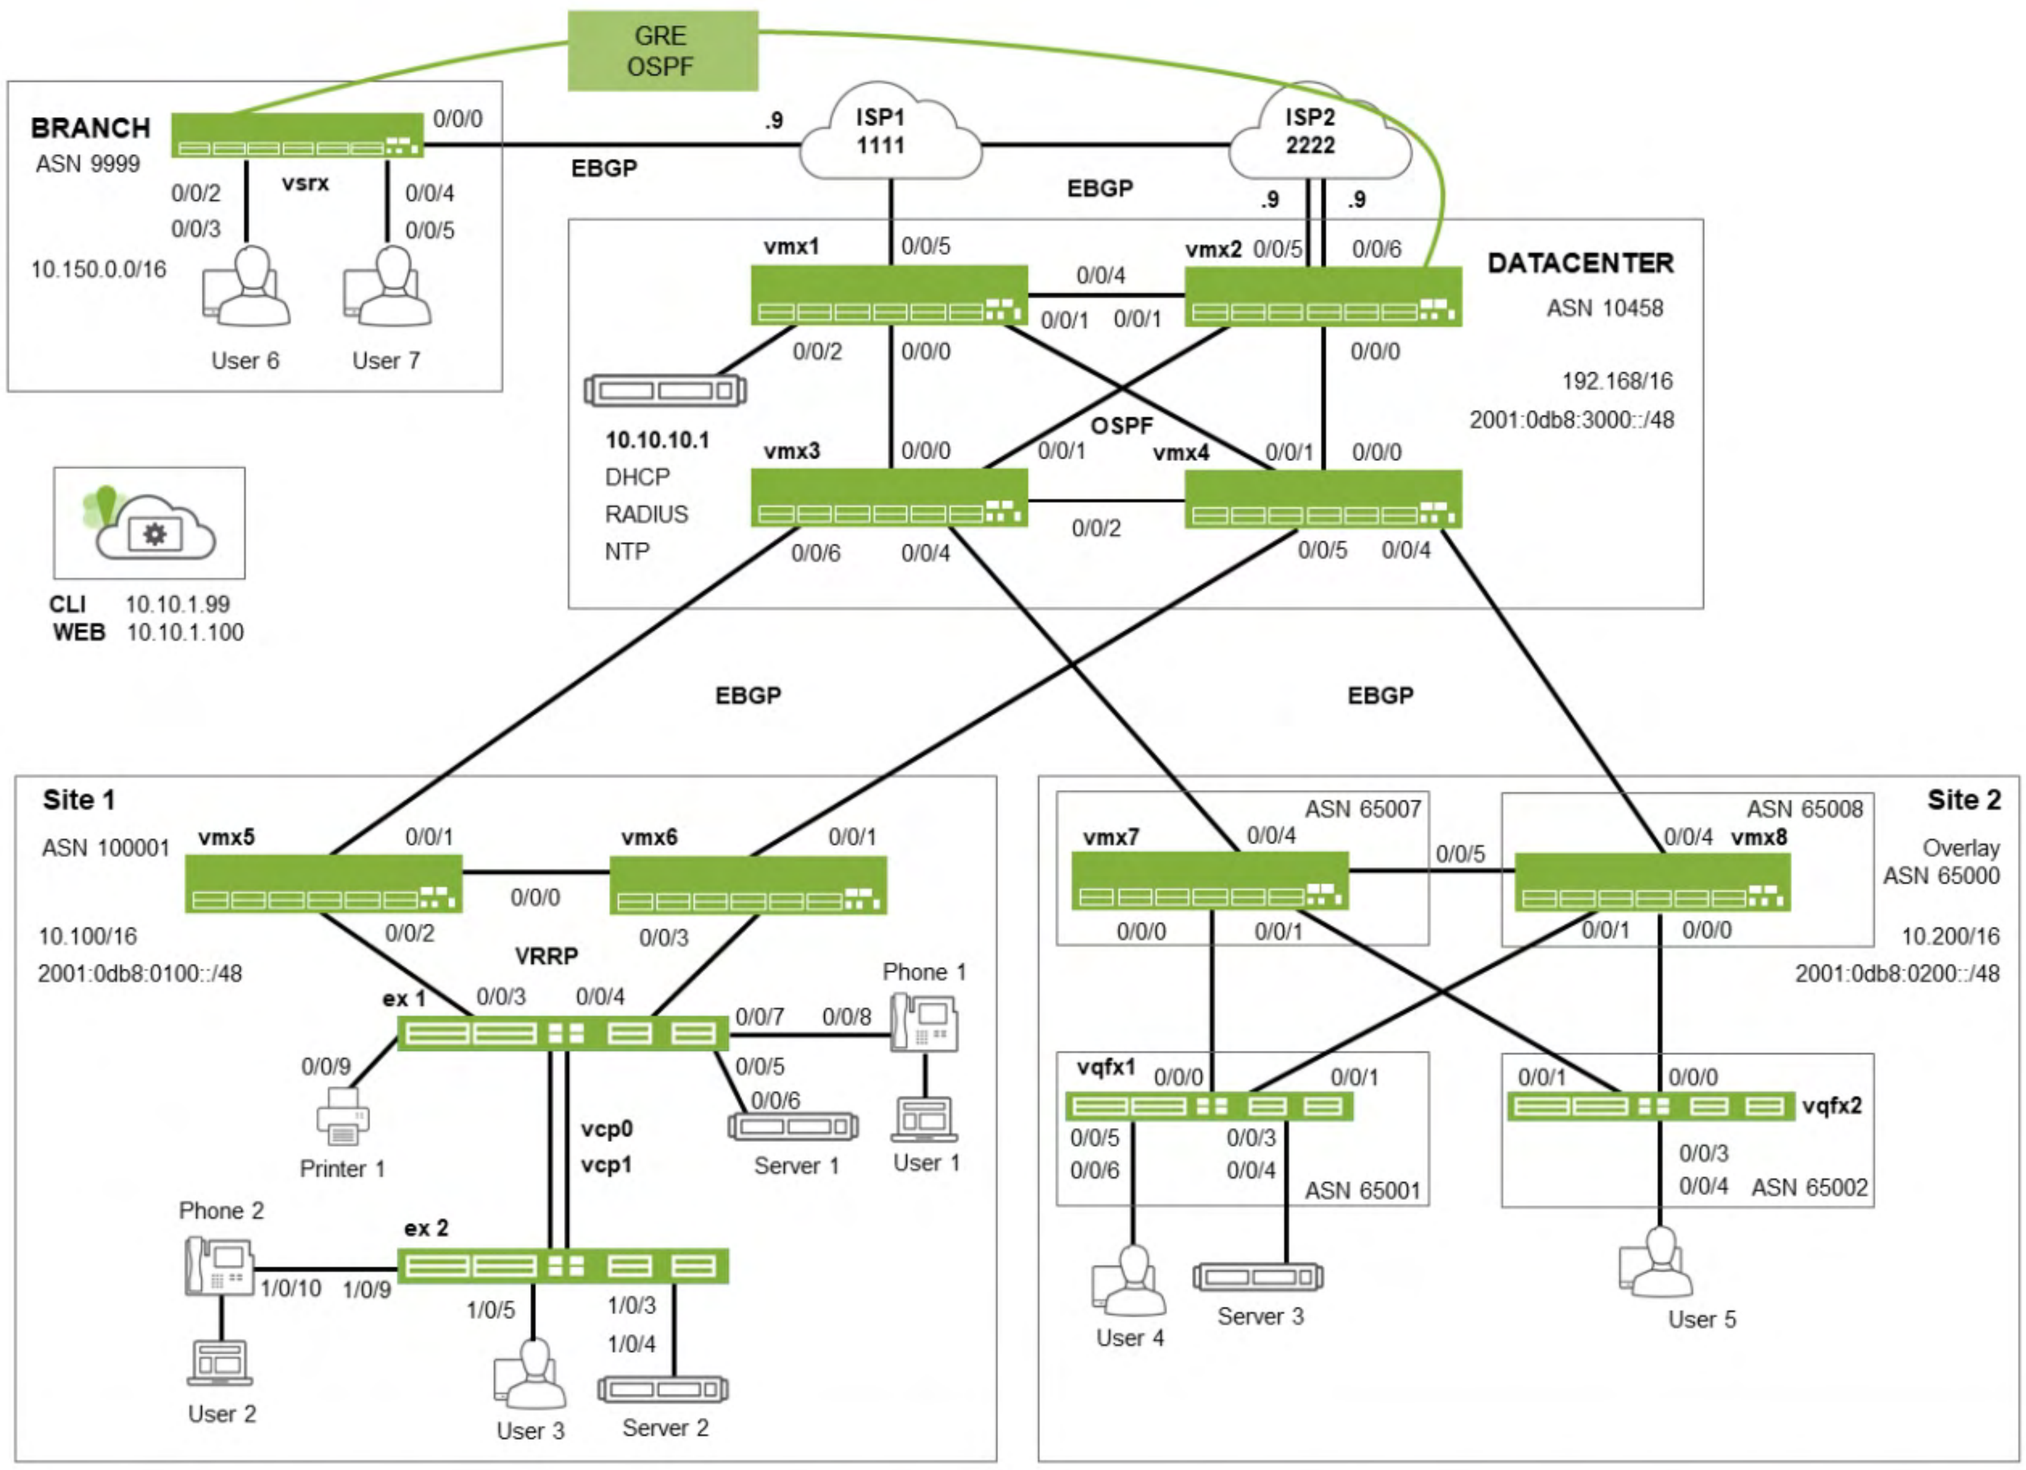

In this task, you will configure the layer 2 connectivity in Site 2 using the EVPN/VXLAN technology.

1) Configure the EBGP underlay in Site 2 to exchange device loopback reachability. Use the AS numbers as shown in the topology. Use the interface addresses to establish the EBGP peerings.
2) Configure the IBGP EVPN overlay for Site 2. vMX7 and vMX8 are router reflectors for the EVPN overlay. Leaf devices should only peer with Spine devices.

© 2020 Juniper Networks, Inc. All rights reserved.

3) Ensure that virtual gateway addresses respond to ICMP packets.
4) Configure one EVPN VLAN aware service on vMX7 and vMX8 for VLAN 100 and 101.
5) Ensure the devices in Site 2 can reach each other, the device in Site1 and the Internet.

Task 7: Class of Service

This task is focused on Class of Services. You will configure and monitor MF and BA Classifiers, Two-color Policers, different Forwarding classes with respected Queues and Schedulers, DSCP Rewrite markers, RED Drop profiles and Class based forwarding.

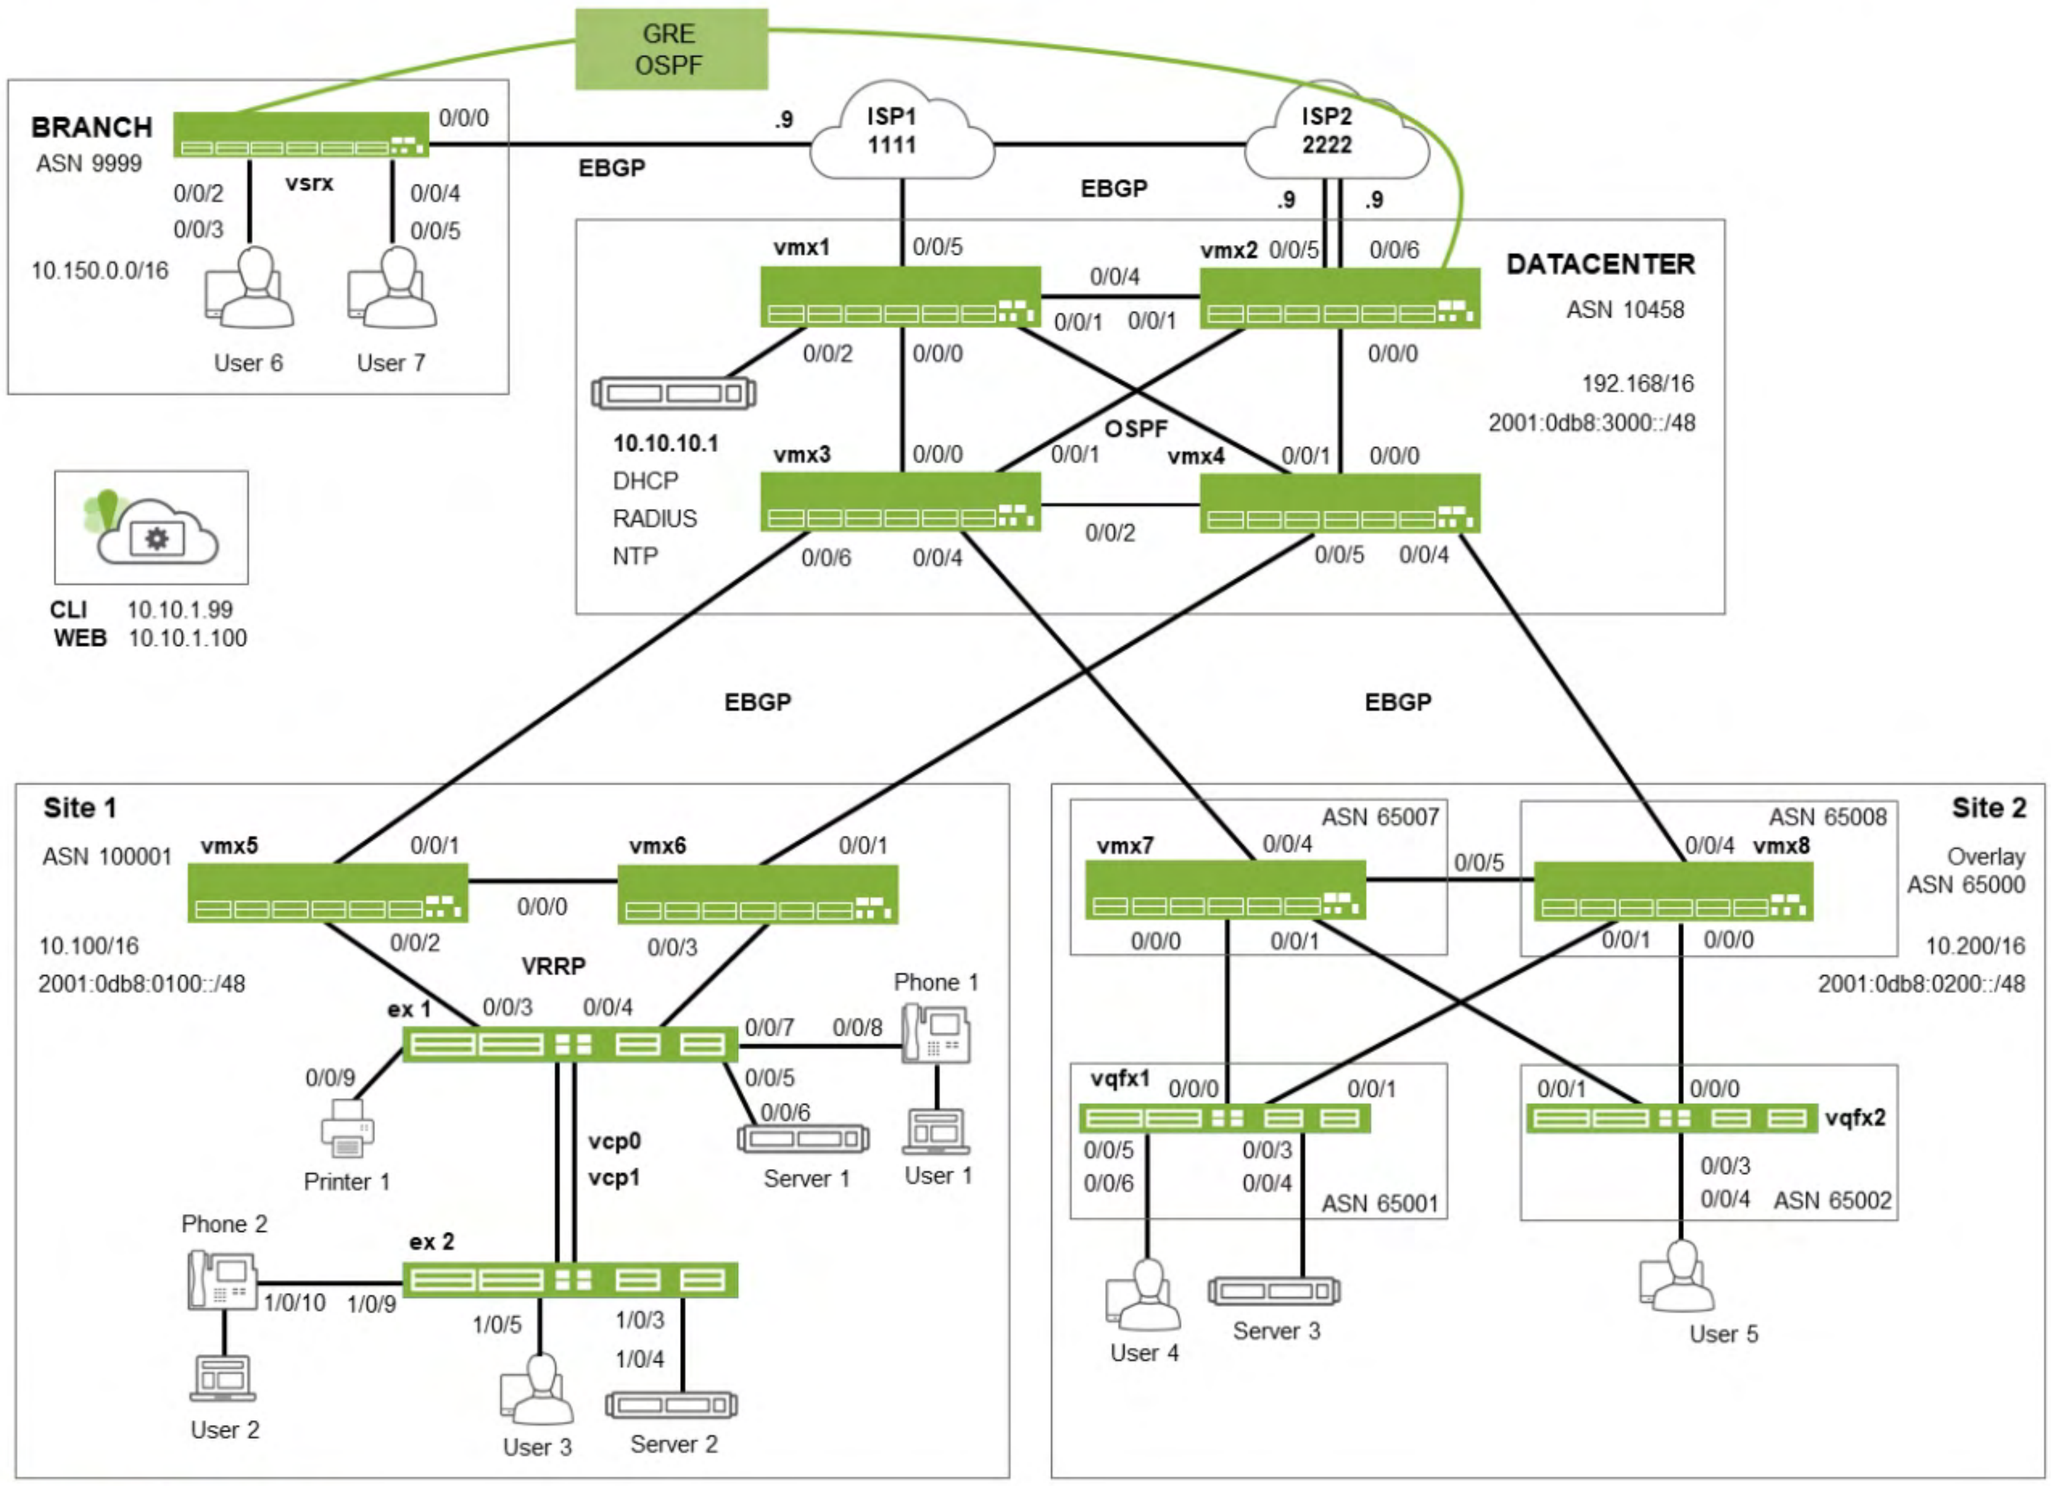

Figure 4 Topology

© 2020 Juniper Networks, Inc. All rights reserved.

In this task you will configure all vMX to support Class of Service.

1) Configure Forwarding Classes and map them to the output Queues as indicated in the table below.

Forwarding Class	Queue	Scheduler
control	7	control-sched
voice	5	voice-sched
server-gold	3	server-gold-sched
server-silver	2	server-silver-sched
server-bronze	1	server-bronze-sched
best-effort	0	best-effort-sched

2) Configure schedulers with the parameters in Table 18.

© 2020 Juniper Networks, Inc. All rights reserved.

Scheduler	Parameter	Value
control-sched	priority	medium-high
	transmit rate	5%
	buffer-size	10%
voice-sched	priority	strict-high
	transmit rate	10%
	buffer-size	20 (milliseconds)
server-gold-sched	priority	high
	transmit rate	25%
	buffer-size	10% exact
server-silver-sched	priority	medium-high
	transmit rate	35%
	buffer-size	50% exact
	drop profile high	relaxed
server-bronze-sched	priority	medium-low
	transmit rate	remainder
	buffer-size	remainder
	drop profile low	relaxed
	drop profile high	aggressive
best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder
	drop profile any	aggressive

Table 18

3) Configure two interpolated Drop profiles listed in Table 19.

Drop profile	Fill level (%)	Drop probability (%)
relaxed	75	0
	90	75
	100	100
aggressive	50	0
	75	50
	90	75
	100	100

Table 19

© 2020 Juniper Networks, Inc. All rights reserved.

4) Apply the schedulers to all physical gigabit ethernet interfaces.

5) Configure IPv4 and IPv6 Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

6) Apply the classifiers to all (logical) interfaces.

7) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

8) Apply the rewrite rules on all logical interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Super Lab 2

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

Lab Setup Information

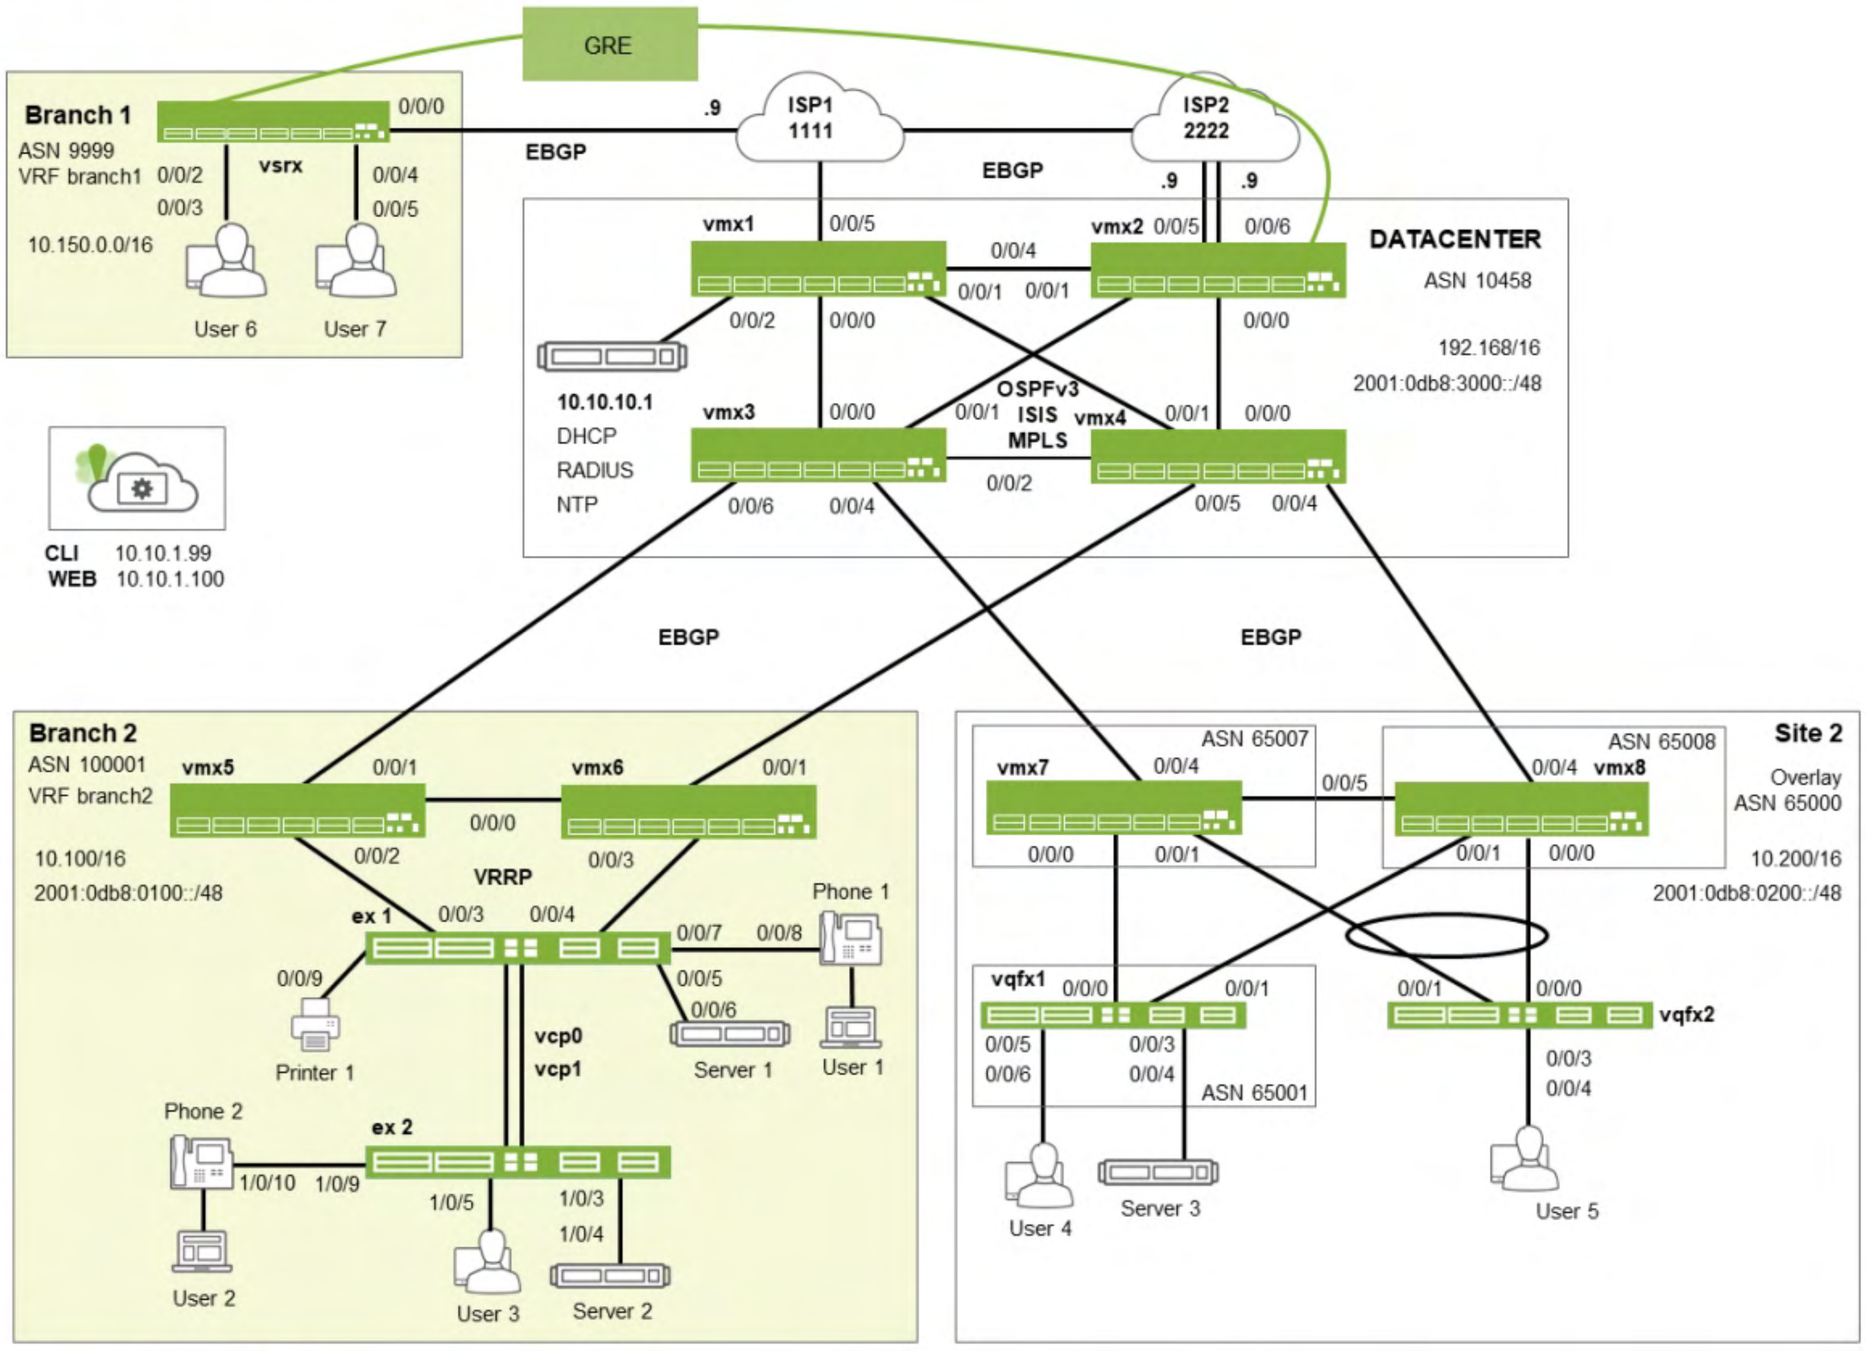

Figure 5 Topology

In the figure above, the topology is shown that is used for this super lab. There are four vMX devices in the data center. vMX1 and vMX2 are both connected to an ISP. vMX3 and vMX4 are connected to vMX5 and vMX6 in Branch2. Branch2 contains two EX4300 devices (EX1 and EX2) that need to be configured as a Virtual Chassis. The Virtual Chassis has several devices attached, such as phones, printers, servers and users. Please note that some of these devices are simulated using routing-instances on the EX4300

© 2020 Juniper Networks, Inc. All rights reserved.

switches. For example, server2 on EX2 has a layer3 port (ge-1/0/4) that participates in a routing-instance. There is a physical Ethernet connection between port ge-1/0/4 and ge-1/0/3, where port ge-1/0/3 participates as a layer2 port. Printer1 and User2 are virtual machines that are accessible (see lab information for more information). The data center vMX3 and vMX4 are also connected to Site2 vMX7 and vMX8. vMX7 and vMX8 have connections to vQFX1 and vQFX2. Like the EX4300 switches in Branch 2 the vQFX switches also have users and a server attached. Branch 1 contains two users and is connected to ISP1.

Device	OoB Interface	OoB IPv4 address
vMX1	fxp0	10.10.1.1/24
vMX2	fxp0	10.10.1.2/24
vMX3	fxp0	10.10.1.3/24
vMX4	fxp0	10.10.1.4/24
vMX5	fxp0	10.10.1.5/24
vMX6	fxp0	10.10.1.6/24
vMX7	fxp0	10.10.1.7/24
vMX8	fxp0	10.10.1.8/24
VR	fxp0	10.10.1.9/24
vSRX	fxp0	10.10.1.10/24
EX1	ge-0/0/0	10.10.1.11/24
EX2	ge-0/0/0	10.10.1.12/24
vQFX1	em0	10.10.1.13/24
vQFX2	em0	10.10.1.14/24
Junos Space	N/A	10.10.1.99/24 (CLI) 10.10.1.100/24 (WEB)

Table 2 Management network (out of band)

The table above shows the devices used in the topology and their management IP addresses. These addresses are preconfigured and should not be changed and are listed here for informational purposes only.

Device	Interface	Ipv4 address	Ipv6 address
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128
	ge-0/0/0	192.168.13.1/24	2001:0db8:3000:0013::1/64
	ge-0/0/1	192.168.14.1/24	2001:0db8:3000:0014::1/64
	ge-0/0/2	10.10.10.100/24	N/A
	ge-0/0/4	192.168.12.1/24	2001:0db8:3000:0012::1/64
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/0	192.168.24.2/24	2001:0db8:3000:0024::2/64
	ge-0/0/1	192.168.23.2/24	2001:0db8:3000:0023::2/64
	ge-0/0/4	192.168.12.2/24	2001:0db8:3000:0012::2/64
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64
	ge-0/0/6	172.16.3.2/24	2001:0dba:4000:0003::2/64
	GRE	173.17.99.2/24	
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128
	ge-0/0/0	192.168.13.3/24	2001:0db8:3000:0013::3/64
	ge-0/0/1	192.168.23.3/24	2001:0db8:3000:0023::3/64
	ge-0/0/2	192.168.34.3/24	2001:0db8:3000:0034::3/64
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128
	ge-0/0/0	192.168.24.4/24	2001:0db8:3000:0024::4/64
	ge-0/0/1	192.168.14.4/24	2001:0db8:3000:0014::4/64
	ge-0/0/2	192.168.34.4/24	2001:0db8:3000:0034::4/64
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128
	ge-0/0/0	10.100.56.5/24	2001:0db8:0100:0056::5/64
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64
	ge-0/0/2	10.100.15.5/24	2001:0db8:0100:015::5/64
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128
	ge-0/0/0	10.100.56.6/24	2001:0db8:0100:0056::6/64
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64
	ge-0/0/3	10.100.116.6/24	2001:0db8:0100:0016::6/64
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128
	ge-0/0/0	10.200.17.7/24	N/A
	ge-0/0/1	Layer 2	N/A
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64
	ge-0/0/5	192.168.78.7/24	2001:0db8:3000:0078::7/64
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128
	ge-0/0/0	Layer 2	N/A
	ge-0/0/1	10.200.18.8/24	N/A
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64
	ge-0/0/5	192.168.78.8/24	2001:0db8:3000:0078::8/64
vSRX	lo0	10.150.0.99/32	
	ge-0/0/0	10.1.1.2/30	

© 2020 Juniper Networks, Inc. All rights reserved.

EX1	ge-0/0/2	Layer 2	
	ge-0/0/3	10.150.99.4/24	
	ge-0/0/4	Layer 2	
	ge-0/0/5	10.150.99.5/24	
	GRE	173.17.99.1/24	
	lo0	10.100.0.1/32	2001:0db8:0100:0000::1/128
	ge-0/0/3	10.100.15.1/24	2001:0db8:0100:015::1/64
	ge-0/0/4	10.100.16.1/24	2001:0db8:0100:016::1/64
	ge-0/0/5	Layer 2	
	ge-0/0/6	dhcp	2001:0db8:0100:101::1/64
	ge-0/0/7	Layer 2	
	ge-0/0/8	10.100.201.1/24	2001:0db8:0100:201::1/64
	irb.100	10.100.100.254/24	2001:0db8:0100:100::254/64
	irb.101	10.100.101.254/24	2001:0db8:0100:101::254/64
	irb.156	10.100.156.1/24	2001:0db8:0100:0156::1
	irb.200	10.100.200.254/24	2001:0db8:0100:200::254/64
	irb.201	10.100.201.254/24	2001:0db8:0100:201::254/64
EX2	ge-1/0/3	Layer 2	
	ge-1/0/4	dhcp	
	ge-1/0/5	Layer 2	
	ge-1/0/6	Layer 2	
	ge-1/0/9	Layer 2	
	ge-1/0/10	10.100.201.2/24	2001:0db8:0100:201::2/64
vQFX1	xe-0/0/0	10.200.17.1/24	
	xe-0/0/1	10.200.18.1/24	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.101.1/24	2001:0db8:0200:101::1/64
	xe-0/0/5	Layer 2	
	xe-0/0/6	10.200.100.1/24	2001:0db8:0200:100::1/64
vQFX2	xe-0/0/0	Layer 2	
	xe-0/0/1	Layer 2	
	xe-0/0/3	Layer 2	
	xe-0/0/4	10.200.100.3/24	2001:0db8:0200:100::3/64

The table above shows the IP addressing scheme that is used throughout this super lab.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: General System Features

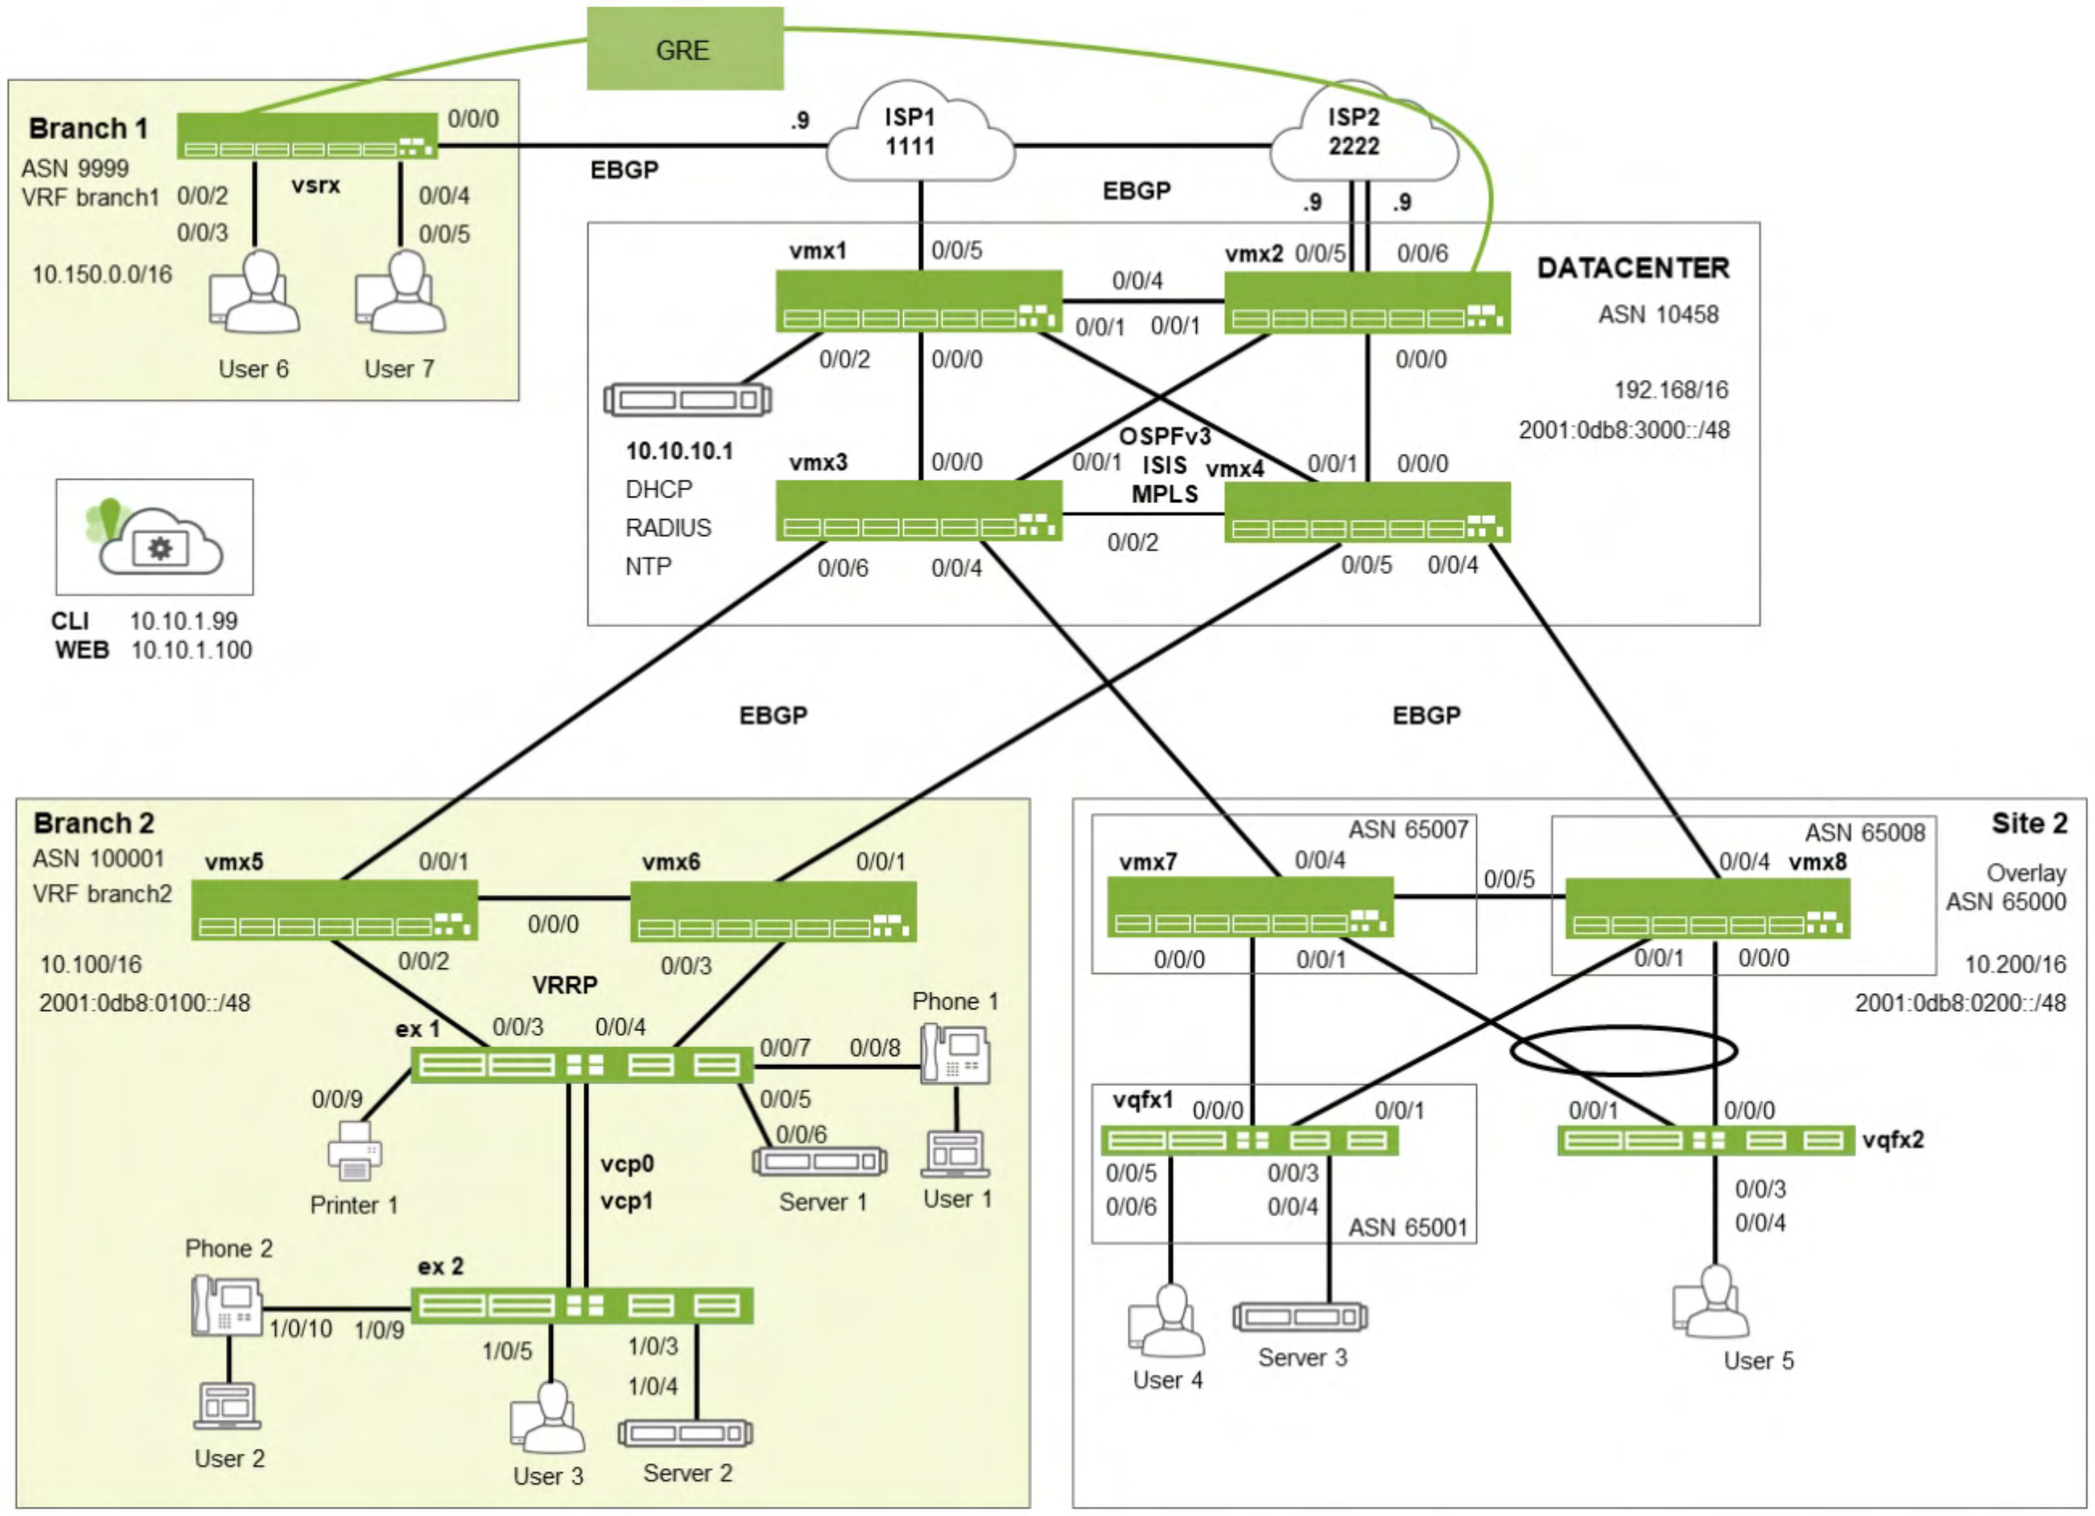

Note: For this task you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place.

1) Configure all vMX and EX devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute. Maximum 32 sessions are allowed per SSH connection.
2) Enable the FTP service.
3) Set the time zone to *America/Los_Angeles* on all vMX devices, EX1.
4) Make sure that all vMX devices, EX1 synchronize the time with NTP server 10.10.10.1.

© 2020 Juniper Networks, Inc. All rights reserved.

NOTE: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1.*

The following tasks should be configured on all vMX devices and EX1.

1) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts 3 and timeout of 3 seconds. Use juniper123 as the RADIUS shared secret.
2) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices and EX1 as a super-user.
3) Configure the following additional users on the devices as defined in the table below:

Username	Password	Privileges
ronly	ronly123	Has permissions “view” and “view-configuration”. Additionally, the user can NOT execute the “file delete” command.
admin1	admin123	Has permissions “all”. Can access the configuration mode only using “configure private” command.
restricted	restricted123	Has permissions “clear” and execute only the “show system uptime”, “show system storage” and “show interfaces terse” commands and nothing else.

The following tasks should be configured on all vMX, vQFX and vSRX devices in the topology.

4) The SNMP read-only community string is jncie. Make sure that only the Junos Space server can poll the devices in the data center, Branch2 , Site2 and the branch device.
5) Configure all devices to send SNMP traps to Junos Space. Make sure that traps are only send using the OoB management interface.

The following tasks should be configured on all vMX devices in Branch 2. Configure RE (Routing Engine) protection using a firewall filter.

© 2020 Juniper Networks, Inc. All rights reserved.

6) Configure an IPv6 firewall filter which allows protocol messages from VRRP and BGP. BGP peers that are configured in the future must be automatically allowed in the firewall filter.

7) Configure the devices to accept IPv6 ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.

8) Configure the IPv6 firewall filter to discard any other traffic and count the packets.

9) Apply the IPv6 firewall filter such as to ensure that it is used for the RE protection.

10) Ensure that the configuration under firewall hierarchy is not displayed when printed out.

Task 2: Layer 2 Configuration and Features

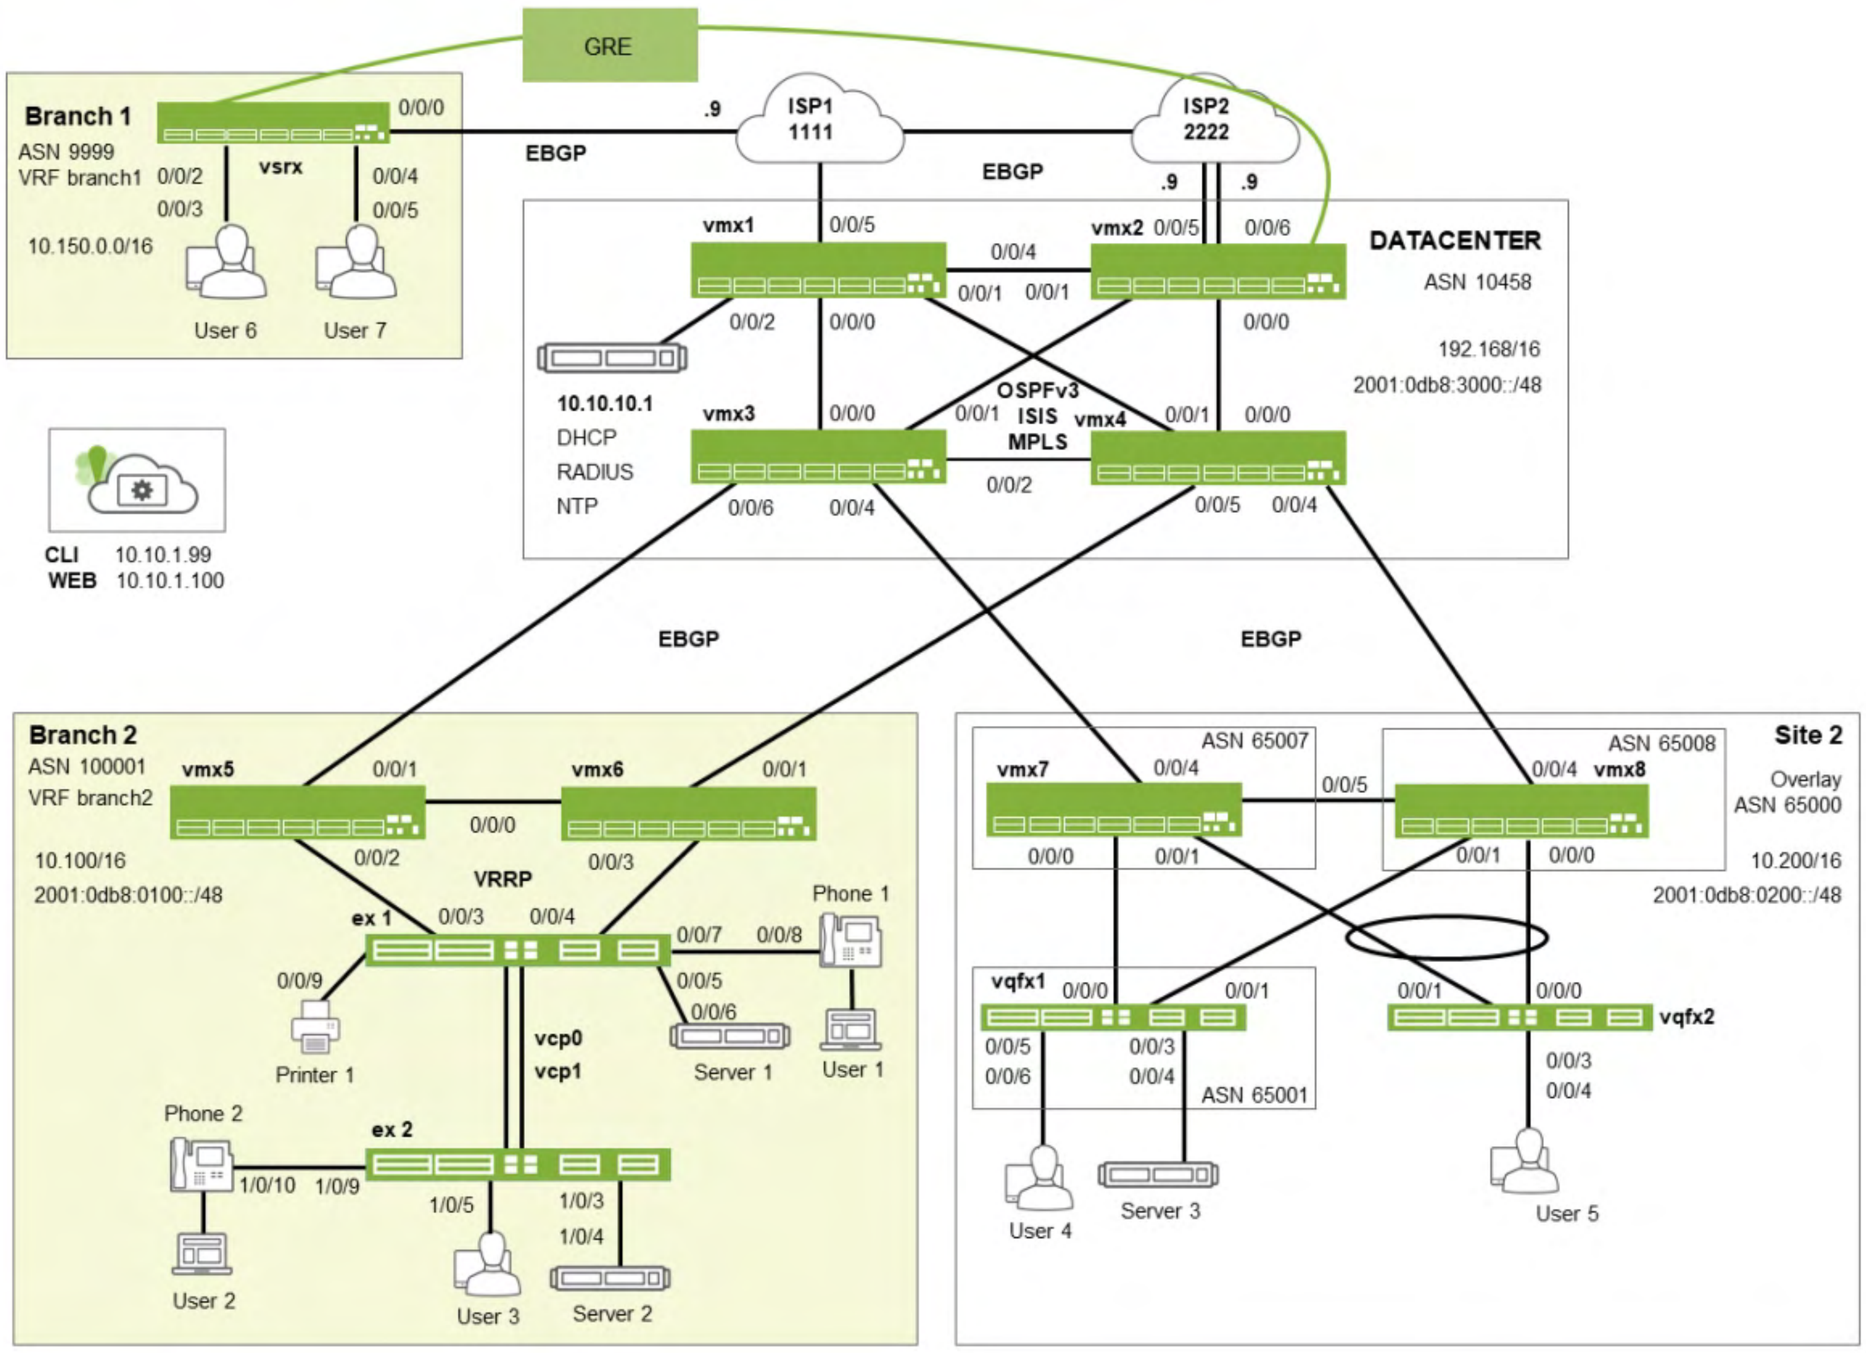

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure EX1 and EX2 as a virtual chassis. Ensure EX1 is FP0 and EX2 is FPC1. Ensure the hostname for the virtual-chassis is EX1.
2) Configure Juniper recommended priority settings for a two-member VC.
3) Ensure that the master device retains mastership by not taking the linecard role in case the backup device fails.

VLAN Name	VLAN ID	Devices	Interface
user	100	EX1	ge-0/0/7 untagged
			ge-0/0/5 untagged
			ge-1/0/9 untagged
		vQFX1	ge-0/0/5 untagged
		vQFX2	ge-0/0/3 untagged
server	101	EX1	ge-0/0/5 untagged
			ge-1/0/3 untagged
		vQFX1	ge-0/0/3 untagged
printer	200	EX1	ge-0/0/9 untagged
phone	201	EX1	ge-0/0/7 tagged
			ge-1/0/9 tagged
guest	666	EX1	Dot1x guest vlan
vmx5vmx6ex1	156	EX1	ge-0/0/3
			ge-0/0/4

Table 20 VLANs Branch2 and Site 2

4) Configure the VLANs and corresponding interfaces in Branch2 and Site 2 as shown in the table above. Protect the vlans configuration hierarchy so that a warning message is displayed when deleted.
5) Configure interface ge-0/0/7 and interface ge-1/0/9 on virtual chassis EX1 for Voice VLAN support. Make sure that tagged frames go to VLAN201 whereas untagged frames go to VLAN100.
6) Configure aggregate ethernet interface (AE15) on vMX7 and vMX8 towards vQFX2. Enable LACP on the AE15 interface. All Site 2 VLANs must be allowed on the AE interfaces. Configure interface AE0 on vQFX2 towards vMX7 and vMX8. Interface AE0 on vQFX2 should not initiate LACP.
7) Ensure the following MAC addresses are in the Ethernet switching table of EX1 for vlan vmx5vmx6ex1:

00:00:5e:00:01:0d
00:00:5e:00:02:63
00:56:56:56:56:05
00:56:56:56:56:06

8) Ensure that only traffic received from the MAC address of server2 is allowed on EX1 port ge-1/0/3.

© 2020 Juniper Networks, Inc. All rights reserved.

9) Ensure Server1 and server 2 receive their IPv4 address from the DHCP server.
10) User3 must be authenticated dynamically and ensure that only the authenticated device is allowed to communicate to the network. If authentication fails or the server is unreachable, the user must be given guest access.
11) Configure interface ge-0/0/7 and interface ge-1/0/9 on EX1. No more than 20 WATT should be allocated the POE ports.
12) Enable Sticky MAC for EX1 interface ge-1/0/9.
13) Configure EX1 to block all data packets on ports if Unknown Unicast packets consume more than 10% of the interface bandwidth. Broadcast and multicast traffic should not be affected by your solution. Control-plane traffic should never be blocked.

Task 3: Protocol Independent Features

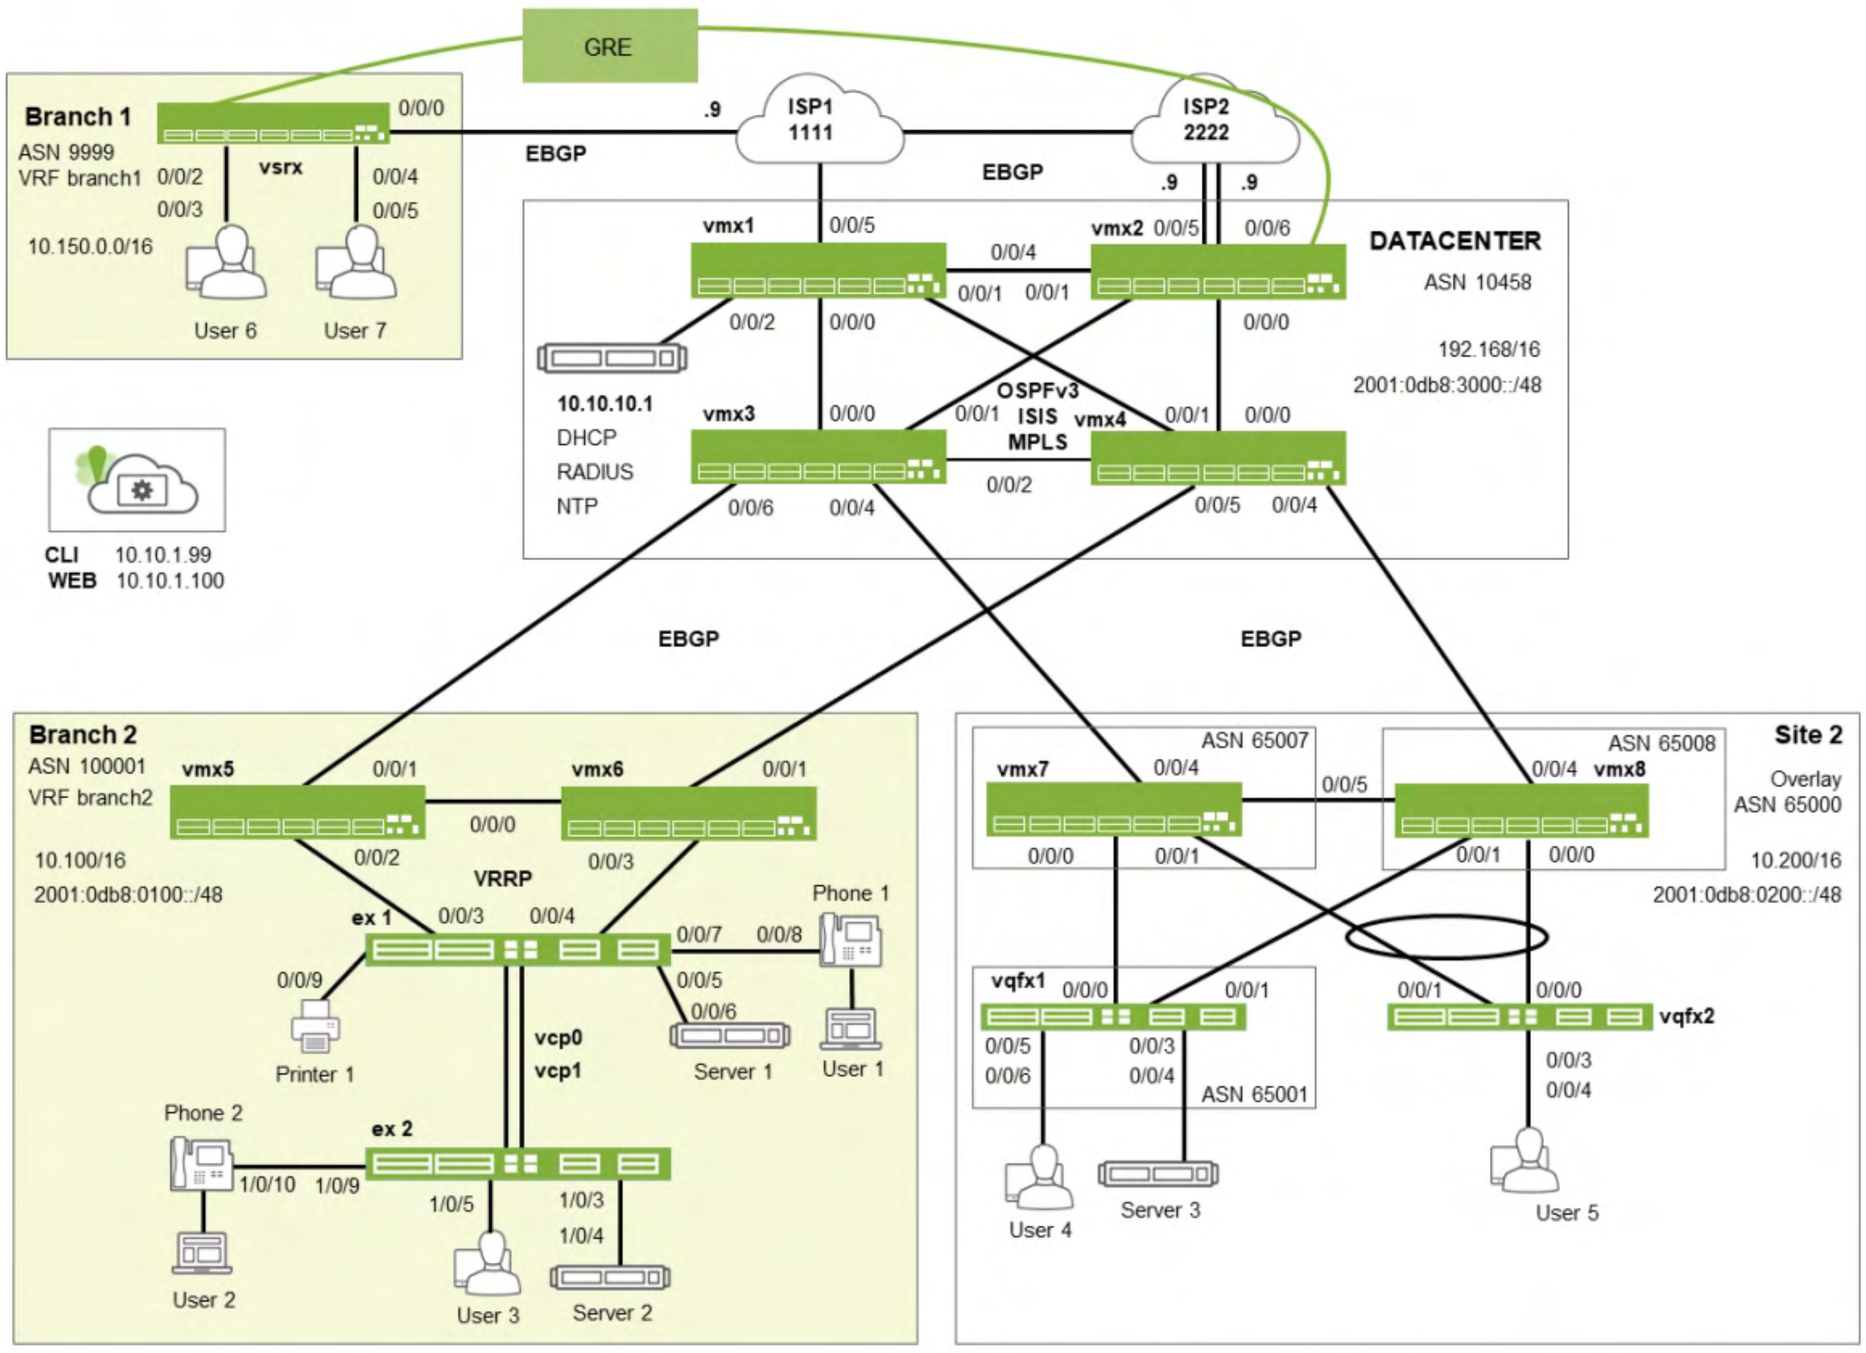

1) Configure a GRE tunnel between Branch 1 and vMX2 in the data center. Use the IP addressing as shown in the lab setup Information chapter. Use the loopback interface as GRE tunnel source address.

© 2020 Juniper Networks, Inc. All rights reserved.

2) vMX2 should forward ping reply packets received from USER3 to the 28.184.130.64/27 using interface ge-0/0/5. If interface ge-0/0/5 goes down, then these packets should be dropped.
3) vMX6 has a commit script stored that needs to be enabled. The commit script provides a warning message if vMX6 has IPv6 control plane filtering disabled. The commit script is currently not working and needs to be fixed.

Branch1:

1) Configure static routing over the ge-0/0/0 interface to ensure vMX5 and vMX6 can reach each other loopback address.
2) Create two static summary routes on vMX5 and vMX6 that encompass the EX1 IRB interfaces.
3) Configure VRRP on vMX5 and vMX6 to provide a redundant default gateway for EX1. Make sure that the virtual address responds to icmp packets.
4) vMX5 should be the VRRP master for IPv4 traffic. vMX6 should be the master for IPv6 traffic. After recovery of a failure a VRRP master should revert back to its role.
5) vMX5 should become VRRP master for IPv6 in case vMX6 interface ge-0/0/1 fails.
6) Create two static default routes on EX1 towards the VRRP virtual address (.254) of vMX5 and vMX6.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: IGP and MPLS

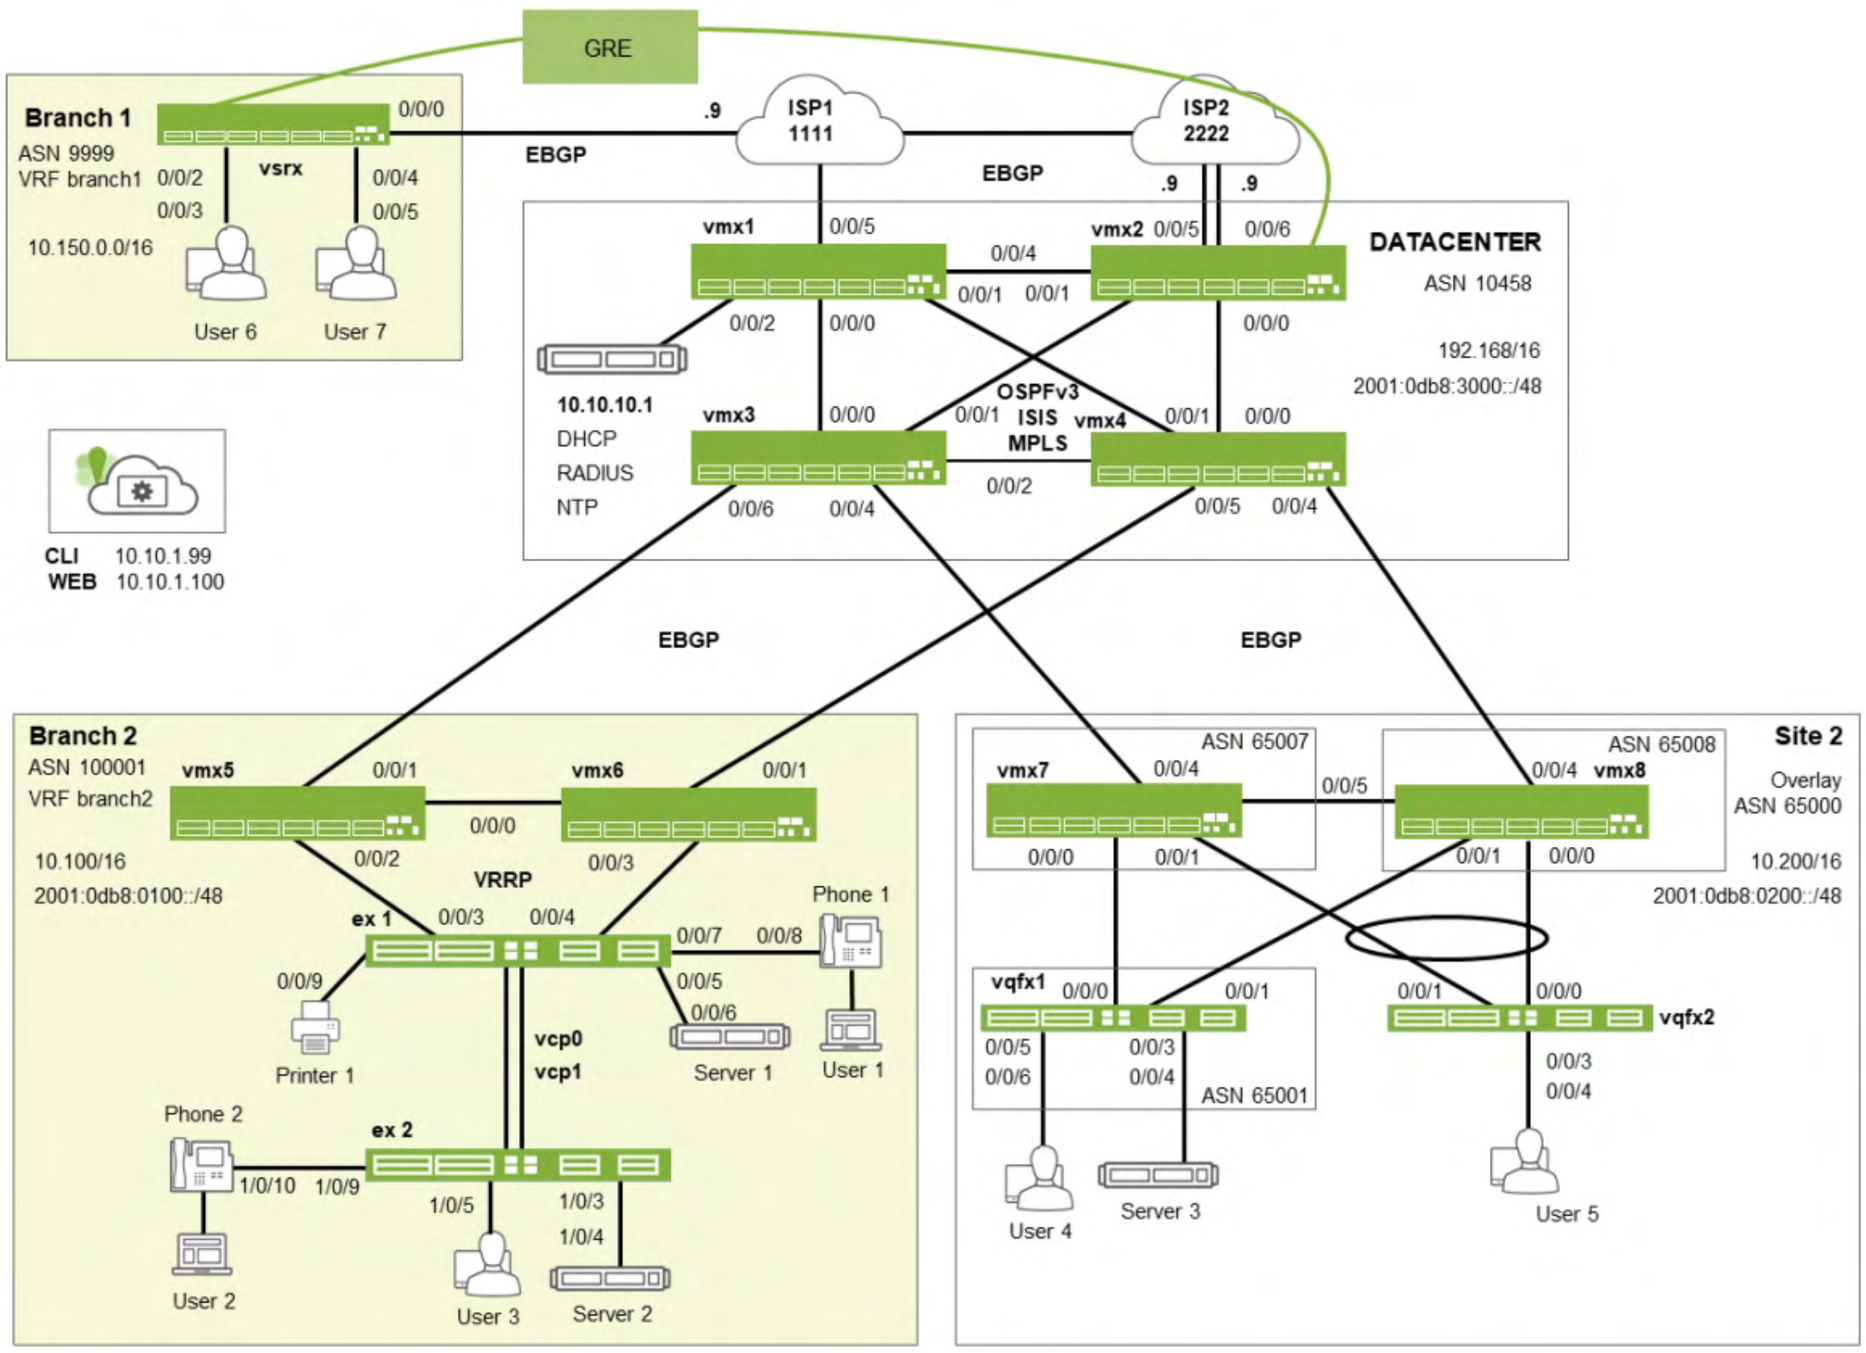

Data Center:

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX2	lo0	0.0.0.0
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX3	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1

© 2020 Juniper Networks, Inc. All rights reserved.

vMX4	ge-0/0/1	0.0.0.1
	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1

Table 21 Data Center OSPF Information

1) Configure OSPFv3 for IPv6 routing in the data center using the information as shown in the table above.
2) Ensure no LSA type 0x2002 are found in the OSPFv3 database.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four missed hello packets.
4) Ensure that each vMX devices in the Data Center learn each other's IPv6 loopback address.

OSPFv2 for IPv4 routing in the data center and over the GRE tunnel between the data center and the Branch office.

Device	Interface	Routing instance	area
vMX2	GRE interface	branch1	0.0.0.0
vSRX	GRE interface	default	0.0.0.0

Table 22 OSPF Information for GRE Tunnel

5) Enable OSPFv2 over the GRE interface as shown in the table above.
6) Ensure the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.
7) Ensure that vSRX OSPF database contains a default route and the IS-IS destinations received from vMX2. Configure OSPF on vSRX that the default route is the only external LSA installed in its routing table.
8) Configure IS-IS in the data center using the information as shown in the tables below.

Device	Level	Interfaces
vMX1	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
		ge-0/0/2
vMX2	2	lo0

© 2020 Juniper Networks, Inc. All rights reserved.

		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
vMX3	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2
vMX4	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2

Table 23 IS-IS Information

Device	Required IS-IS configuration
vMX1	Sysid: 0168.0000.0001 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled
vMX2	Sysid: 0168.0000.0002 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled

© 2020 Juniper Networks, Inc. All rights reserved.

vMX3	Sysid: 0168.0000.0003 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled
vMX4	Sysid: 0168.0000.0004 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled

9) Do not enable Level 1 routing.
10) Ensure no DIS is elected.
11) Ensure interface ge-0/0/2 on vMX1 participates in IS-IS, but no adjacency can be formed.
12) The maximum number of routes that can be redistributed in IS-IS should be 100.
13) Enable MD5 authentication for level 2. Use “juniper123” as the key.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 5: MPLS/LDP

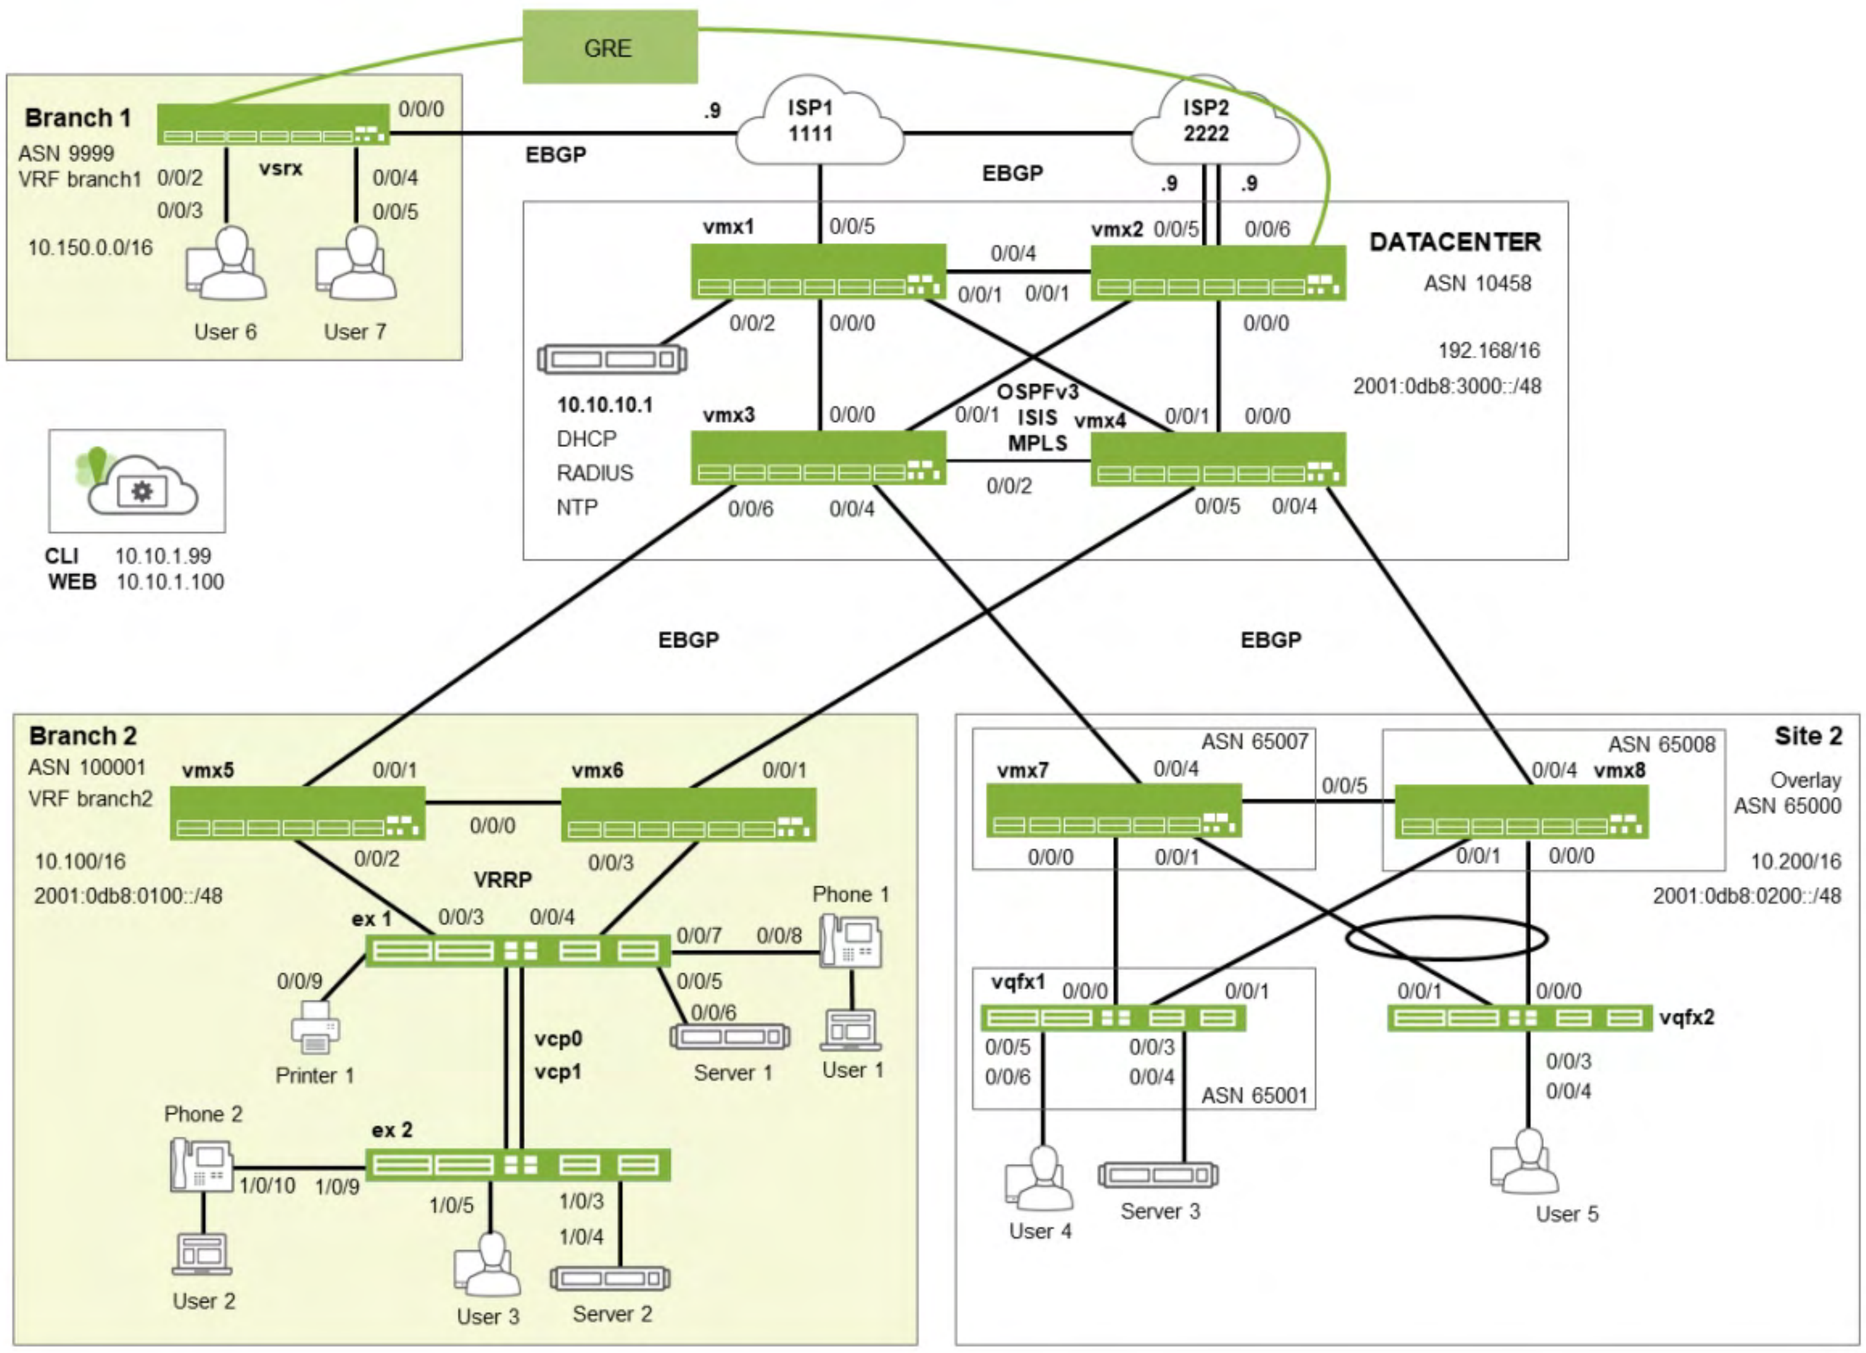

Data Center:

1) Enable MPLS forwarding between the data center vMX devices using the LDP protocol.
2) Penultimate Hop Popping (PHP) should be disabled.

Task 6: BGP

This chapter focuses on the implementation of the BGP. At various points in the topology you need to perform route redistribution between an IGP and BGP for end-to-end connectivity.

© 2020 Juniper Networks, Inc. All rights reserved.

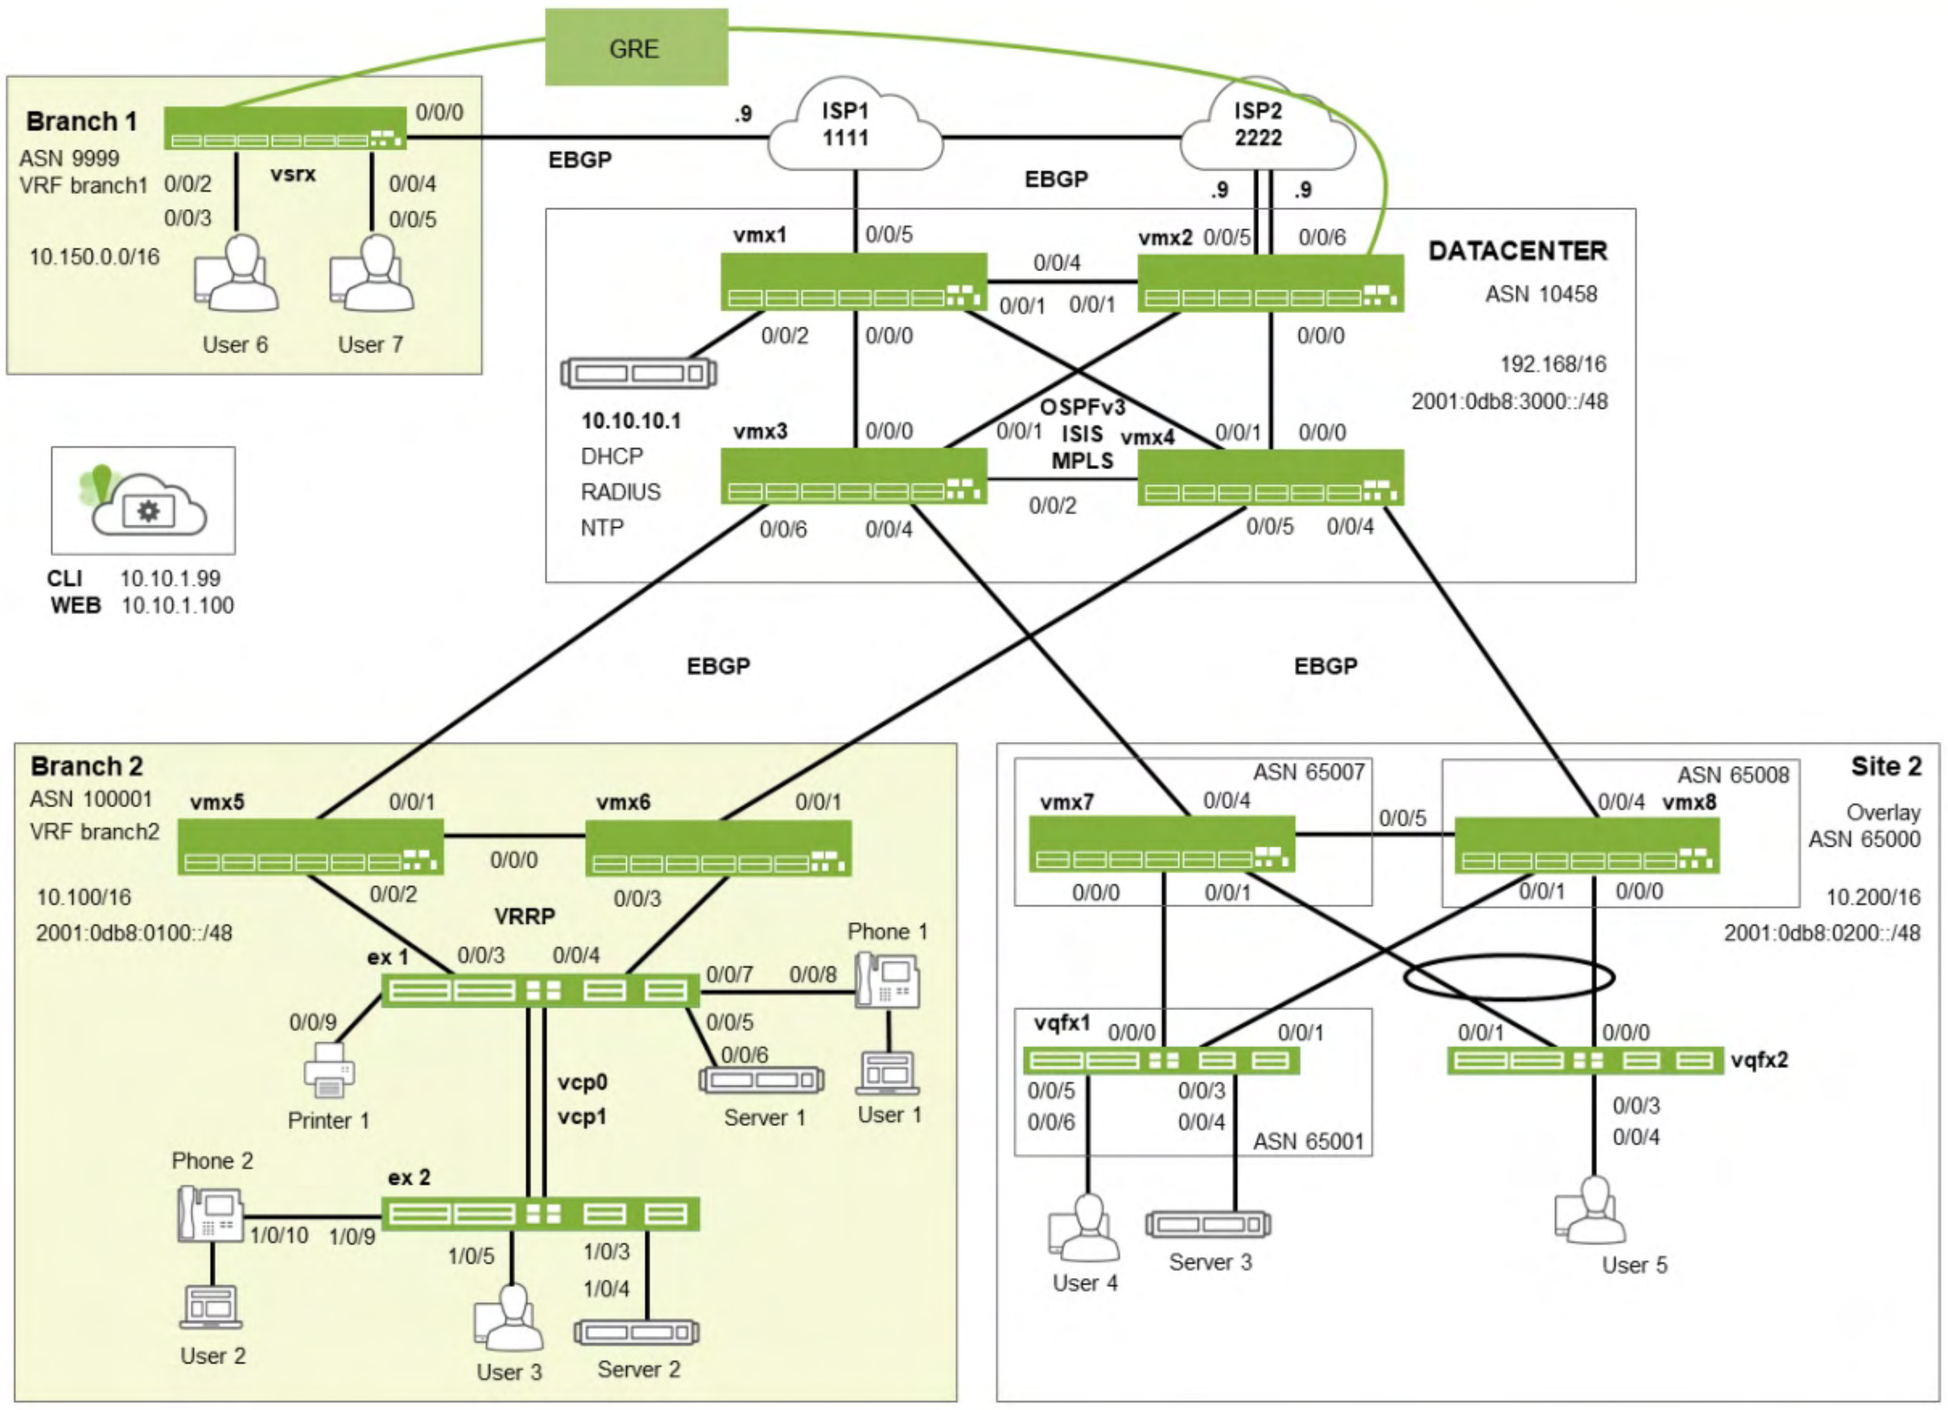

Device	Interface	Ipv4 address	Ipv6 address	Description
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128	IBGP source interface
				v4 EBGP to AS 2222
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64	v6 EBGP to AS 2222
	ge-0/0/6	172.16.3.2/24	Link local	v6 EBGP to AS 2222
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128	IBGP source interface
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64	EBGP to AS 64513
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64	EBGP to AS 100001
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128	IBGP source interface
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64	EBGP to AS 64513
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64	EBGP to AS 100001
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458

© 2020 Juniper Networks, Inc. All rights reserved.

vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table 24 BGP peering information

IBGP Data Center:

3) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have a IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and also accept routes received from each other.
4) All IBGP sessions on vMX1 and vMX2 should be configured without a neighbor statement.
5) Ensure IPv4, IPv6 and MPLS L3VPN NLRI is exchanged. Use separate peerings for IPv4 and IPv6 NLRI exchange.
6) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS.
7) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down. vMX2 should not advertise the default route if the BGP peering with ISP2 goes down.
8) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the data center AS.

IBGP Branch 2 and Site 2:

9) Configure IPv4 and IPv6 IBGP peerings in Branch 2 between vMX5 and vMX6. Peer using the loopback interfaces. The ASN is 100001.

EBGP Data Center:

10) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
11) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.

© 2020 Juniper Networks, Inc. All rights reserved.

12) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.
13) Advertise the IPv4 summary routes received from Branch2 and Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs.
14) Advertise a single /32 route to both ISPs that encompasses the IPv6 networks of the data center, Branch 2, and Site 2.
15) Configure vMX1 so that ISP2 is the primary path inbound to your AS.
16) Ensure the ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.
17) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2.
18) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Branch 2 and Site 2.
19) Site 2 should learn the IPv4 and IPv6 default route from the data center. Site 2 should learn the management network 10.10.10/24 that originated in the data center and the IPv4 summarized networks from Branch 2 and the data center.
20) Ensure the devices in Branch1 and Site2 have full reachability to each other and the Internet.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 7: VPN

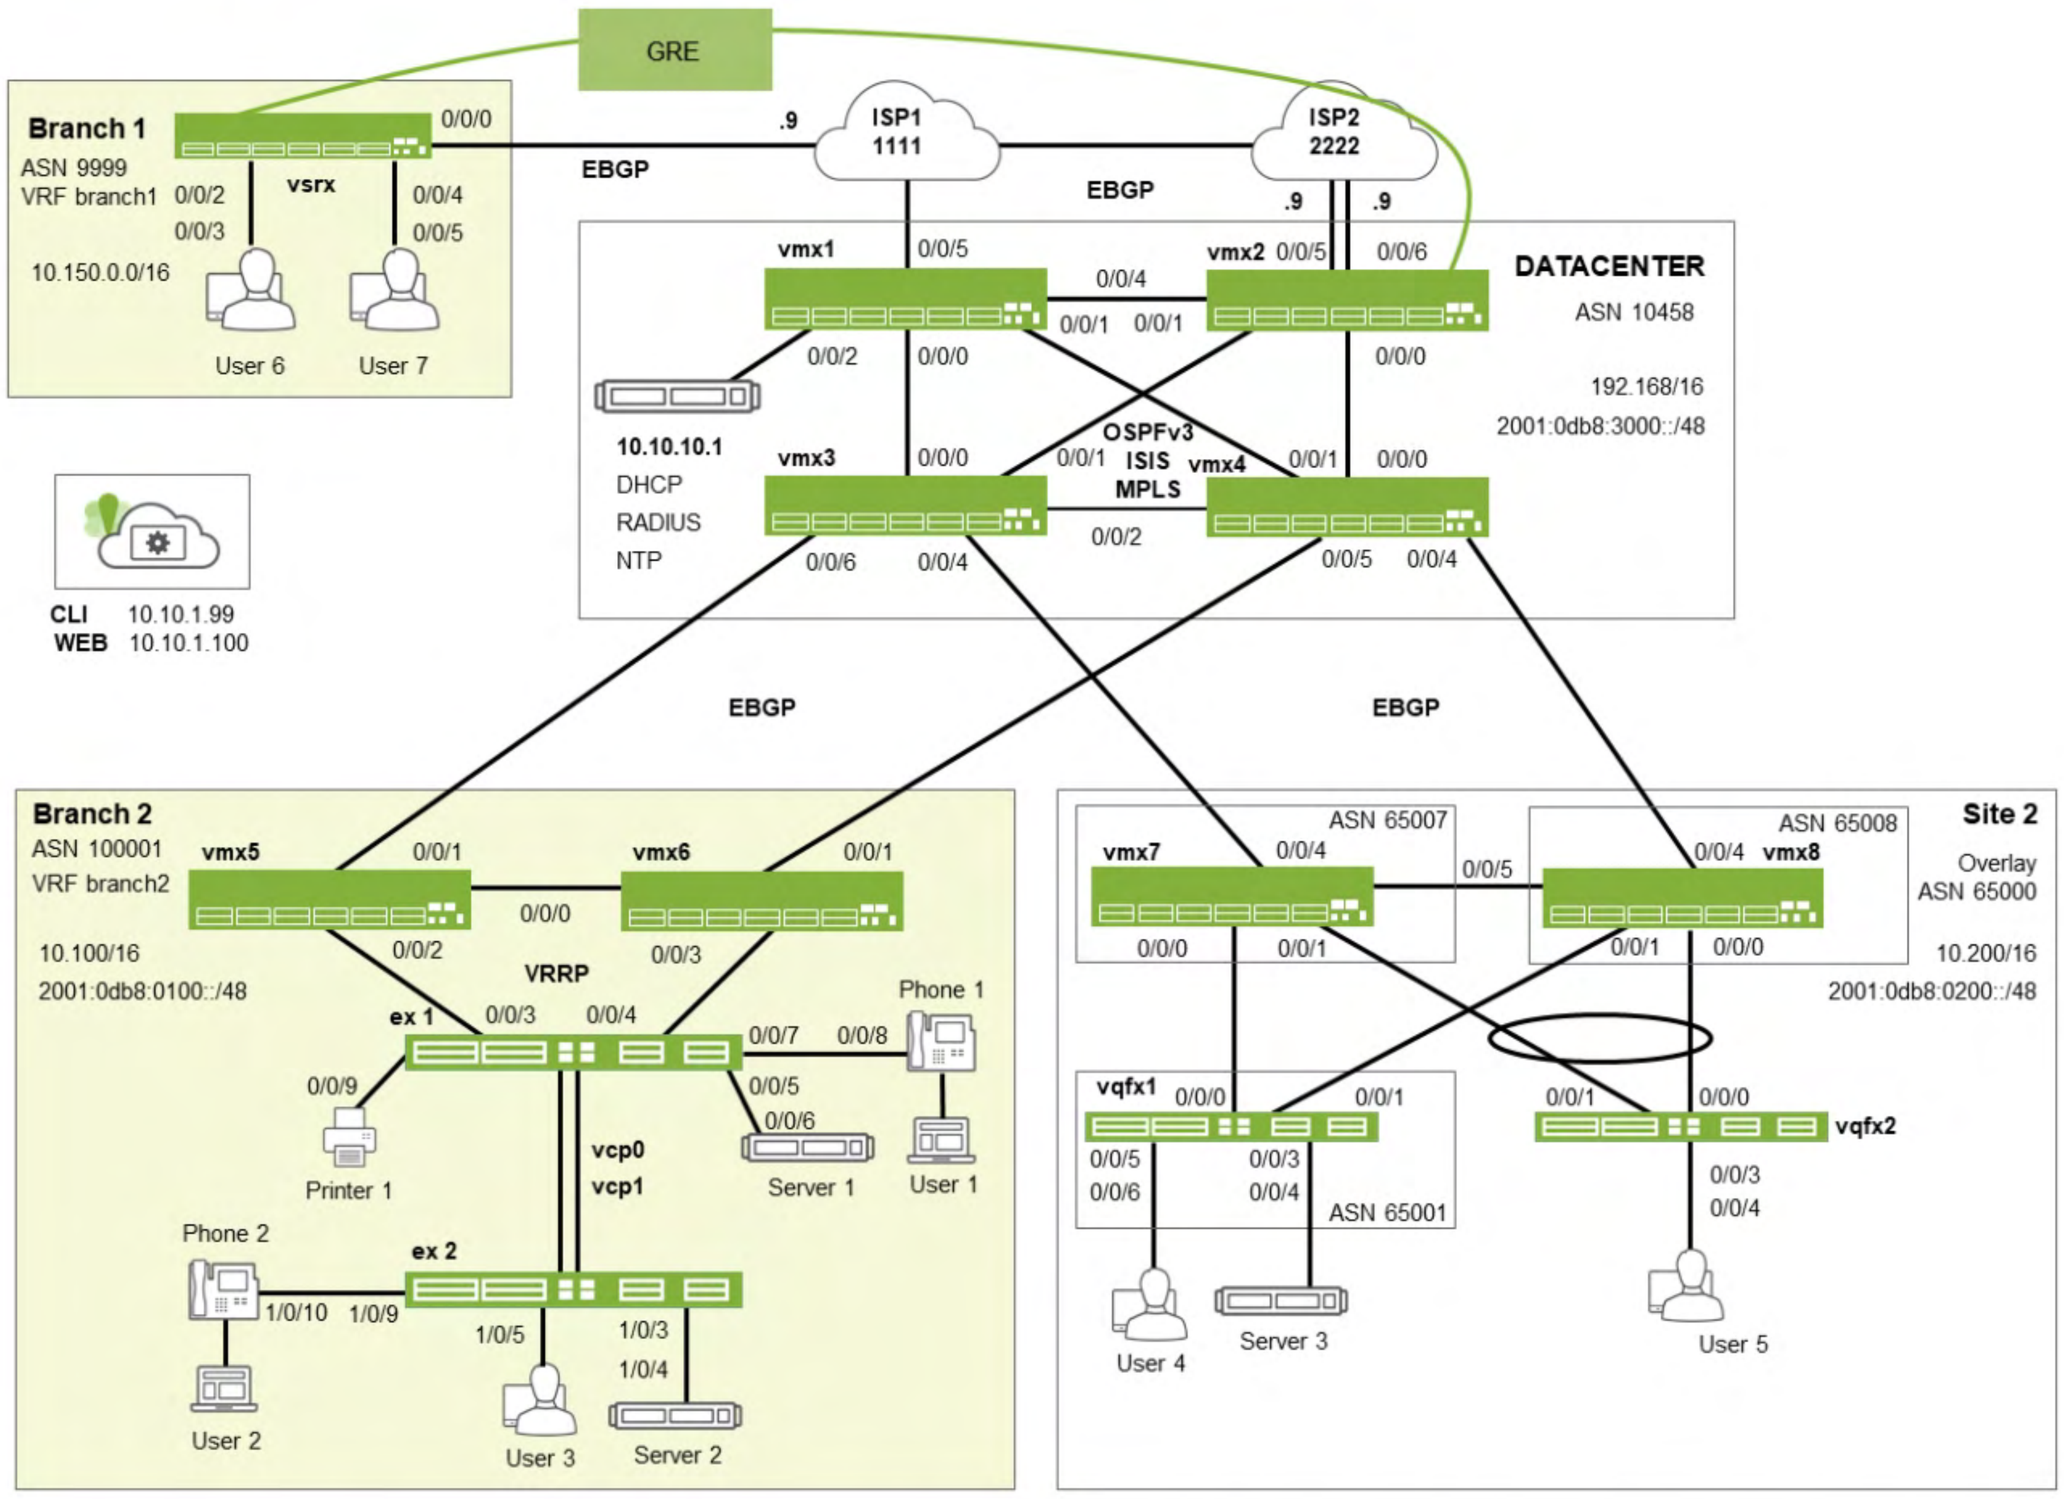

1) Configure a MPLS L3VPN on vMX2, vMX3 and vMX4 connecting Branch1 and Branch2 for IPv4 routing. Use the information as shown in the table below.

Router	Interface	Routing Instance	PE-CE Routing
vMX2	GRE tunnel	branch1	OSPFv2
vMX3	ge-0/0/6.0	branch2	EBGP
vMX4	ge-0/0/5.0	branch2	EBGP

2) Configure vMX2 to ensure Branch1 can access Site 2 and the Internet for IPv4 traffic using inet.0 routing table.
3) Configure vMX3 and vMX4 to ensure Branch 2 can access Site 2 and the Internet for IPv4 and IPv6 traffic using the inet6.0 routing table.

© 2020 Juniper Networks, Inc. All rights reserved.

4) Branch 1 should only learn the IPv4 and IPv6 default route from the data center.

Task 8: EVPN/VXLAN

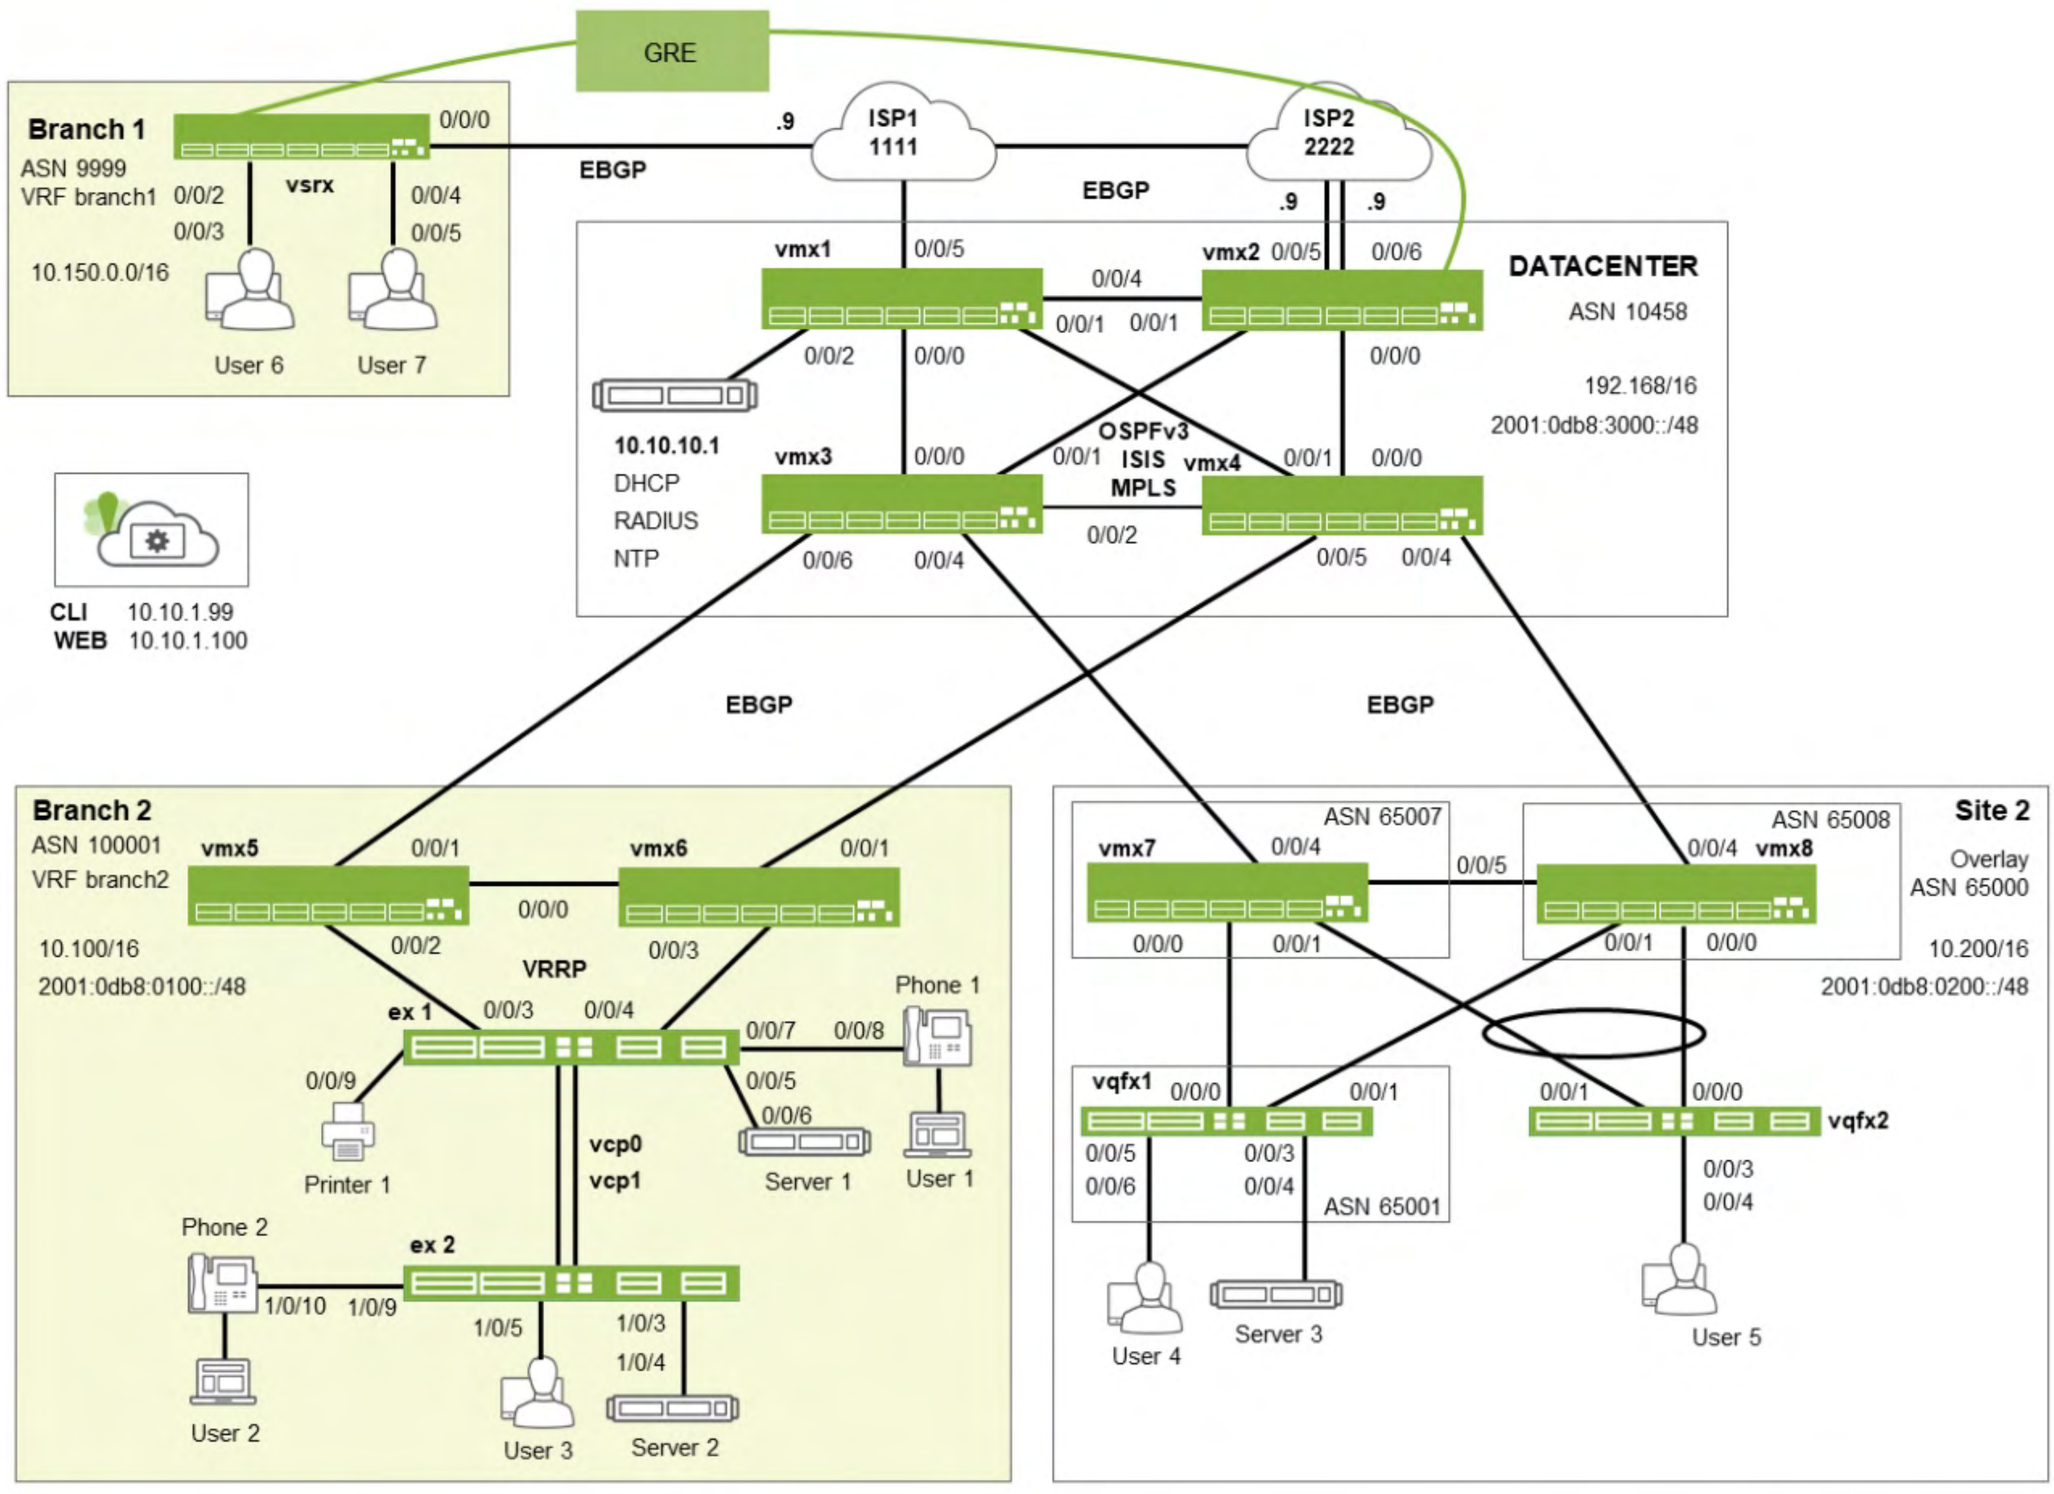

1) Configure the EBGP underlay in Site 2 to exchange device loopback reachability. Use the AS numbers as shown in the topology. Use the interface addresses to establish the EBGP peerings.
2) Configure the IBGP EVPN overlay for Site 2. vMX7 and vMX8 are router reflectors for the EVPN overlay.
3) Ensure that virtual gateway addresses respond to ICMP packets.
4) Configure one EVPN VLAN aware service on vMX7 and vMX8 for VLAN 100 and 101
5) Ensure the devices in Site 2 can reach each other, the devices in Branch2 and the Internet.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 9: Class of Service

In this Task you will configure virtual chassis EX1 in Branch 2 to support Class of Service.

1) Configure the following Forwarding Classes and map them to the output Queues as indicated in Table 9.

Forwarding Class	Queue	Schedulers L2 and L3 core Interfaces	Schedulers L2 access ports
control	7	control-sched-core	control-sched-user
voice	5	voice-sched-core	voice-sched-user
server-gold	3	server-gold-sched	server-gold-sched
server-silver	2	server-silver-sched	server-silver-sched
server-bronze	1	server-bronze-sched	server-bronze-sched
best-effort	0	best-effort-sched	best-effort-sched

Table 25

2) Configure schedulers with the parameters as described in Table 10 below.

Scheduler	Parameter	Value
control-sched-core	priority	strict-high
	shaping rate	5%
	buffer-size	10%
control-sched-user	priority	strict-high
	shaping rate	2%
	buffer-size	10%
voice-sched-core	priority	strict-high
	shaping rate	10%
	buffer-size	5% exact
voice-sched-user	priority	strict-high
	shaping rate	5%
	buffer-size	5% exact
server-gold-sched	priority	strict-high
	shaping rate	25%
	buffer-size	10% exact
server-silver-sched	priority	low
	transmit rate	35%
	buffer-size	50%
server-bronze-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder

© 2020 Juniper Networks, Inc. All rights reserved.

best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder

Table 26

3) Apply the schedulers to the physical interfaces on virtual chassis EX1. Make sure that OoB interfaces are not affected.
4) Configure IPv4, IPv6 and dot1p Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	Dot1p	Dot1p Loss-priority	DSCP value	DSCP Loss-priority
control	111	low	110000	low
			111000	low
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 27

5) Apply the classifiers to all interfaces as described in Table 12. Make sure that the OoB interfaces are not affected.

Virtual chassis	Classifiers	Interfaces
EX1	BA IPv4, IPv6	All L3 interfaces
	Dot1p	All L2 interfaces

Table 28

6) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	Dot1p	Dot1p Loss-priority	DSCP value	DSCP Loss-priority
control	111	low	110000	low
			111000	low
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low

© 2020 Juniper Networks, Inc. All rights reserved.

			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 29

7) Apply the rewrite-rules to all interfaces as described in the table below. Make sure that the OoB interfaces are not affected.

Virtual chassis	Rewrite-rules	Interfaces
EX1	IPv4, IPv6	All L3 interfaces
	Dot1p rewrite rules	All L2 interfaces

Table 30

© 2020 Juniper Networks, Inc. All rights reserved.

Appendices

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

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Additional Theory

MPLS Overview

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network header is analyzed, and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. The packet is then forwarded to the next provider switch in the LSP. This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is forwarded based on its IP routing information.

This topic describes:

MPLS Label-Switched Paths and MPLS Labels

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the MPLS label (32 bits). Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

The figure below shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

© 2020 Juniper Networks, Inc. All rights reserved.

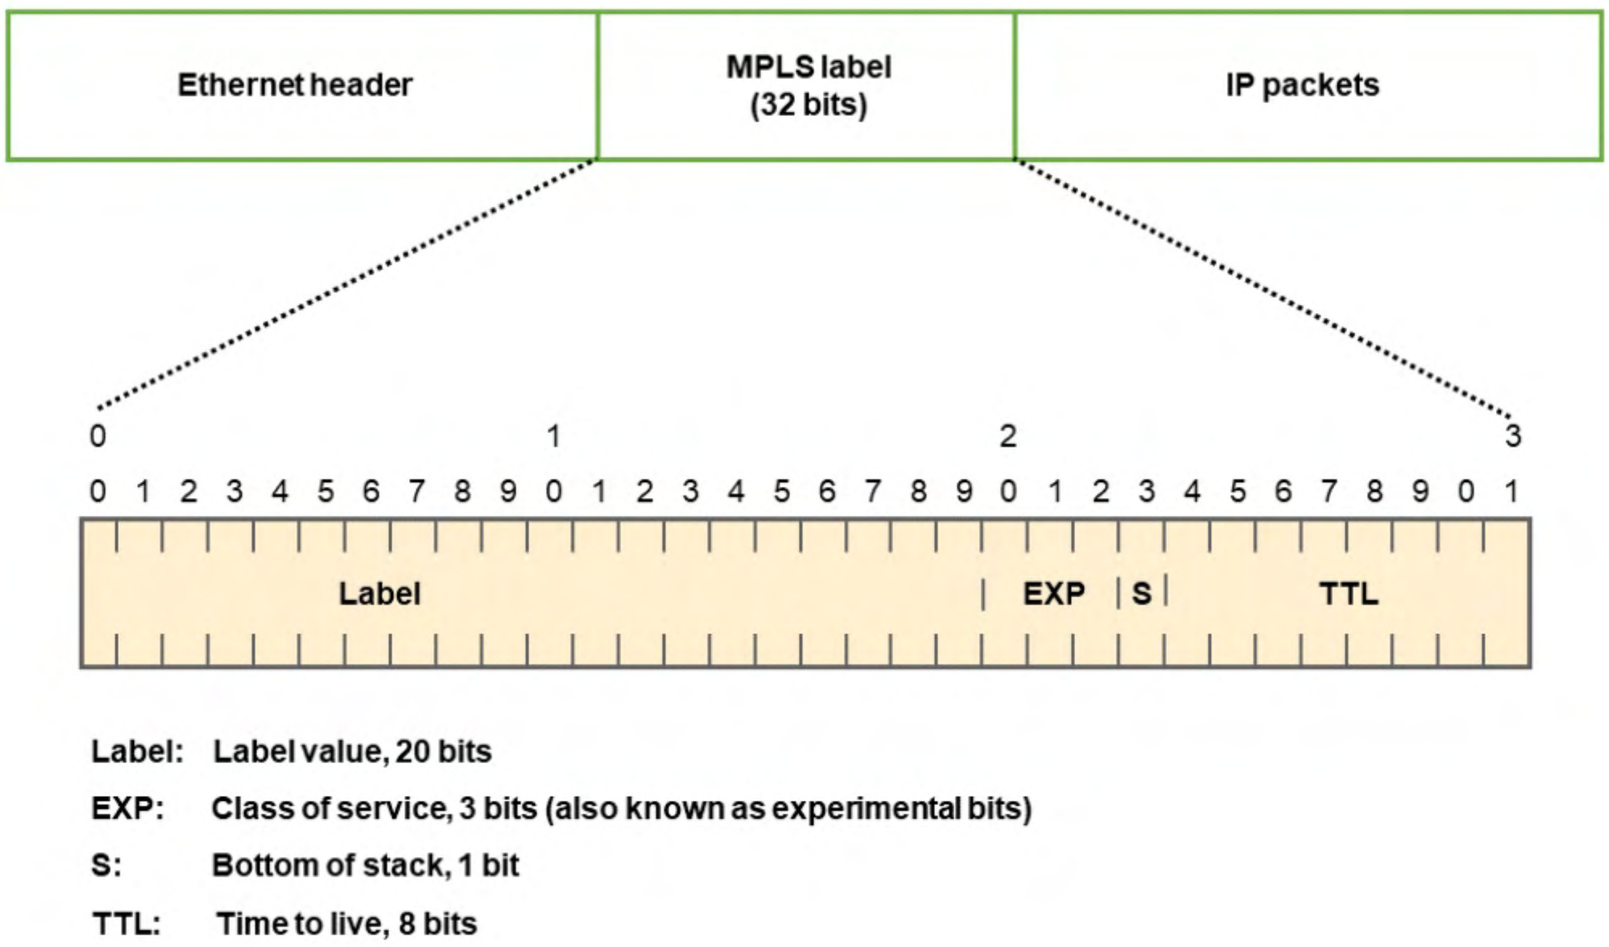

Label: Label value, 20 bits
EXP: Class of service, 3 bits (also known as experimental bits)
S: Bottom of stack, 1 bit
TTL: Time to live, 8 bits

Figure: Label Encoding

Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the range 0 through 15) have well-defined meanings. The following reserved labels are used by QFX Series and EX4600 devices:

0, IPv4 Explicit Null label—This value is valid only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 4 (IPv4) packet.

1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.

3, Implicit Null label—This label is used in the signaling protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

MPLS Label Operations:

Push
Pop

© 2020 Juniper Networks, Inc. All rights reserved.

Swap

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet.

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

Incoming interface
Label
Label forwarding table

The figure below shows an IP packet without a label arriving on the customer edge interface (ge-0/0/1) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination as the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (ge-0/0/5). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface ge-0/0/5 with label 100. The provider switch swaps label 100 with label 200 and forwards the MPLS packet through its core interface (ge-0/0/7) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (ge-0/0/7), removes the MPLS label, and sends the IP packet out of its customer edge interface (ge-0/0/1) to a destination that is beyond the tunnel.

MPLS Label Swapping

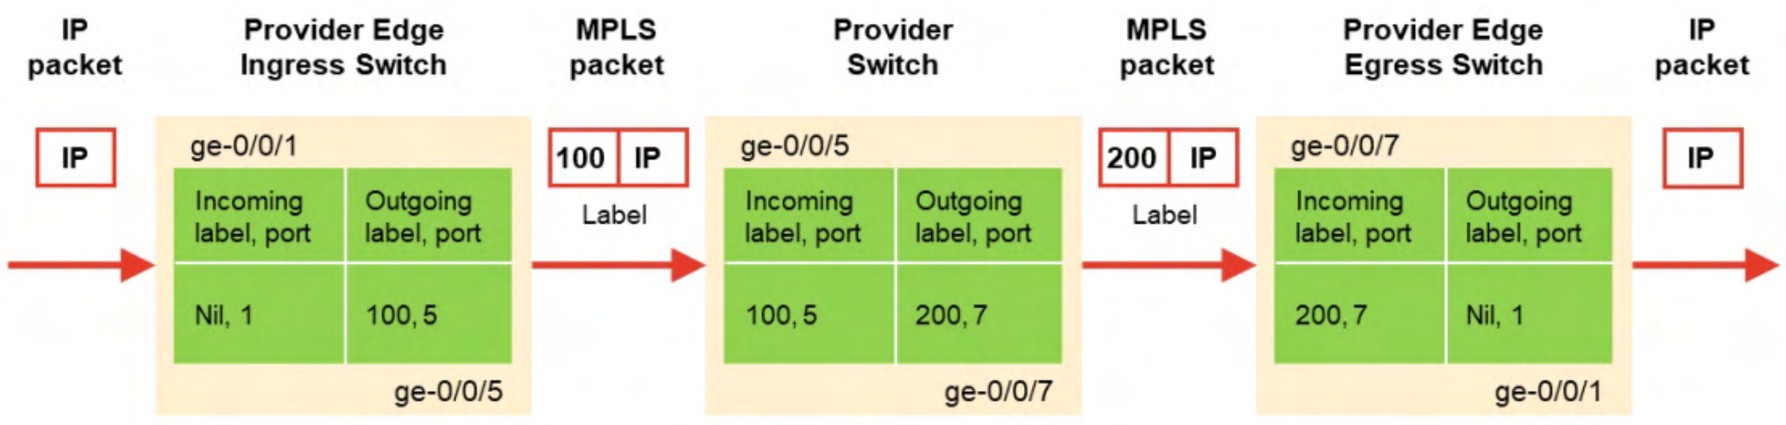

The figure above shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to

© 2020 Juniper Networks, Inc. All rights reserved.

travel in the reverse direction. Thus, each PE switch operates as both an ingress switch and an egress switch.

Penultimate-Hop Popping and Ultimate-Hop Popping

The switches enable penultimate-hop popping (PHP) by default with IP over MPLS configurations. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping the MPLS label.

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop switch removes the label and sends the packet to the egress PE switch.

If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised and the egress PE switch of the LSP removes the label.

Layer 3 VPN Overview

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. Instead of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. VPNs are a cost-effective alternative to expensive dedicated lines. The type of VPN is determined by the connections it uses and whether the customer network or the provider network performs the virtual tunneling.

You can configure a router running Junos OS to participate in several types of VPNs. This topic discusses MPLS VPNs.

This topic contains the following sections:

MPLS VPN Topology

There are many ways to set up an MPLS VPN and direct traffic through it. The figure below shows a typical MPLS VPN topology.

© 2020 Juniper Networks, Inc. All rights reserved.

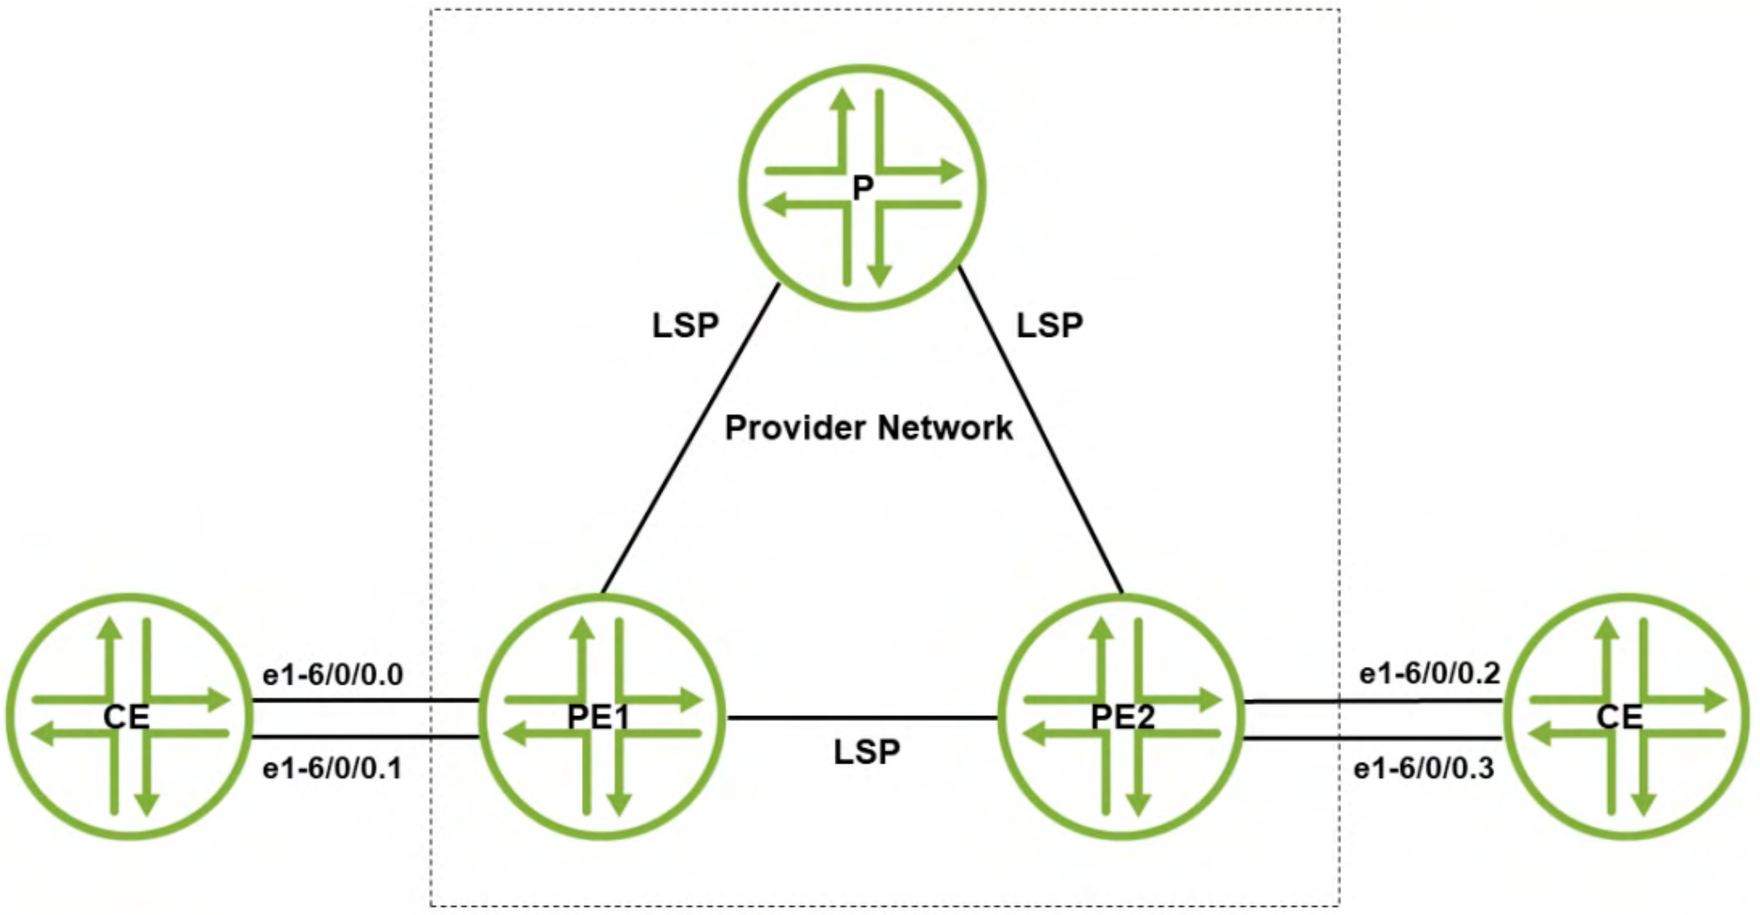

Figure 1: Typical VPN Topology

There are three primary types of MPLS VPNs: Layer 2 VPNs (such as EVPN and VPLS), Layer 2 circuits, and Layer 3 VPNs. All types of MPLS VPNs share certain components:

- The provider edge (PE) routers in the provider's network connect to the customer edge (CE) routers located at customer sites. PE routers support VPN and MPLS label functionality. Within a single VPN, pairs of PE routers are connected through a virtual tunnel, typically a label-switched path (LSP).
- Provider routers within the core of the provider's network are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support LSP functionality as part of the tunnel support, but do not support VPN functionality.
- CE routers are the routers or switches located at the customer site that connect to the provider's network. CE routers are typically IP routers, but they can also be Asynchronous Transfer Mode (ATM), Frame Relay, or Ethernet switches.

All VPN functions are performed by the PE routers. Neither CE routers nor provider routers are required to perform any VPN functions.

MPLS VPN Routing

VPNs tunnel traffic as follows from one customer site to another customer site, using a public network as a transit network, when certain requirements are met:

© 2020 Juniper Networks, Inc. All rights reserved.

1) Traffic is forwarded by standard IP forwarding from the CE routers to the PE routers.
2) The PE routers establish an LSP through the provider network.
3) The inbound PE router receives traffic, and it performs a route lookup. The lookup yields an LSP next hop, and the traffic is forwarded along the LSP.
4) The traffic reaches the outbound PE router, and the PE router pops the MPLS label and forwards the traffic with standard IP routing.

VRF Instances

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of MPLS VPNs, each VPN has a VPN routing and forwarding (VRF) instance.

A VRF instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation.

A separate VRF table is created for each VPN that has a connection to a CE router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF instance, and with routes received from other PE routers in the same VPN.

Route Distinguishers

Because a typical transit network is configured to handle more than one VPN, the provider routers are likely to have multiple VRF instances configured. As a result, depending on the origin of the traffic and any filtering rules applied to the traffic, the BGP routing tables can contain multiple routes for a particular destination address. Because BGP requires that exactly one BGP route per destination be imported into the forwarding table, BGP must have a way to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs.

A route distinguisher is a locally unique number that identifies all route information for a particular VPN. Unique numeric identifiers allow BGP to distinguish between routes that are otherwise identical.

Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

- *as-number*:*number*, where *as-number* is an autonomous system (AS) number (a 2–byte value) in the range 1 through 65,535, and *number* is any 4–byte value. We

© 2020 Juniper Networks, Inc. All rights reserved.

recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.

- *ip-address*:*number*, where *ip-address* is an IP address (a 4–byte value) and *number* is any 2–byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the router-id statement, which is a public IP address in your assigned prefix range.

The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VRF table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes.

The PE router then exports the route in IBGP sessions to the other provider routers. Route export is governed by any routing policy that has been applied to the particular VRF table. To propagate the routes through the provider network, the PE router must also convert the route to VPN format, which includes the route distinguisher.

When the outbound PE router receives the route, it strips off the route distinguisher and advertises the route to the connected CE router, typically through standard BGP IPv4 route advertisements.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Chapter 1: Junos System Features

This chapter focuses on initial system configuration and general system features. You will configure various tasks, such as configuring user authentication and authorization, NTP, SNMP, Syslog and RE protection.

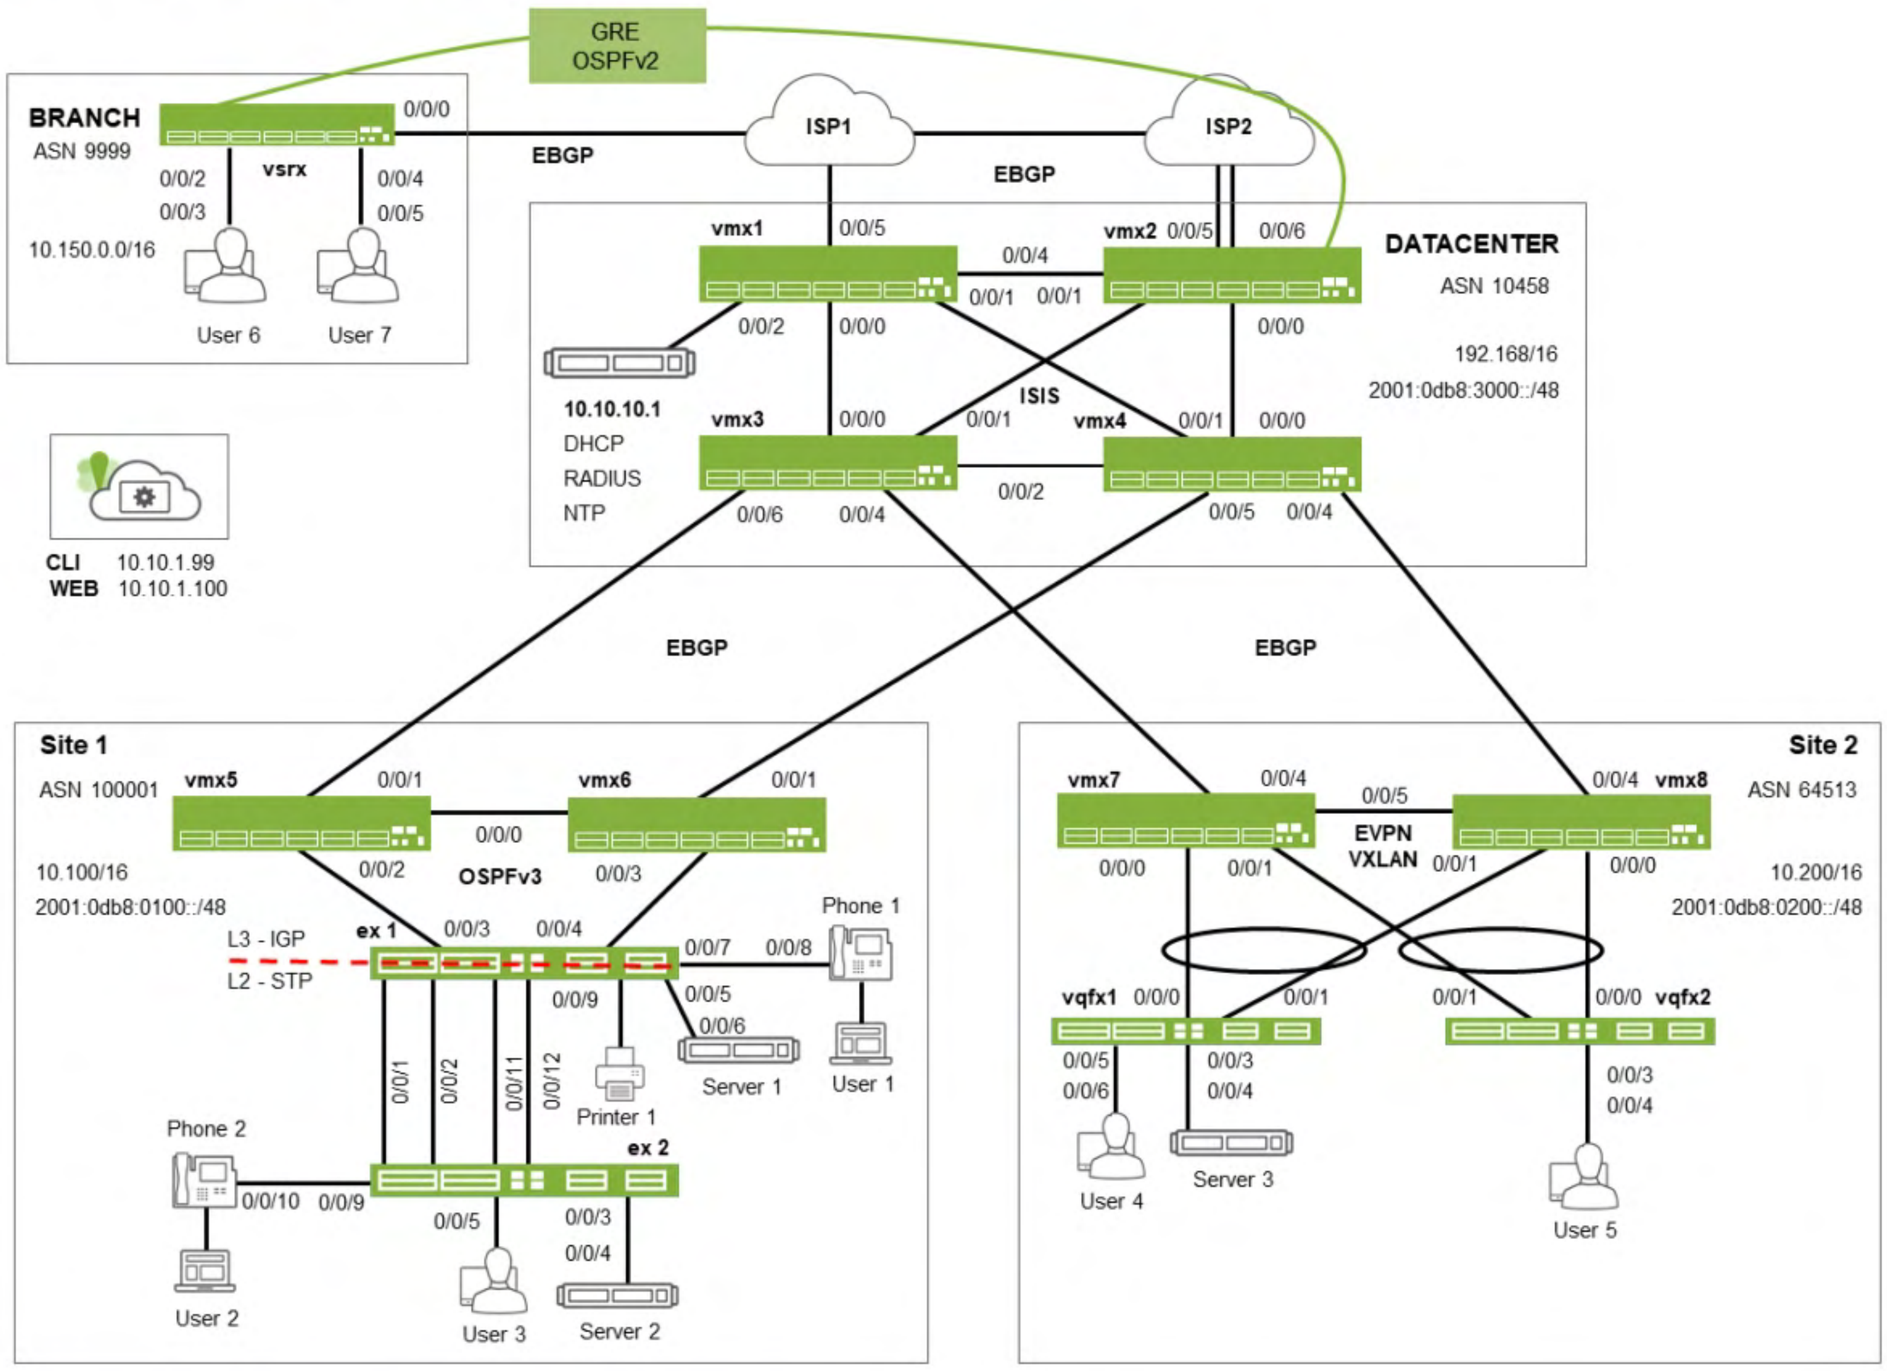

Task 1: General System Services

Note: For many of these tasks you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place. You will configure this starting in Chapter 4.

Note: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1*

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure all vMX and EX devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute.

Configure the following on all vMX and EX devices:

```
[edit system services]
lab@vmx1# show
ssh {
    root-login allow;
    connection-limit 4;
    rate-limit 10;
}
telnet {
    connection-limit 4;
    rate-limit 10;
}
```

2) Configure server 10.10.10.1 as the DNS server.

Configure the following on all vMX and EX devices:

```
[edit system]
lab@vmx1# show
name-server {
    10.10.10.1;
}
...
```

3) Set the time zone to *America/Los_Angeles* on all vMX devices and EX1

Configure the following on all vMX devices and EX1:

```
[edit system]
lab@vmx1# show
time-zone America/Los_Angeles;
...
```

4) Make sure that all vMX devices, and EX1 synchronize the time with NTP server 10.10.10.1. Communication should be sourced from the device loopback.

Configure on all vMX devices and EX1:

```
[edit system ntp]
lab@vmx1# show
server 10.10.10.1;
source-address 192.168.0.1;
```


© 2020 Juniper Networks, Inc. All rights reserved.

5) Configure all vMX devices and EX1 transfer their configuration to the FTP server 10.10.10.1 each time the configuration is committed. You can use username *lab* and password *lab123* for the FTP server access.

Configure the following on all vMX devices and EX1:

```
[edit system]
lab@vmx1# show
archival {
    configuration {
        transfer-on-commit;
        archive-sites {
            "ftp://lab@10.10.10.1" password "$9$y8HrvLx7VsYoLXUjq.5T"; ## SECRET-DATA
        }
    }
}
```

Configurations for Task Completion:

vMX1:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.1
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX2:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.2
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX3:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.3
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX4:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.4
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX5:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.5
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX6:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.6
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX7:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.7
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

vMX8:

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.8
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

EX1

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.1
set system archival configuration transfer-on-commit
set system archival configuration archive-sites "ftp://lab@10.10.10.1" password
"$9$y8HrvLx7VsYoLXUjq.5T"
```

EX2

```
set system services ssh connection-limit 4 rate-limit 10
set system services telnet connection-limit 4 rate-limit 10
set system name-server 10.10.10.1
```

Task 2: User Authentication and Authorization

In this part you will configure new users allowed to access the devices and define their privileges and permissions.

The following tasks should be configured on all vMX devices and EX1:

1) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts

© 2020 Juniper Networks, Inc. All rights reserved.

3 and timeout of 3 seconds. Use juniper123 as the RADIUS shared secret. Use the local loopback IP address as your source address.

Configure on all vMX devices and EX1:

```
[edit system]
lab@vmx1# show
...
authentication-order [ radius password ];
...
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.1;
    }
}
...
```

2) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices, EX1 as a super-user.

Configure the following on all vMX devices and EX1:

```
lab@vmx1# show system login
user admin {
    uid 2003;
    class super-user;
}
...
```

3) Configure additional users on all the devices as defined in Table 3.

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

Username	Password	Privileges
noc	noc123	Permissions “all”. Additionally, the user cannot execute any of the “clear”, “configure”, “edit” or “start shell” commands.
ops	ops123	Can view “show interfaces” with all options and can execute “ping” and “traceroute” commands.

Table 3

Configure the following on all vMX devices and EX1:

```
[edit]
lab@vmx1# show system login
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
class limited {
    permissions network;
    allow-commands "show interfaces *";
    deny-commands "(telnet)|(ssh)";
}
class privileged {
    permissions all;
    deny-commands "(clear)|(configure)|(edit)|(start shell)";
}
...
user noc {
    uid 2000;
    class privileged;
    authentication {
        encrypted-password "$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"; ## SECRET-DATA
    }
}
user ops {
    uid 2002;
    class limited;
    authentication {
        encrypted-password "$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"; ## SECRET-DATA
    }
}
```

4) Make sure that for all users in Table 3 sessions are automatically logout after 5 minutes of idle time-out. Ensure that new user classes will automatically inherit this setting.

Configure the following on all vMX devices and EX1:

```
[edit]
lab@vmx1# show groups
class-logout {
    system {
        login {
            class <*> {
                idle-timeout 5;
            }
        }
    }
}

[edit]
lab@vmx1# show system login
apply-groups class-logout;
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Configurations for Task Completion:

vMX1:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.1
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX2:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.2
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX3:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.3
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX4:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.4
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX5:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.5
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX6:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.6
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX7:

```
set system authentication-order [radius password]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.7
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

vMX8:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.8
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

EX1:

```
set system authentication-order [radius password]
set system radius-server 10.10.10.1 secret juniper123
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system radius-server 10.10.10.1 timeout 3 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.1
set system login user admin uid 2003 class super-user
set system login class limited permissions network
set system login class limited allow-commands "show interfaces *"
set system login class limited deny-commands "(telnet)|(ssh)"
set system login class privileged permissions all
set system login class privileged deny-commands "(clear)|(configure)|(edit)|(start
shell)"
set system login user noc uid 2000
set system login user noc class privileged
set system login user noc authentication encrypted-password
"$1$nf5Qzj6o$Fn5FUAMpN.dJ9YOBs8hUQ1"
set system login user ops uid 2002
set system login user ops class limited
set system login user ops authentication encrypted-password
"$1$ClUahiKL$xc4EecO0NzZuYaSgM6.ii/"
set groups class-logout system login class <*> idle-timeout 5
set system login apply-groups class-logout
```

Task 3: Syslog Configuration

Ensure that all vMX devices and EX1 have the following Syslog configuration:

1) All messages regardless of the facility with severity level of “info” and higher are sent to the default syslog file. Make sure that the repeated messages are individually logged.
2) A file named “interactive-commands” for command audit tracking receives records about the users and commands they execute.
3) All messages with severity level “info” and higher are sent to the syslog server 10.10.10.1. Log messages should be in a structured format.
4) Ensure that the maximum archived log files are 10. Each logfile can contain up to 5 Mbyte of log information.

Note: Ensure that messages sent to a remote syslog server are sourced from the local loopback address of each device.

Configure on all vMX devices and EX1:

```
[edit system syslog]
lab@vmx1# show
archive size 5m files 10 world-readable;
user * {
    any emergency;
}
host 10.10.10.1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    any info;
    log-prefix vmx1;
}
file messages {
    any notice;
    authorization info;
}
file interactive-commands {
    interactive-commands any;
}
time-format year millisecond;
source-address 192.168.0.1;
```

Configurations for Task Completion.

vMX1:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx1
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 192.168.0.1
```

vMX2:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx2
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 192.168.0.2
```

vMX3:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx3
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 192.168.0.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX4:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx4
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 192.168.0.4
```

vMX5:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx5
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 10.100.0.5
```

vMX6:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx6
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 10.100.0.6
```

vMX7:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix vmx7
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 10.200.0.7
```

vMX8:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system syslog host 10.10.10.1 log-prefix vmx8
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 10.200.0.8
```

EX1:

```
set system syslog archive size 5m files 10 world-readable
set system syslog host 10.10.10.1 any info
set system syslog host 10.10.10.1 log-prefix ex1
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set system syslog time-format year millisecond
set system syslog source-address 10.100.0.1
```

Task 4: SNMP Configuration

In this task you will configure SNMPv2.

The following tasks should be configured on all vMX, vQFX, and vSRX devices in the topology:

1) The SNMP read-only community string is jncie. Make sure that only the Junos Space server can poll the devices in the data center, Site1, Site2, and the branch device.
2) Configure all devices to send SNMP traps to Junos Space. in case of link and/or chassis events. Make sure that traps are only sent using the OoB management interface.

Configure on all vMX, vQFX and vSRX devices:

```
[edit]
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.1;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

Configurations for Task Completion.

vMX1:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.1
set snmp trap-group traps targets 10.10.1.99
```

vMX2:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.2
set snmp trap-group traps targets 10.10.1.99
```

vMX3:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.3
set snmp trap-group traps targets 10.10.1.99
```

vMX4:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.4
set snmp trap-group traps targets 10.10.1.99
```

vMX5:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.5
set snmp trap-group traps targets 10.10.1.99
```

vMX6:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.6
set snmp trap-group traps targets 10.10.1.99
```

vMX7:

```
set snmp client-list list 10.10.1.99/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.7
set snmp trap-group traps targets 10.10.1.99
```

vMX8:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.8
set snmp trap-group traps targets 10.10.1.99
```

vSRX:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.10
set snmp trap-group traps targets 10.10.1.99
```

vQFX1:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.13
set snmp trap-group traps targets 10.10.1.99
```

vQFX2:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.14
set snmp trap-group traps targets 10.10.1.99
```

Task 5: Junos Space

In this task, you will configure the Junos Space Network Director.

1) You have a Junos Space Network Director instance running for device management. You must discover all vMX devices using Junos Space Network Director with SNMPv2 and the community string jncie. Ensure that the devices are present under the Inventory window.

Note: Junos Space can be accessed by browsing to the management address using HTTPS from your student desktop.

Log in to Junos Space with username: super and password: lab123

© 2020 Juniper Networks, Inc. All rights reserved.

In the left column, navigate to “Devices / Device Discovery” and select “Device Discovery Profiles”.

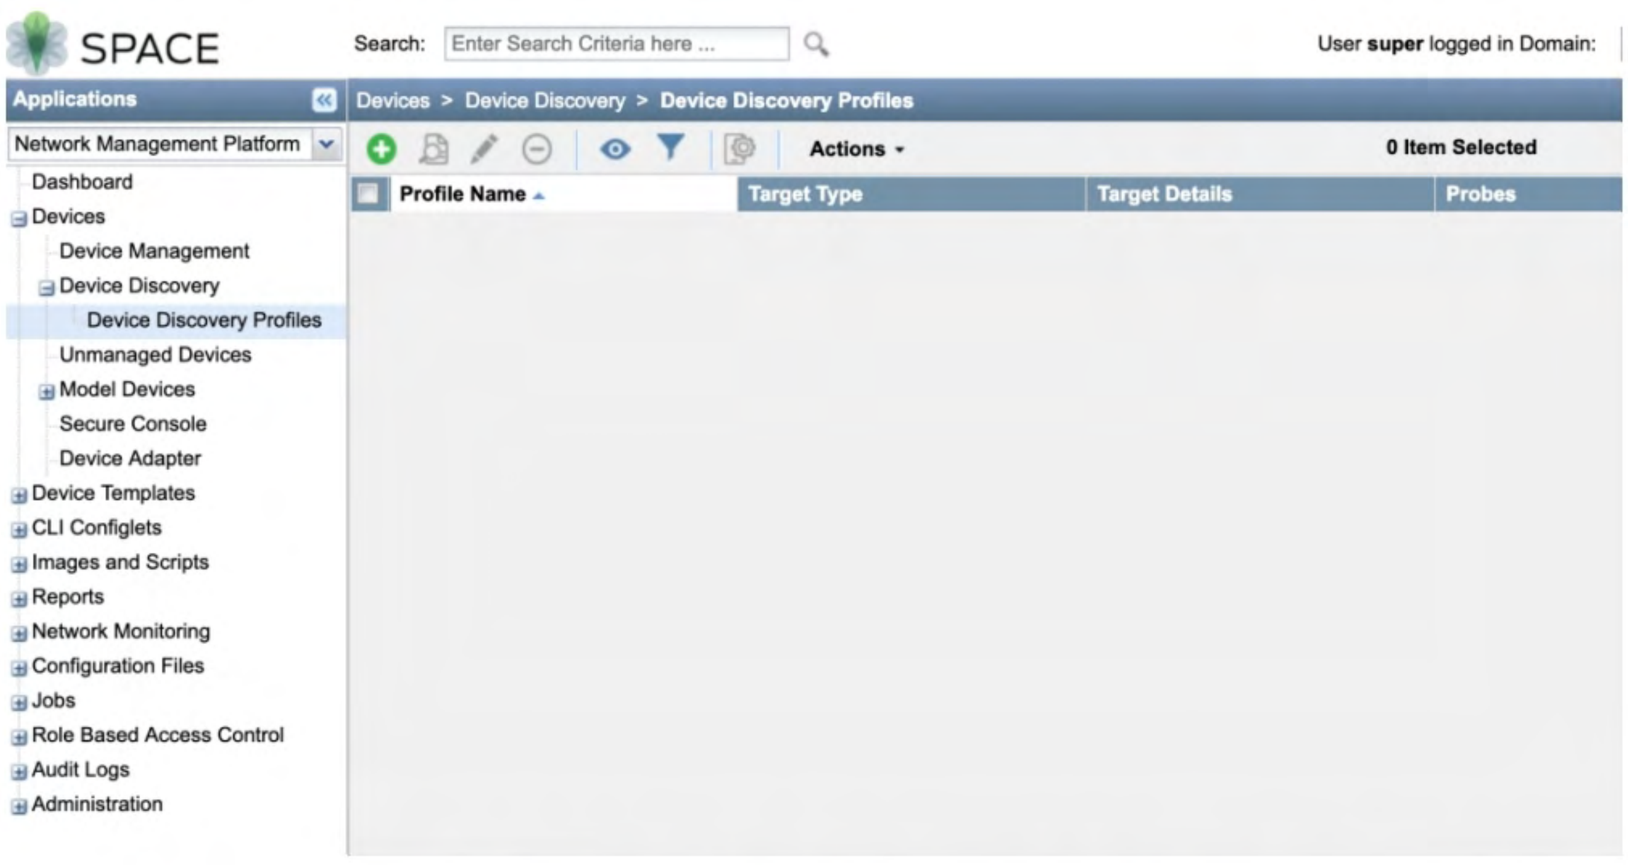

Specify a Device Profile Name. In this case we used the name “mydevices”. The task requirement is that all vMXs in the topology should be discovered. The IPv4 management/ OOB range of the vMXs is 10.10.1.1 to 10.1.1.8. Select “Next”.

© 2020 Juniper Networks, Inc. All rights reserved.

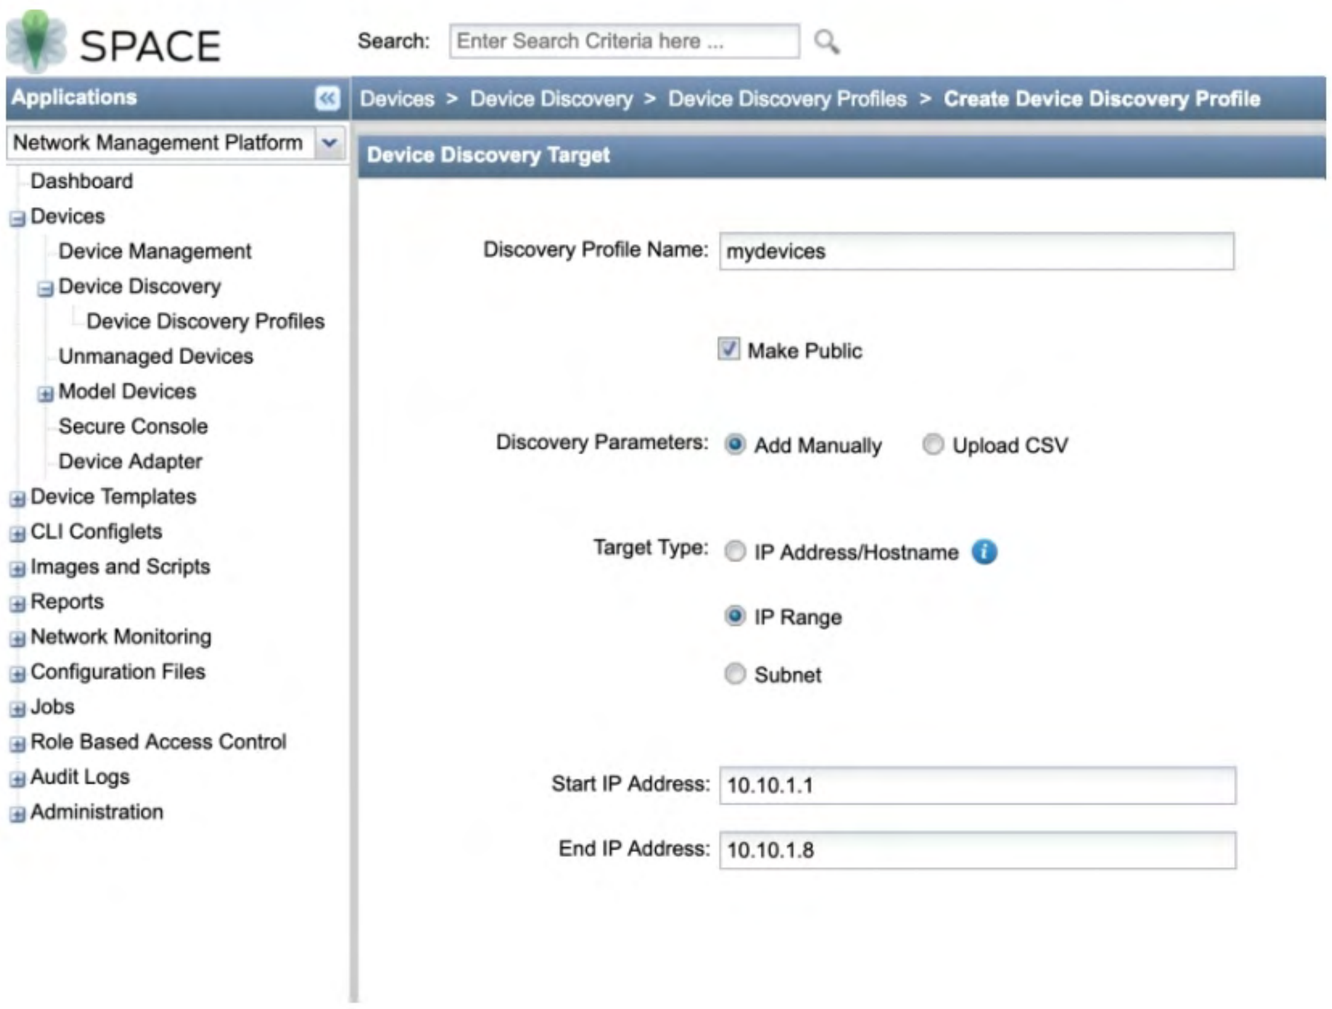

Select and configure the probing mechanisms so that the vMX devices can be discovered. Make sure the specify the correct community name ‘jncie’ as configured on the vMX devices. Select “Next”.

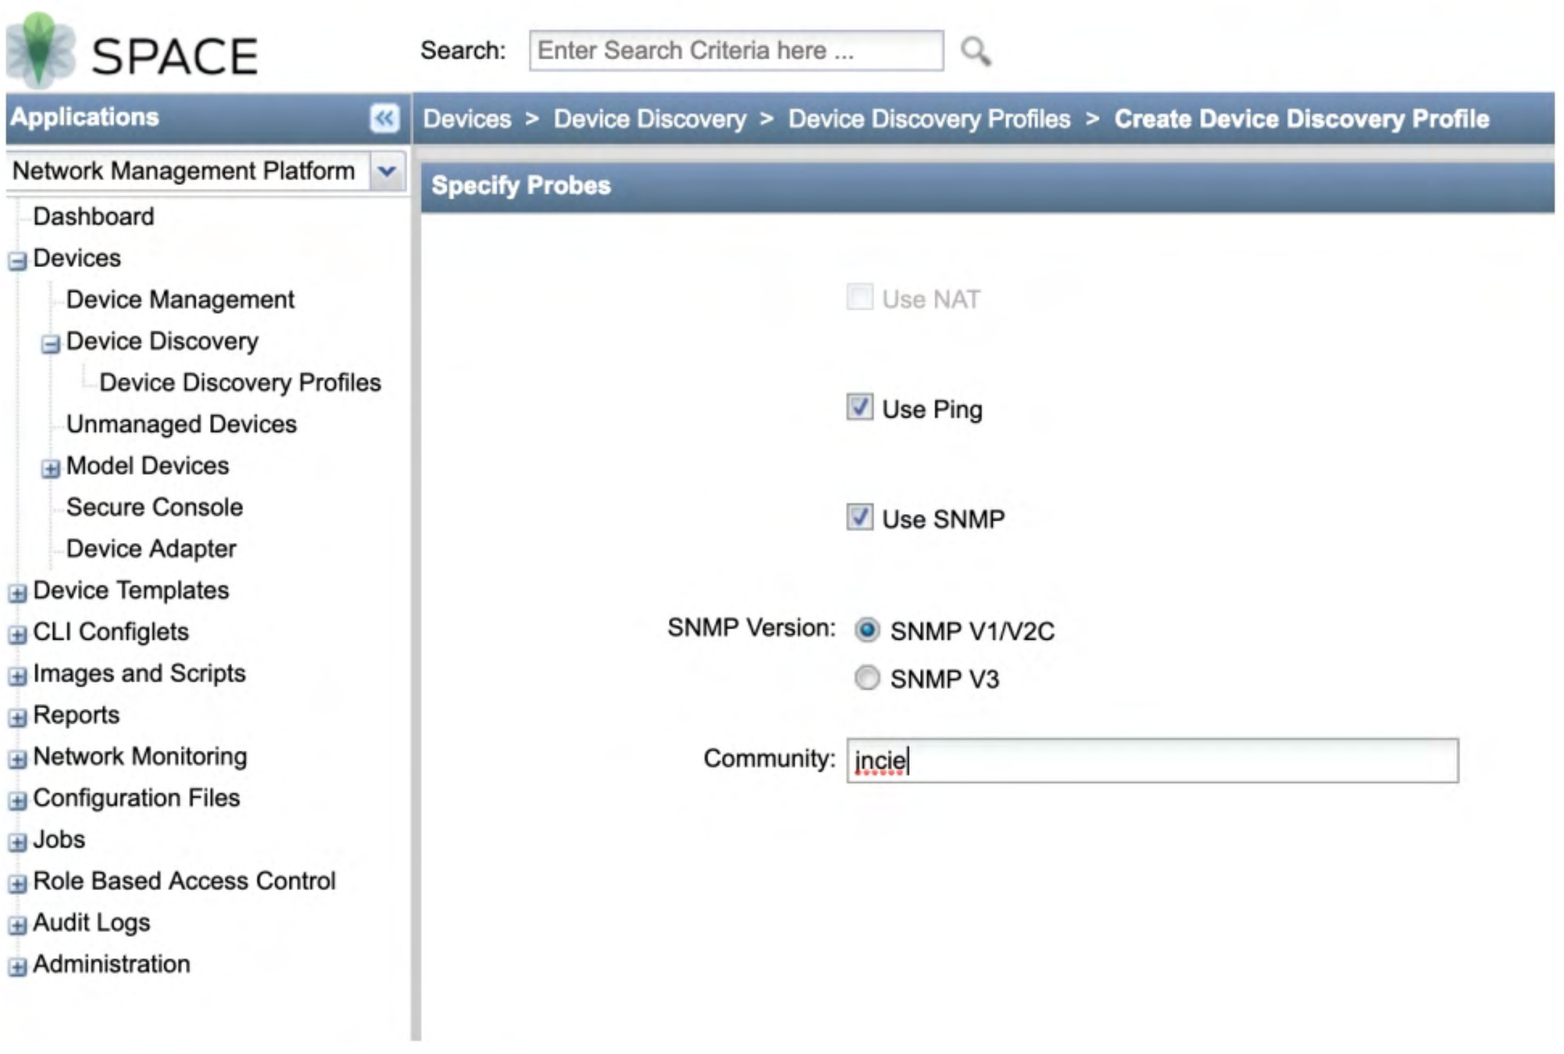

© 2020 Juniper Networks, Inc. All rights reserved.

Specify a user with super-user authorization for the vMX devices.

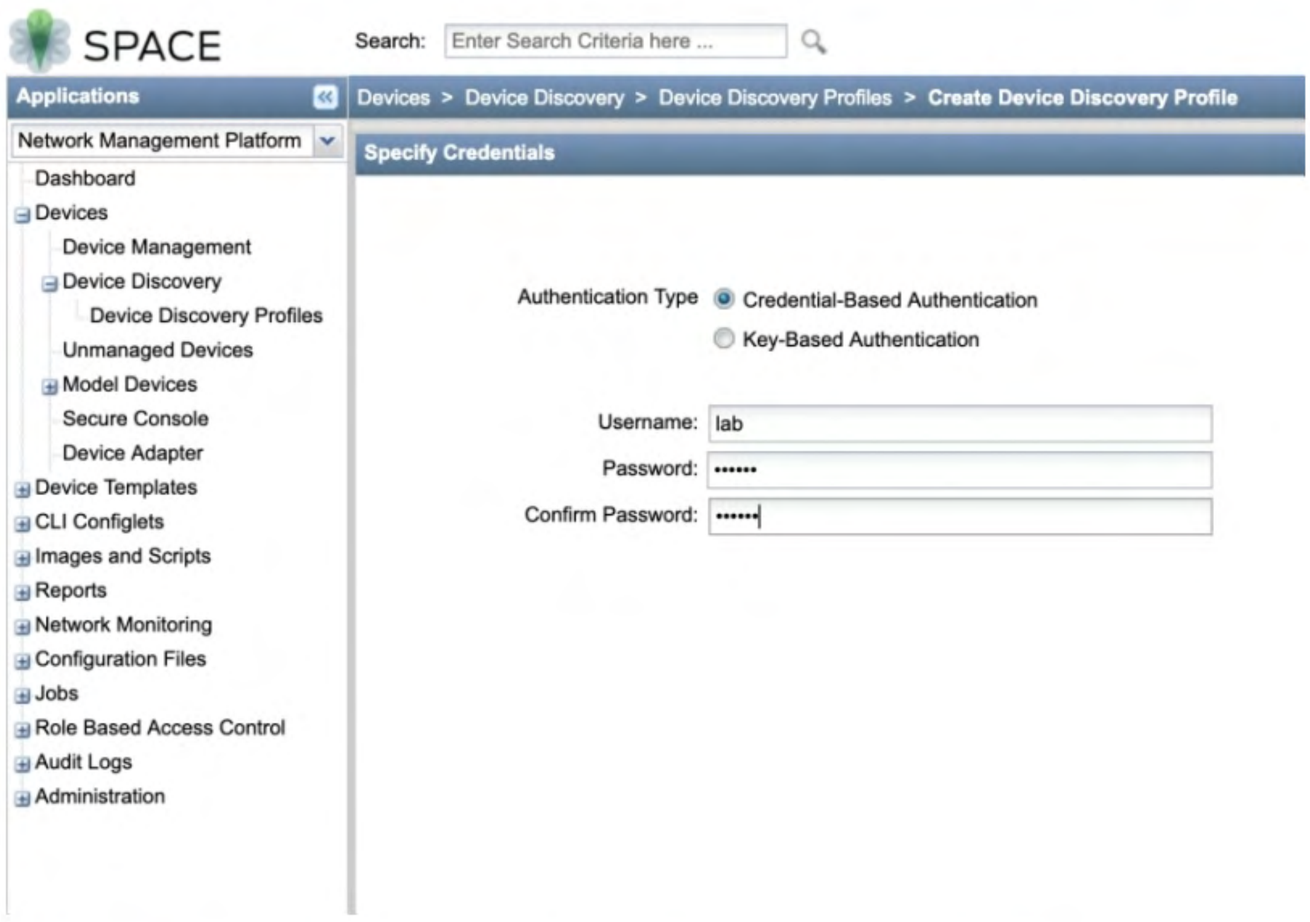

It is not needed to specify a ‘fingerprint’ as it is optional and not required by the task. Select ‘Next’.

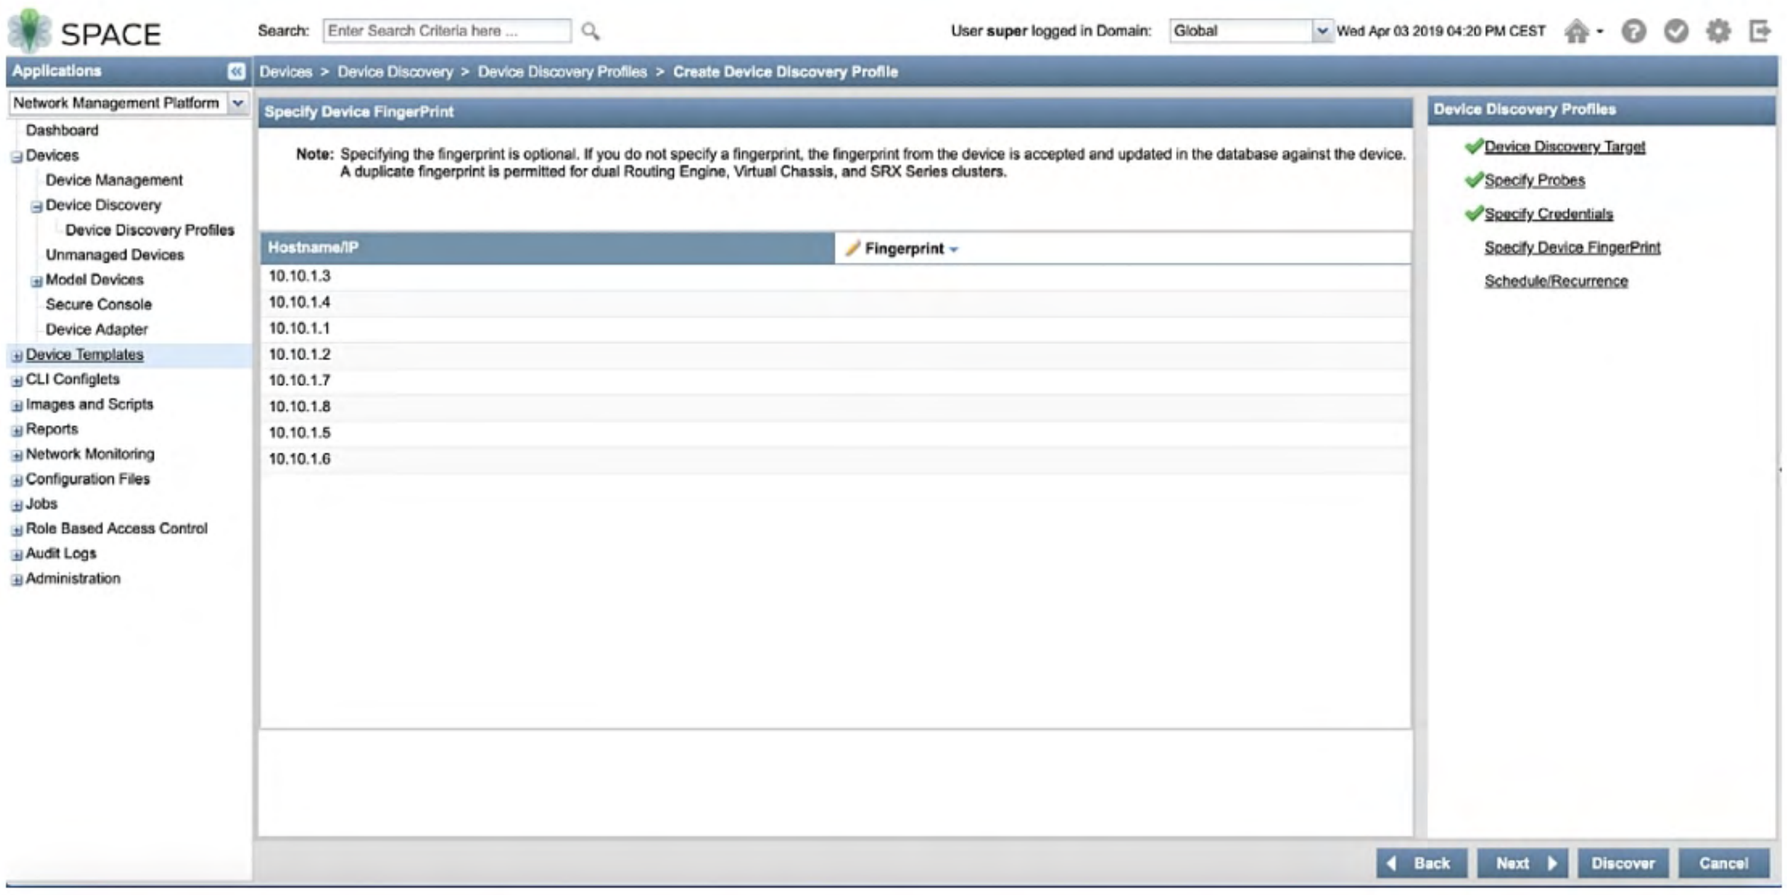

Junos Space provides the option to schedule device discovery at a later time. This is not needed. Select ‘Discover’.

© 2020 Juniper Networks, Inc. All rights reserved.

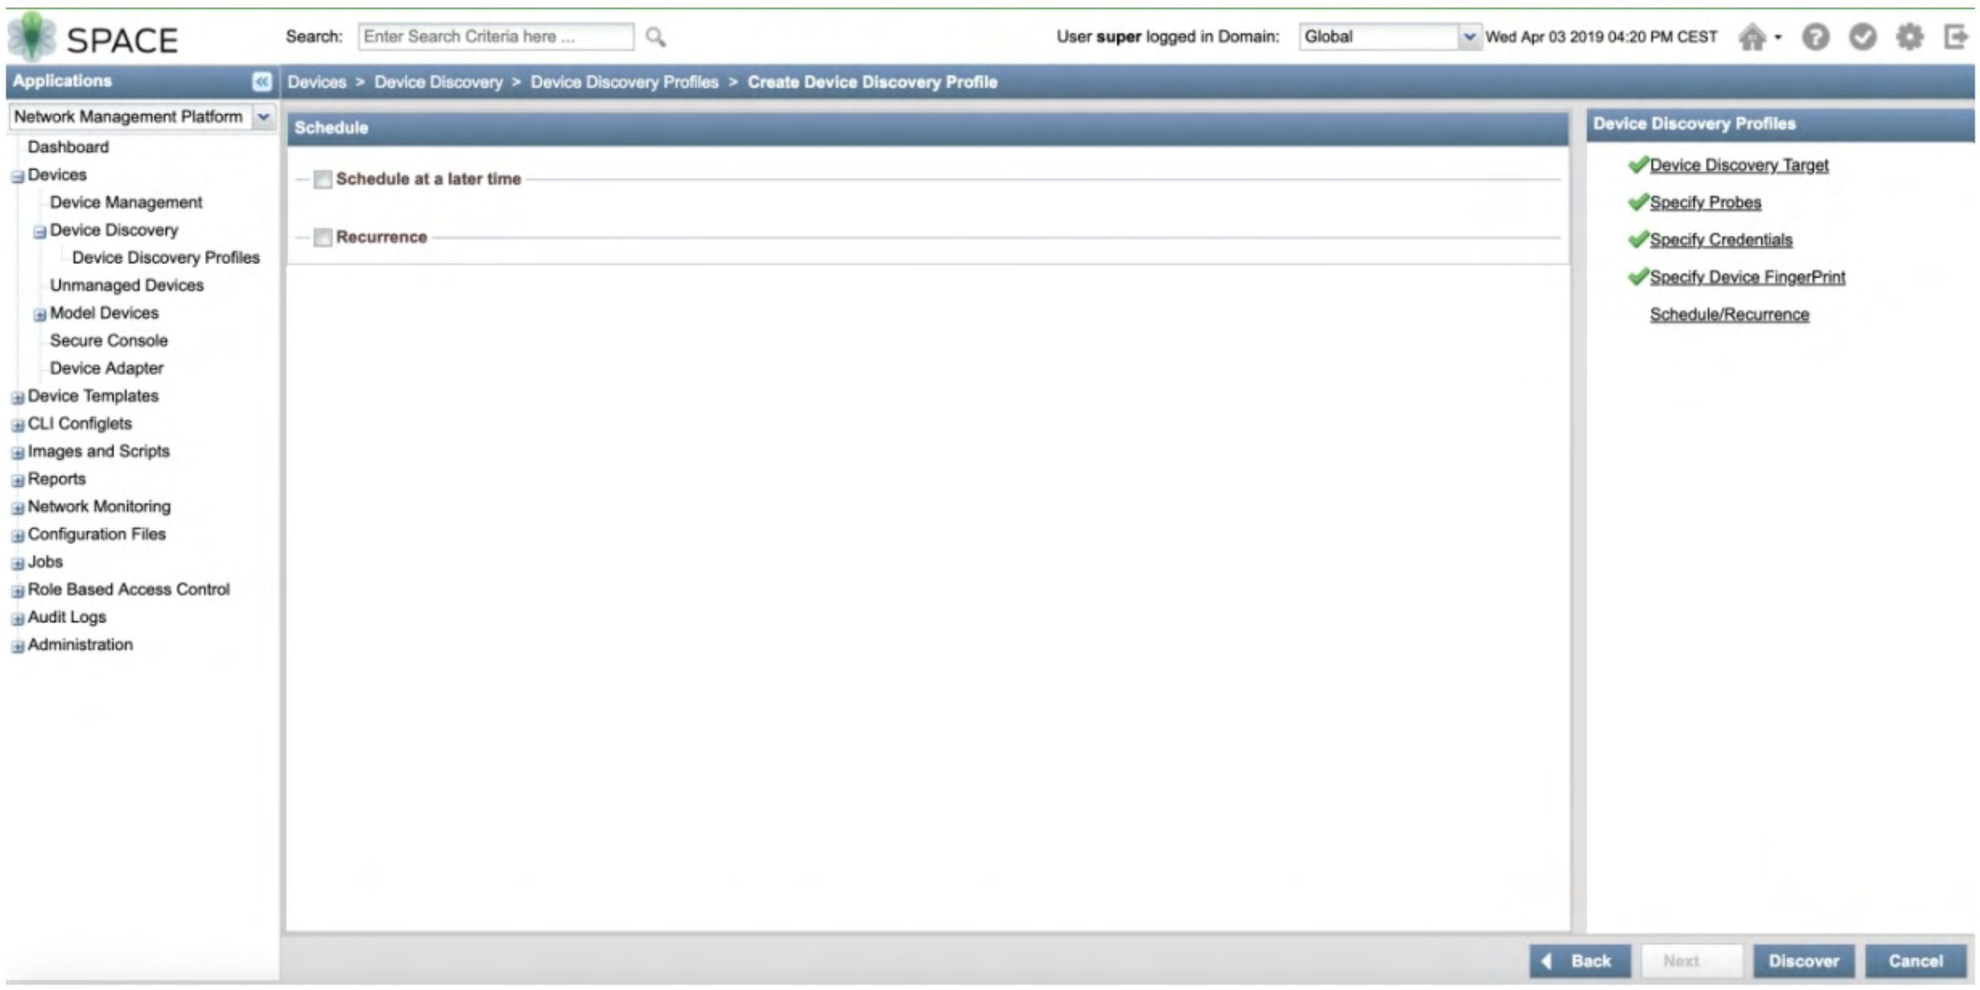

Now Junos Space start the actual discovery process. The Discovery Status shows if the devices where discovered of if the discovery process failed. Once the devices are fully discovered they will become “managed” as shown in the following figure.

Note: This complete process can take few minutes, so be patient.

© 2020 Juniper Networks, Inc. All rights reserved.

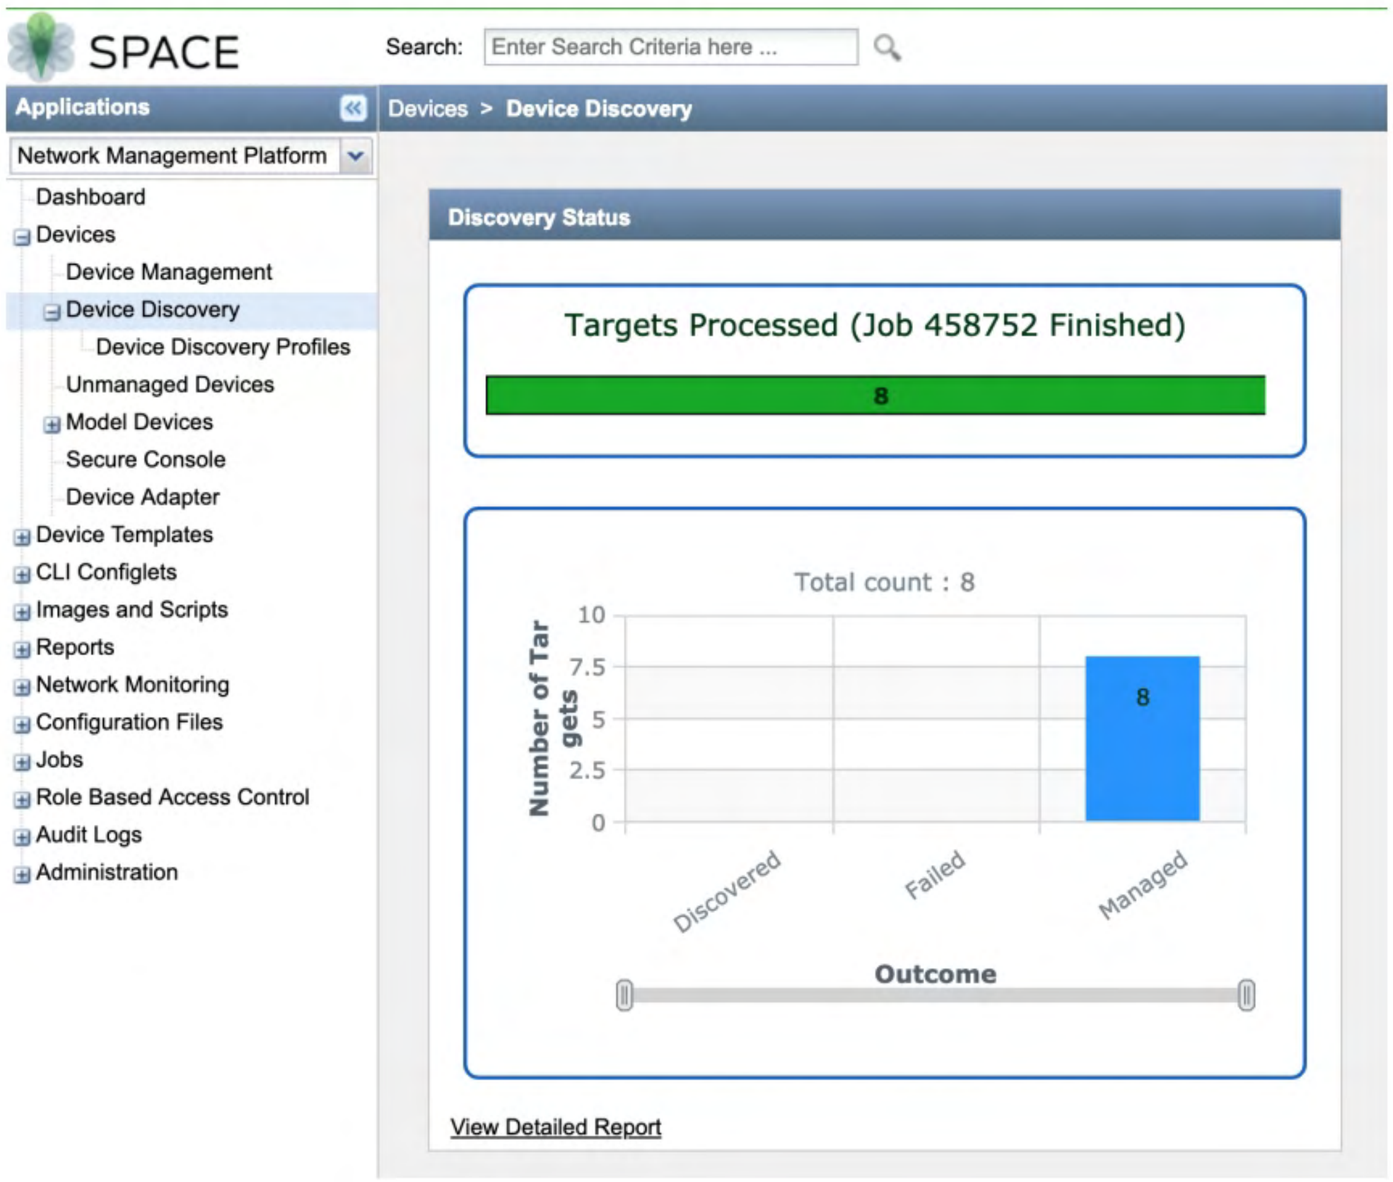

In the left column navigate to “Devices/ Device Management” to verify if all vMX where successfully discovered. Like the figure below you should see all vMXs discovered and in Sync.

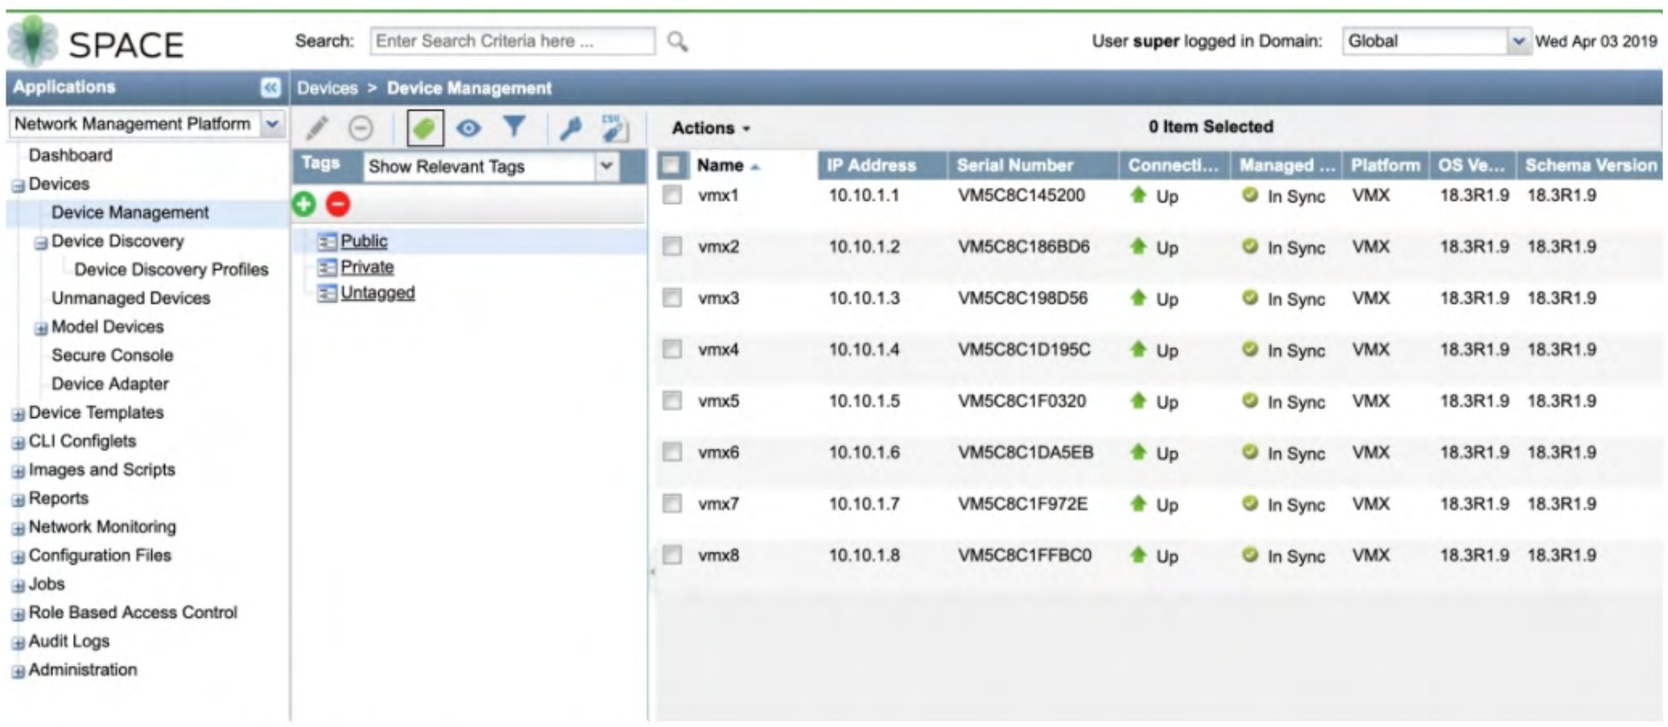

2) Back up all vMX configurations onto Junos Space Network Director.

© 2020 Juniper Networks, Inc. All rights reserved.

Now that all vMX devices are discovered it is possible to backup the configuration files. In the left column navigate to “Configuration Files” and select “Click here to backup devices configuration files”.

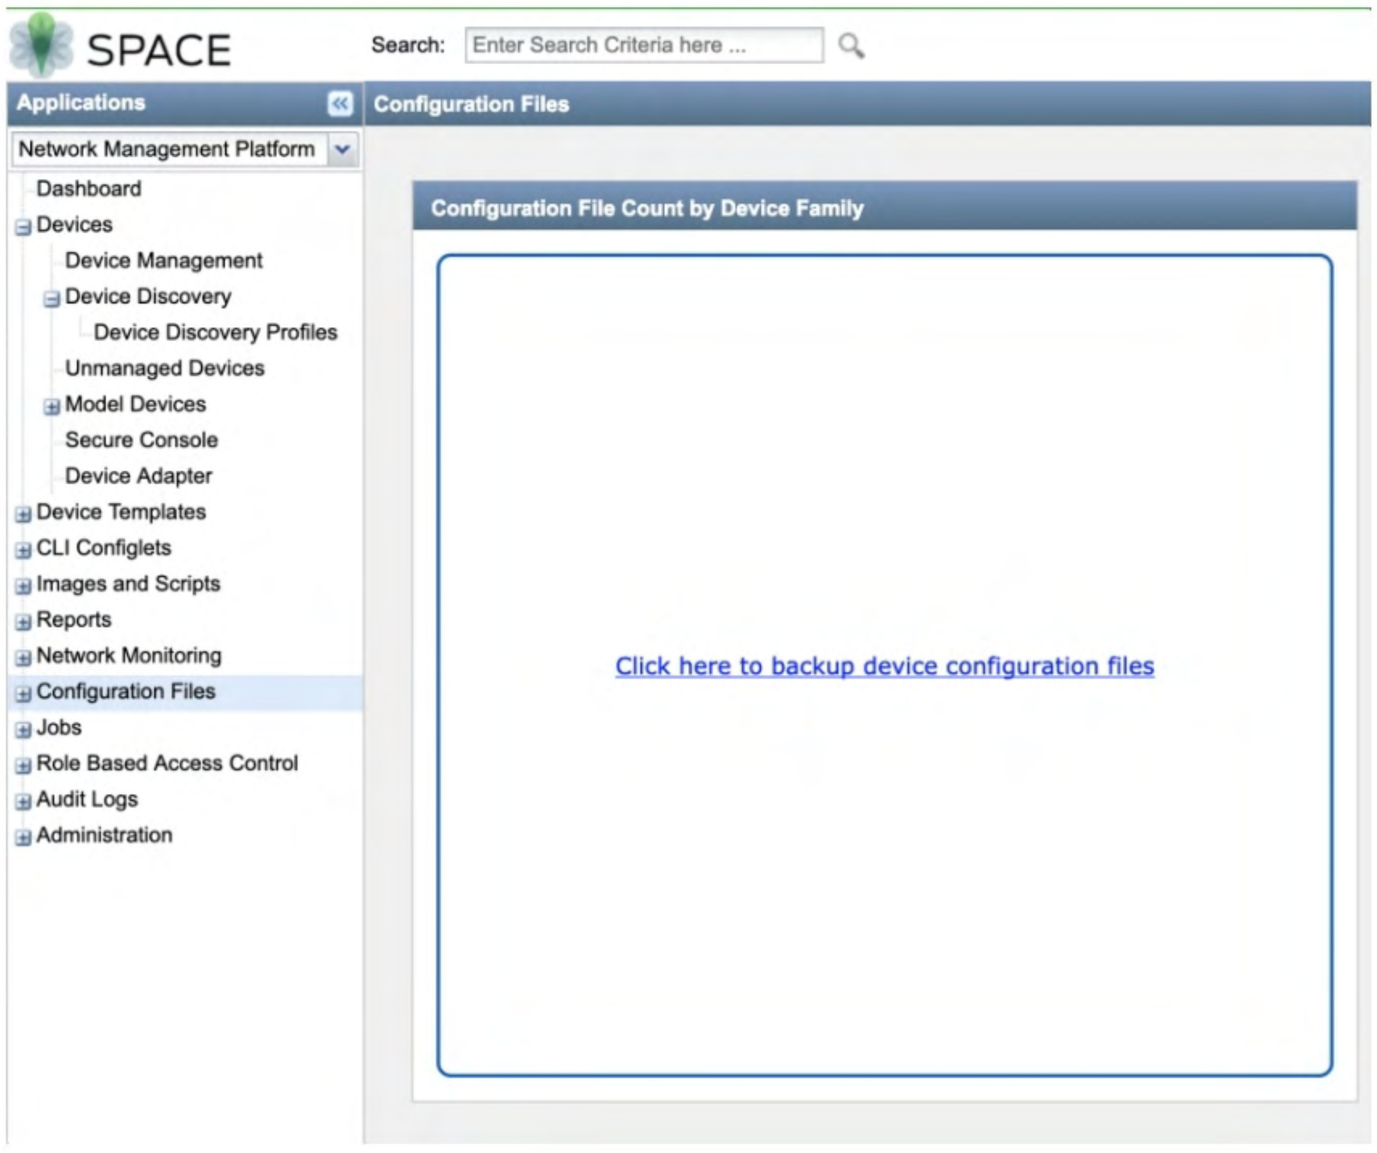

© 2020 Juniper Networks, Inc. All rights reserved.

Next Junos Space provides a list of the devices that are eligible for configuration backup. You should see all vMX devices in the list. Scroll down and select “Backup”.

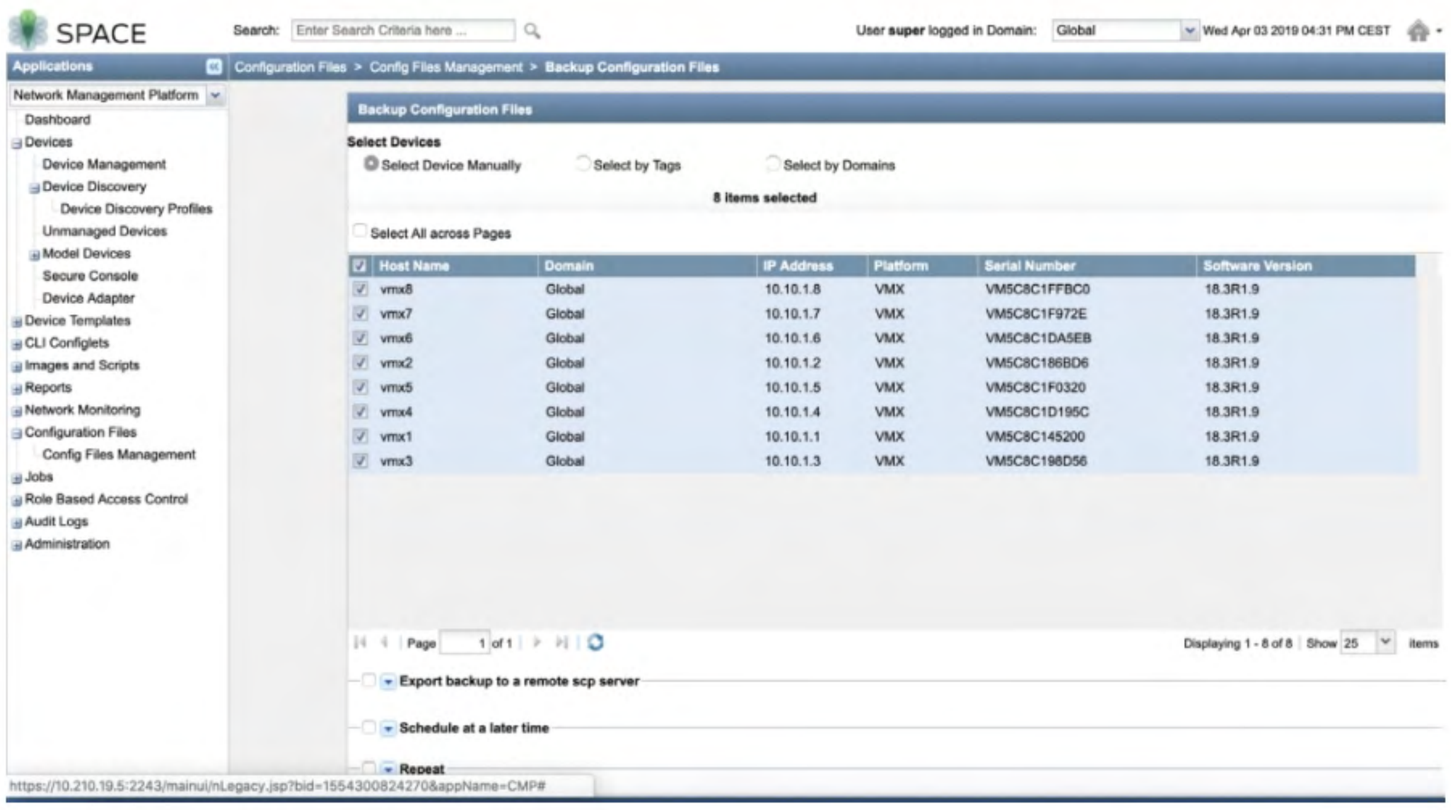

Junos Space will create a Job that runs in the background to download the devices configurations of all vMX devices. Each job is referenced by a job ID. By clicking on the Job ID you can see the details of each job if needed.

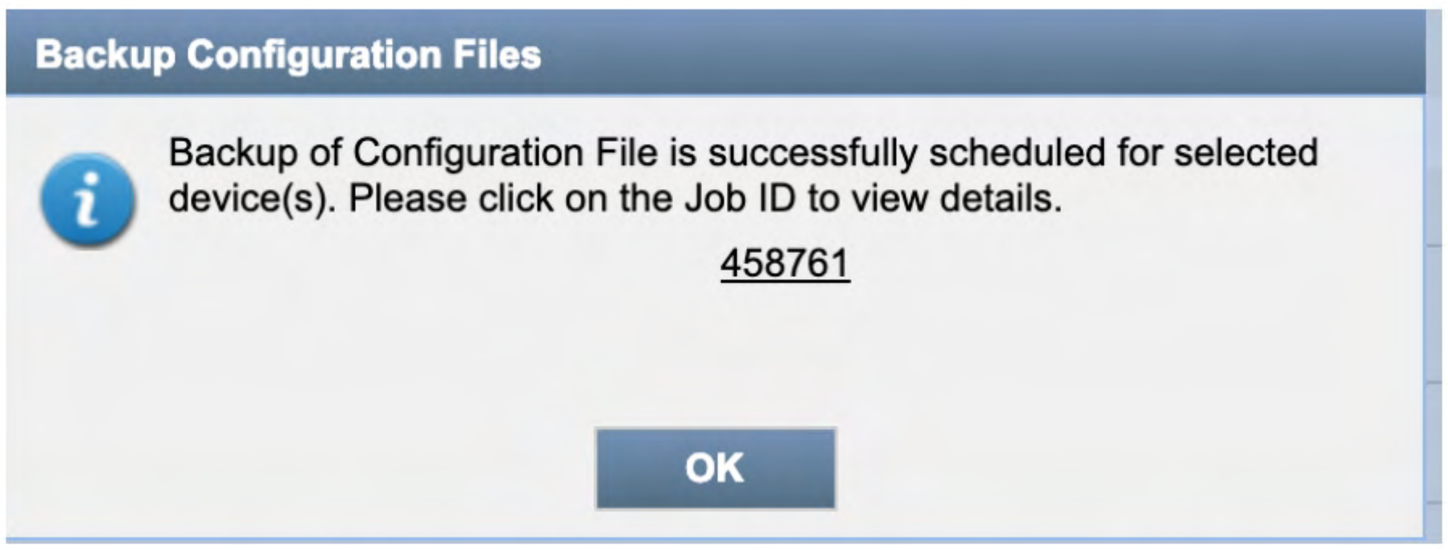

In the left column navigate to “Configuration Files” and select “Config File Management”. You should see the configurations of all vMX devices.

© 2020 Juniper Networks, Inc. All rights reserved.

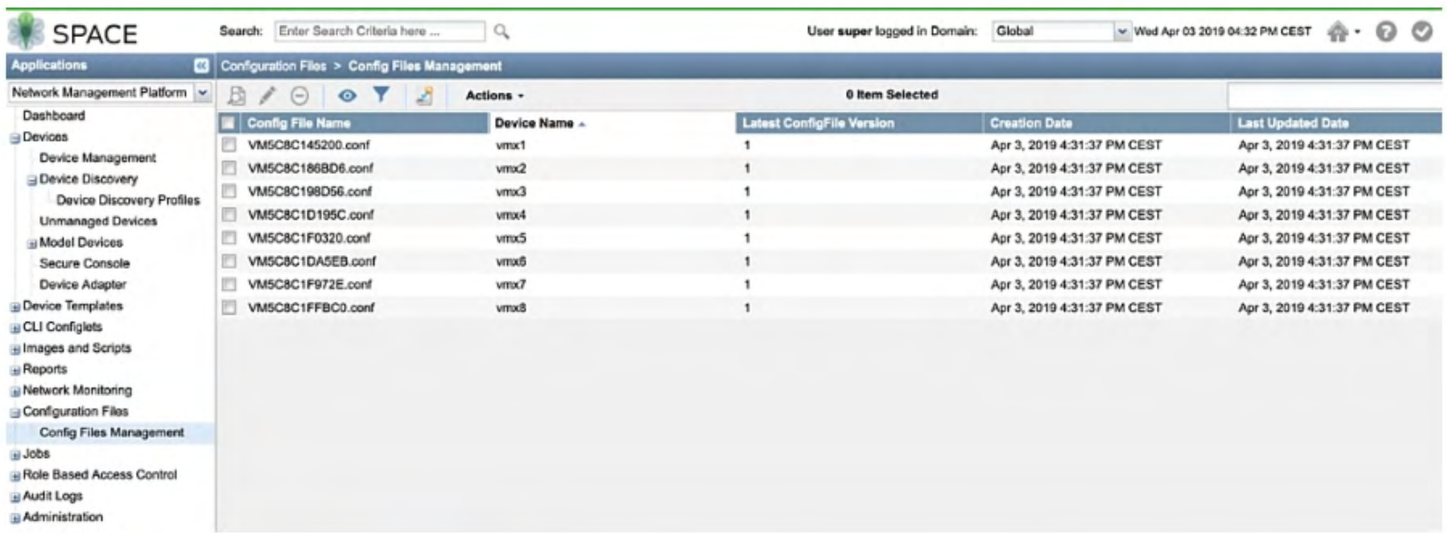

Task 6: Firewall Filters

In this task you will configure RE (Routing Engine) protection using a firewall filter for all vMX devices in the data center.

1) Configure an IPv4 firewall filter which allows protocol messages from the BFD, VRRP, OSPF, BGP. Make sure that for the protocols VRRP and OSPF messages from directly connected subnets are accepted.
2) Configure the firewall filter to accept RADIUS, DNS, SNMP management protocols only from the 10.10.10/24 network.
3) Accept NTP messages only from the 10.10.10/24 and 10.10.1/24 networks.
4) Accept SSH and Telnet protocols only from the 10.10.1/24 management network.
5) Configure the devices to accept ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer, which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
6) Protect the devices from TCP syn/rst/fin attacks. Use a policer with the following parameters: bandwidth 500kbps and burst size 125kB. The excess traffic must be dropped.
7) Configure the firewall filter to discard any other traffic and count the packets.
8) Apply the firewall filter on vMX1 such as to ensure that it is used for the RE protection.
9) Configure an IPv6 firewall filter which allows protocol messages from OSPF and BGP.
10) Configure the devices to accept IPv6 ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
11) Configure the IPv6 firewall filter to discard any other traffic and count the packets.
12) Apply the IPv6 firewall filter on VMX1 such as to ensure that it is used for the RE protection.

© 2020 Juniper Networks, Inc. All rights reserved.

13) Ensure that the configuration under firewall hierarchy is not displayed when printed out.

Configure the following on vMX1, vMX2, vMX3 and vMX4:

```
lab@vmx1# show firewall | display omit
family inet {
    filter protect-re-ipv4 {
        apply-flags omit;
        term tcp-flood {
            from {
                protocol tcp;
                tcp-flags "(syn & !ack) | fin | rst";
            }
            then {
                policer tcp-flood-limit;
                accept;
            }
        }
        term bfd {
            from {
                source-prefix-list {
                    direct-subnets;
                }
                protocol udp;
                destination-port 3784;
            }
            then accept;
        }
        term vrrp {
            from {
                source-prefix-list {
                    direct-subnets;
                }
                protocol vrrp;
            }
            then accept;
        }
        term ospf {
            from {
                source-prefix-list {
                    direct-subnets;
                }
                protocol ospf;
            }
            then accept;
        }
        term bgp-source {
            from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            protocol tcp;
            source-port bgp;
        }
        then accept;
    }
    term bgp-destination {
        from {
            protocol tcp;
            destination-port bgp;
        }
        then accept;
    }
    term ntp {
        from {
            source-address {
                10.10.10.0/24;
            }
            source-prefix-list {
                direct-subnets;
            }
            protocol udp;
            source-port ntp;
        }
        then accept;
    }
    term ntp-local {
        from {
            source-prefix-list {
                direct-subnets;
            }
            protocol udp;
            destination-port ntp;
        }
        then accept;
    }
    term snmp {
        from {
            source-address {
                10.10.10.0/24;
                10.10.1.0/24;
            }
            protocol udp;
            destination-port snmp;
        }
        then accept;
    }
    term radius {
        from {
            source-address {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            10.10.10.0/24;
        }
        protocol udp;
        source-port radius;
    }
    then accept;
}
term dns {
    from {
        source-address {
            10.10.10.0/24;
        }
        protocol udp;
        source-port domain;
    }
    then accept;
}
term ssh {
    from {
        source-address {
            10.10.1.0/24;
        }
        protocol tcp;
        destination-port ssh;
    }
    then accept;
}
term ssh-local {
    from {
        source-address {
            10.10.1.0/24;
        }
        protocol tcp;
        source-port ssh;
    }
    then accept;
}
term telnet {
    from {
        source-address {
            10.10.1.0/24;
        }
        protocol tcp;
        destination-port telnet;
    }
    then accept;
}
term telnet-local {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            source-address {
                10.10.1.0/24;
            }
            protocol tcp;
            source-port telnet;
        }
        then accept;
    }
    term dhcp-1 {
        from {
            protocol udp;
            source-port dhcp;
        }
        then accept;
    }
    term dhcp-2 {
        from {
            protocol udp;
            destination-port dhcp;
        }
        then accept;
    }
    term icmp {
        from {
            protocol icmp;
        }
        then {
            policer re-policer;
            accept;
        }
    }
    term traceroute {
        from {
            protocol udp;
            destination-port 33434-33534;
        }
        then {
            policer re-policer;
            accept;
        }
    }
    term last {
        then {
            count dropped-ipv4;
            discard;
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
family inet6 {
    filter protect-re-ipv6 {
        apply-flags omit;
        term ospf {
            from {
                next-header ospf;
            }
            then accept;
        }
        term bgp-source {
            from {
                next-header tcp;
                source-port bgp;
            }
            then accept;
        }
        term bgp-destination {
            from {
                next-header tcp;
                destination-port bgp;
            }
            then accept;
        }
        term vrrp {
            from {
                next-header vrrp;
            }
            then accept;
        }
        term icmp {
            from {
                next-header icmp6;
            }
            then {
                policer re-policer;
                accept;
            }
        }
        term traceroute {
            from {
                next-header udp;
                destination-port 33434-33534;
            }
            then {
                policer re-policer;
                accept;
            }
        }
        term last {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            then {
                count dropped-ipv6;
                discard;
            }
        }
    }
}
policer re-policer {
    apply-flags omit;
    if-exceeding {
        bandwidth-limit 100k;
        burst-size-limit 25k;
    }
    then discard;
}
policer tcp-flood-limit {
    apply-flags omit;
    if-exceeding {
        bandwidth-limit 500k;
        burst-size-limit 125k;
    }
    then discard;
}

[edit]
lab@vmx1# show policy-options prefix-list direct-subnets
apply-path "interfaces <*> unit <*> family inet address <*>";
```

Activate the IPv4 and IPv5 firewall filter on vMX1.

vMX1:

```
lab@vmx1# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re-ipv4;
        }
    }
    family inet6 {
        filter {
            input protect-re-ipv6;
        }
}
```

Configurations for Task Completion:

On vMX1, vMX2, vMX3, and vMX4:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 apply-flags omit
set firewall family inet filter protect-re-ipv4 term tcp-flood from protocol tcp
set firewall family inet filter protect-re-ipv4 term tcp-flood from tcp-flags "(syn &
!ack) | fin | rst"
set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-
limit
set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
set firewall family inet filter protect-re-ipv4 term bfd from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term bfd from protocol udp
set firewall family inet filter protect-re-ipv4 term bfd from destination-port 3784
set firewall family inet filter protect-re-ipv4 term bfd then accept
set firewall family inet filter protect-re-ipv4 term vrrp from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term vrrp from protocol vrrp
set firewall family inet filter protect-re-ipv4 term vrrp then accept
set firewall family inet filter protect-re-ipv4 term ospf from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ospf from protocol ospf
set firewall family inet filter protect-re-ipv4 term ospf then accept
set firewall family inet filter protect-re-ipv4 term bgp-source from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp-source from source-port bgp
set firewall family inet filter protect-re-ipv4 term bgp-source then accept
set firewall family inet filter protect-re-ipv4 term bgp-destination from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp-destination from destination-
port bgp
set firewall family inet filter protect-re-ipv4 term bgp-destination then accept
set firewall family inet filter protect-re-ipv4 term ntp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term ntp from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term ntp from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp from source-port ntp
set firewall family inet filter protect-re-ipv4 term ntp then accept
set firewall family inet filter protect-re-ipv4 term ntp-local from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ntp-local from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp-local from destination-port ntp
set firewall family inet filter protect-re-ipv4 term ntp-local then accept
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term snmp from protocol udp
set firewall family inet filter protect-re-ipv4 term snmp from destination-port snmp
set firewall family inet filter protect-re-ipv4 term snmp then accept
set firewall family inet filter protect-re-ipv4 term radius from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term radius from protocol udp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term radius from source-port radius
set firewall family inet filter protect-re-ipv4 term radius then accept
set firewall family inet filter protect-re-ipv4 term dns from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term dns from protocol udp
set firewall family inet filter protect-re-ipv4 term dns from source-port domain
set firewall family inet filter protect-re-ipv4 term dns then accept
set firewall family inet filter protect-re-ipv4 term ssh from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh from destination-port ssh
set firewall family inet filter protect-re-ipv4 term ssh then accept
set firewall family inet filter protect-re-ipv4 term ssh-local from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh-local from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh-local from source-port ssh
set firewall family inet filter protect-re-ipv4 term ssh-local then accept
set firewall family inet filter protect-re-ipv4 term telnet from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet from destination-port telnet
set firewall family inet filter protect-re-ipv4 term telnet then accept
set firewall family inet filter protect-re-ipv4 term telnet-local from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet-local from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet-local from source-port
telnet
set firewall family inet filter protect-re-ipv4 term telnet-local then accept
set firewall family inet filter protect-re-ipv4 term dhcp-1 from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp-1 from source-port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp-1 then accept
set firewall family inet filter protect-re-ipv4 term dhcp-2 from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp-2 from destination-port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp-2 then accept
set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp
set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer
set firewall family inet filter protect-re-ipv4 term icmp then accept
set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp
set firewall family inet filter protect-re-ipv4 term traceroute from destination-port
33434-33534
set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
set firewall family inet filter protect-re-ipv4 term traceroute then accept
set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4
set firewall family inet filter protect-re-ipv4 term last then discard
set firewall family inet6 filter protect-re-ipv6 apply-flags omit
set firewall family inet6 filter protect-re-ipv6 term ospf from next-header ospf
set firewall family inet6 filter protect-re-ipv6 term ospf then accept
set firewall family inet6 filter protect-re-ipv6 term bgp-source from next-header tcp
set firewall family inet6 filter protect-re-ipv6 term bgp-source from source-port bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet6 filter protect-re-ipv6 term bgp-source then accept
set firewall family inet6 filter protect-re-ipv6 term bgp-destination from next-header
tcp
set firewall family inet6 filter protect-re-ipv6 term bgp-destination from destination-
port bgp
set firewall family inet6 filter protect-re-ipv6 term bgp-destination then accept
set firewall family inet6 filter protect-re-ipv6 term vrrp from next-header vrrp
set firewall family inet6 filter protect-re-ipv6 term vrrp then accept
set firewall family inet6 filter protect-re-ipv6 term icmp from next-header icmp6
set firewall family inet6 filter protect-re-ipv6 term icmp then policer re-policer
set firewall family inet6 filter protect-re-ipv6 term icmp then accept
set firewall family inet6 filter protect-re-ipv6 term traceroute from next-header udp
set firewall family inet6 filter protect-re-ipv6 term traceroute from destination-port
33434-33534
set firewall family inet6 filter protect-re-ipv6 term traceroute then policer re-policer
set firewall family inet6 filter protect-re-ipv6 term traceroute then accept
set firewall family inet6 filter protect-re-ipv6 term last then count dropped-ipv6
set firewall family inet6 filter protect-re-ipv6 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set firewall policer tcp-flood-limit apply-flags omit
set firewall policer tcp-flood-limit if-exceeding bandwidth-limit 500k
set firewall policer tcp-flood-limit if-exceeding burst-size-limit 125k
set firewall policer tcp-flood-limit then discard
set policy-options prefix-list direct-subnets apply-path "interfaces <*> unit <*> family
inet address <*>"
set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
set interfaces lo0 unit 0 family inet6 filter input protect-re-ipv6
```

Appendix – Chapter 2: Layer 2 Configuration and Features

Spanning Tree Overview

Multiple Spanning Tree Protocol (MSTP) maps multiple VLANs into a spanning tree instance, with each instance having a spanning tree topology independent of other spanning tree instances.

Ethernet networks are susceptible to broadcast storms if loops are introduced. However, an Ethernet network needs to include loops because they provide redundant paths in case of a link failure. Spanning-tree protocols address both of these issues because they provide link redundancy while simultaneously preventing undesirable loops.

Spanning-tree protocols intelligently avoid loops in a network by creating a tree topology (spanning tree) of the entire bridged network with only one available path between the tree root and a leaf. All other paths are forced into a standby state. The tree root is a switch

© 2020 Juniper Networks, Inc. All rights reserved.

within the network elected by the STA (spanning-tree algorithm) to use when computing the best path between bridges throughout the network and the root bridge. Frames travel through the network to their destination– a leaf. A tree branch is a network segment, or link, between bridges. Switches that forward frames through an STP spanning-tree are called designated bridges.

Juniper Networks devices provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). This topic explains MSTP.

MSTP is an extension of RSTP that maps multiple independent spanning-tree instances onto one physical topology. Each spanning-tree instance (STI) includes one or more VLANs. Unlike in STP and RSTP configurations, a port might belong to multiple VLANs and be dynamically blocked in one spanning-tree instance but forwarding in another. This behavior significantly improves network resource utilization by load-balancing across the network and maintaining switch CPU loads at moderate levels. MSTP also leverages the fast reconvergence time of RSTP when a network, switch, or port failure occurs within a spanning-tree instance.

MSTP creates a common and internal spanning tree (CIST) to interconnect and manage all MSTP regions and even individual devices that run RSTP or STP, which are recognized as distinct spanning-tree regions by MSTP. The CIST views each MSTP region as a virtual bridge, regardless of the actual number of devices participating in the MSTP region and enables multiple spanning-tree instances (MSTIs) to link to other regions. The CIST is a single topology that connects all switches (STP, RSTP, and MSTP devices) through an active topology, ensuring connectivity between LANs and devices within a bridged network. This functionality provided by MSTP enables you to better utilize network resources while remaining backward-compatible with older network devices.

When enabling MSTP, you define one or more MSTP regions. An MSTP region defines a logical domain where multiple spanning-tree instances (MSTIs) can be administered independently of MSTIs in other regions, setting the boundary for bridge protocol data units (BPDUs) sent by one MSTI. An MSTP region is a group of switches that is defined by three parameters:

Region name: User-defined alphanumeric name for the region.

Revision level: User-defined value that identifies the region.

Mapping table: Numerical digest of VLAN-to-instance mappings.

An MSTP region can support up to 64 MSTIs, and each MSTI can support from 1 to 4094 VLANs. When you define a region, MSTP automatically creates an internal spanning-tree

© 2020 Juniper Networks, Inc. All rights reserved.

instance (IST instance 0) that provides the root switch for the region and includes all currently configured VLANs that are not specifically assigned to a user-defined MSTI. An MSTI includes all static VLANs that you specifically add to it. The switch places any dynamically created VLANs in the IST instance by default, unless you explicitly map them to another MSTI. Once you assign a VLAN to a user-defined MSTI, the switch removes the VLAN from the IST instance.

OAM Overview

Juniper Networks Junos operating system (Junos OS) for Juniper Networks allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported:

Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending

© 2020 Juniper Networks, Inc. All rights reserved.

periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

LACP Overview

Link Aggregation Control Protocol (LACP) is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. By default, Ethernet links do not exchange LACP protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit LACP PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the Ethernet link receives them from the remote end. The LACP mode can be active or passive. The transmitting link is known as the actor, and the receiving link is known as the partner. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on

© 2020 Juniper Networks, Inc. All rights reserved.

aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, Aggregation of Multiple Link Segments.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the LAG without user intervention.
- Link monitoring to check whether both ends of the bundle are connected to the correct group.

In a scenario where a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server might not be able to exchange LACP PDUs. In such a situation, you can configure an interface to be in the up state even if no PDUs are exchanged. Use the force-up statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When PDUs are not received, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help prevent communication failure.

When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.

When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

© 2020 Juniper Networks, Inc. All rights reserved.

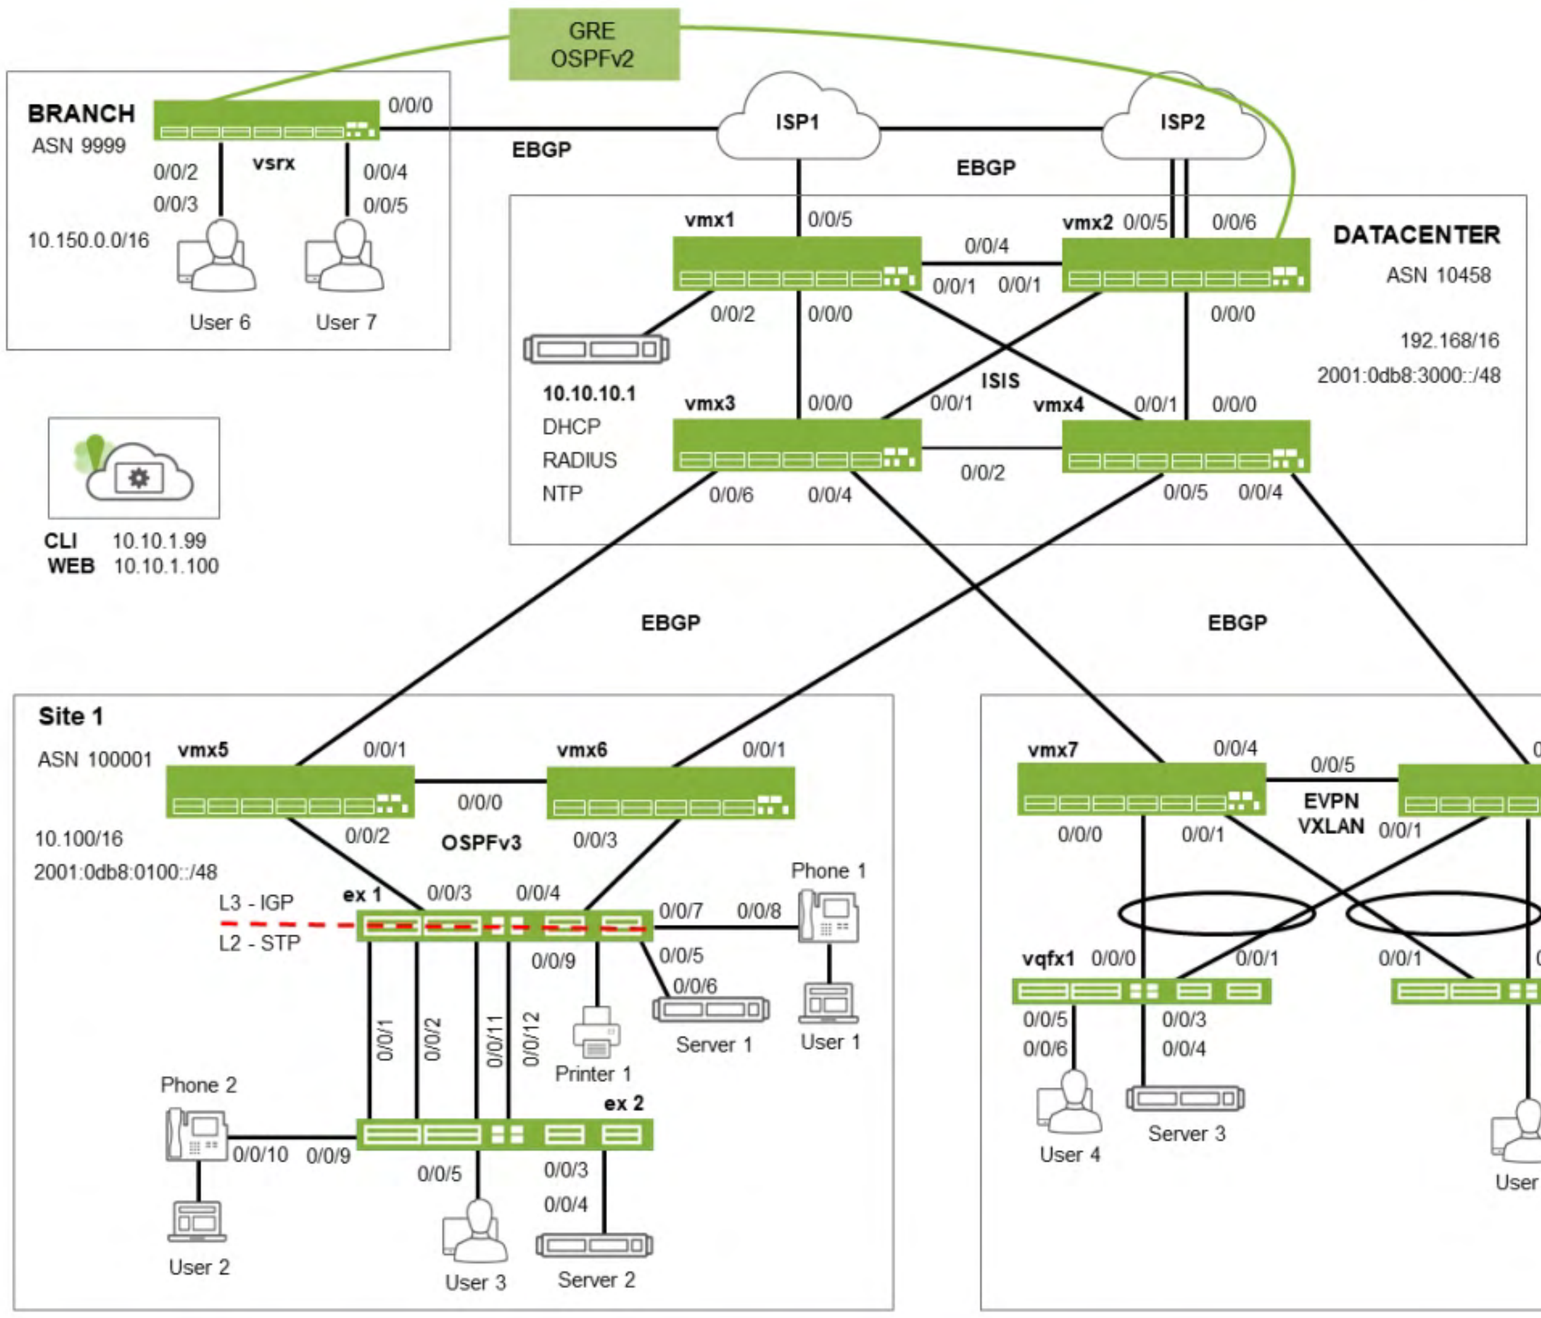

VLAN Name	VLAN ID	Devices	Interface
user	100	EX1	ge-0/0/7 untagged
		EX2	ge-0/0/5 untagged
		EX2	ge-0/0/9 untagged
		vQFX1	ge-0/0/5 untagged
		vQFX2	ge-0/0/3 untagged
server	101	EX1	ge-0/0/5 untagged
		EX2	ge-0/0/3 untagged
		vQFX1	ge-0/0/3 untagged
printer	200	EX1	ge-0/0/9 untagged
phone	201	EX1	ge-0/0/7 tagged
		EX2	ge-0/0/9 tagged
guest	666	EX1, EX2	Dot1x guest vlan

Table 31 VLANs Site1 and Site 2

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: Layer 2 Switching

1) Configure the VLANs and corresponding interfaces in Site1 and Site 2 as shown in the previous table.

EX1:

```
lab@ex1# show vlans
guest {
    vlan-id 666;
}
phone {
    vlan-id 201;
    l3-interface irb.201;
}
printer {
    vlan-id 200;
    l3-interface irb.200;
}
server {
    vlan-id 101;
    l3-interface irb.101;
}
user {
    vlan-id 100;
    l3-interface irb.100;
}

lab@ex1# show interfaces ge-0/0/5
description "to server1";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members server;
        }
    }
}

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/7
description "to phone1 user1";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/9
description "to printer1";
unit 0 {
    family ethernet-switching {
        vlan {
            members printer;
        }
    }
}
```

EX2:

```
lab@ex2# show vlans
phone {
    vlan-id 201;
}
printer {
    vlan-id 200;
}
server {
    vlan-id 101;
}
user {
    vlan-id 100;
}

lab@ex2# show interfaces ge-0/0/3
description "to server2";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members server;
        }
    }
}

{master:0}[edit]
lab@ex2# show interfaces ge-0/0/5
description user3;
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
}

{master:0}[edit]
lab@ex2# show interfaces ge-0/0/9
description "phone2 user2";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```

vQFX1:

```
lab@vQFX1# show vlans
server {
    vlan-id 101;
}
user {
    vlan-id 100;
}

{master:0}[edit]
lab@vQFX1# show interfaces xe-0/0/3
description "to server3";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members server;
        }
    }
}

{master:0}[edit]
lab@vQFX1# show interfaces xe-0/0/5
description "to user4";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vQFX2:

```
lab@vQFX2# show vlans
server {
    vlan-id 101;
}
user {
    vlan-id 100;
}

lab@vQFX2# show interfaces xe-0/0/3
description "to user5";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```

2) Configure aggregate ethernet interface (AE10) between EX1 and EX2. Ensure that the aggregate interface is shutdown if the bandwidth is less than 2Gbit. EX2 should only respond to LACP messages.

EX1:

```
lab@ex1# show chassis
aggregated-devices {
    ethernet {
        device-count 11;
    }
}
...

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/11
gigether-options {
    802.3ad ae10;
}

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/12
gigether-options {
    802.3ad ae10;
}

{master:0}[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@ex1# show interfaces ae10
aggregated-ether-options {
    minimum-links 2;
    lacp {
        active;
    }
}
unit 0;
```

EX2:

```
lab@ex2# show chassis
aggregated-devices {
    ethernet {
        device-count 11;
    }
}

lab@ex2# show interfaces ge-0/0/11
gigether-options {
    802.3ad ae10;
}

{master:0}[edit]
lab@ex2# show interfaces ge-0/0/12
gigether-options {
    802.3ad ae10;
}

lab@ex2# show interfaces ae10
aggregated-ether-options {
    minimum-links 2;
    lacp {
        passive;
    }
}
unit 0;
```

3) The interfaces between EX1 and EX2 should be configured to transport all VLANs as shown in the table above.

EX1 & EX2:

```
{master:0}[edit]
lab@ex1# show interfaces ge-0/0/1
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        vlan {
            members all;
        }
    }
}

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/2
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members all;
        }
    }
}

{master:0}[edit]
lab@ex1# show interfaces ae10
aggregated-ether-options {
    minimum-links 2;
    lacp {
        active;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members all;
        }
    }
}
```

4) Configure aggregate ethernet interface (ae0) on vQFX1 and vQFX2 towards the vMX devices. vQFX1 and vQFX2 should only respond to LACP message.

vQFX1 & vQFX2:

```
{master:0}[edit]
lab@vQFX1# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}

lab@vQFX1# show interfaces xe-0/0/0
gigether-options {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    802.3ad ae0;
}

{master:0}[edit]
lab@vQFX1# show interfaces xe-0/0/1
gigether-options {
    802.3ad ae0;
}

{master:0}[edit]
lab@vQFX1# show interfaces ae0
aggregated-ether-options {
    lacp {
        passive;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members all;
        }
    }
}
```

5) Configure aggregate ethernet interface (AE0) on vMX7 and vMX8 towards vQFX1. Configure aggregate ethernet interface (AE1) on vMX7 and vMX8 towards vQFX2. Enable LACP on the AE interfaces. All Site 2 VLANs must be allowed on the AE0 interfaces.

vMX7:

```
[edit]
lab@vmx7# show chassis aggregated-devices
ethernet {
    device-count 2;
}

[edit]
lab@vmx7# show interfaces ge-0/0/0
gigether-options {
    802.3ad ae0;
}

[edit]
lab@vmx7# show interfaces ge-0/0/1
gigether-options {
    802.3ad ae1;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

[edit]
lab@vmx7# show interfaces ae0
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}

[edit]
lab@vmx7# show interfaces ae1
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 101 100 ];
    }
}
```

vMX8:

```
lab@vmx8# show chassis aggregated-devices
ethernet {
    device-count 2;
}

[edit]
lab@vmx8# show interfaces ge-0/0/0
gigether-options {
    802.3ad ae1;
}

[edit]
lab@vmx8# show interfaces ge-0/0/1
gigether-options {
    802.3ad ae0;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx8# show interfaces ae0
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```

```
[edit]
lab@vmx8# show interfaces ae1
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 101 100 ];
    }
}
```

6) Ensure that VLAN 1 is transported untagged on all trunk interfaces in Site 1. This VLAN must be named ‘default’.

EX1 & EX2:

```
lab@ex1# show vlans default
vlan-id 1;
```

```
lab@ex1# show interfaces ge-0/0/1 native-vlan-id
native-vlan-id 1;

{master:0}[edit]
lab@ex1# show interfaces ge-0/0/2 native-vlan-id
native-vlan-id 1;

{master:0}[edit]
lab@ex1# show interfaces ae10 native-vlan-id
native-vlan-id 1;
```


© 2020 Juniper Networks, Inc. All rights reserved.

7) Configure interface ge-0/0/7 on EX1 and interface ge-0/0/9 on EX2 for Voice VLAN support. Make sure that tagged frames go to VLAN 201 whereas untagged frames go to VLAN 100.

EX1:

```
lab@ex1# show interfaces ge-0/0/7
description "to phone1 user1";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}

lab@ex1# show vlans user
vlan-id 100;
l3-interface irb.100;
switch-options {
    interface ge-0/0/7.0;
}

lab@ex1# show vlans phone
vlan-id 201;
l3-interface irb.201;


lab@ex1# show switch-options
voip {
    interface ge-0/0/7.0 {
        vlan phone;
    }
}
```

Verification:

```
{master:0}[edit]
lab@ex1# run show vlans phone detail

Routing instance: default-switch
VLAN Name: phone                            State: Active
Tag: 201
Internal index: 4, Generation Index: 9, Origin: Static
MAC aging time: 300 seconds
Layer 3 interface: irb.201
VXLAN Enabled : No
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Interfaces:
    ae10.0*,tagged,trunk
    ge-0/0/1.0*,tagged,trunk
    ge-0/0/2.0*,tagged,trunk
    ge-0/0/7.0*,tagged,access
Number of interfaces: Tagged 4    , Untagged 0
Total MAC count: 0

{master:0}[edit]
lab@ex1# run show vlans user detail

Routing instance: default-switch
VLAN Name: user                              State: Active
Tag: 100
Internal index: 7, Generation Index: 12, Origin: Static
MAC aging time: 300 seconds
Layer 3 interface: irb.100
VXLAN Enabled : No
Interfaces:
    ae10.0*,tagged,trunk
    ge-0/0/1.0*,tagged,trunk
    ge-0/0/2.0*,tagged,trunk
    ge-0/0/7.0*,untagged,access
Number of interfaces: Tagged 3    , Untagged 1
Total MAC count: 1
```

EX2:

```
lab@ex2# show interfaces ge-0/0/9
description "phone2 user2";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}

lab@ex2# show vlans phone
vlan-id 201;

{master:0}[edit]
lab@ex2# show vlans user
vlan-id 100;
l3-interface irb.100;
switch-options {
    interface ge-0/0/9.0;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
{master:0}[edit]
lab@ex2# show switch-options
voip {
    interface ge-0/0/9.0 {
        vlan phone;
    }
}
```

Verification:

```
{master:0}[edit]
lab@ex2# run show vlans phone detail

Routing instance: default-switch
VLAN Name: phone                            State: Active
Tag: 201
Internal index: 3, Generation Index: 12, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ae10.0*,tagged,trunk
    ge-0/0/1.0*,tagged,trunk
    ge-0/0/2.0*,tagged,trunk
    ge-0/0/9.0*,tagged,access
Number of interfaces: Tagged 4    , Untagged 0
Total MAC count: 0

{master:0}[edit]
lab@ex2# run show vlans user detail

Routing instance: default-switch
VLAN Name: user                             State: Active
Tag: 100
Internal index: 6, Generation Index: 15, Origin: Static
MAC aging time: 300 seconds
Layer 3 interface: irb.100
VXLAN Enabled : No
Interfaces:
    ae10.0*,tagged,trunk
    ge-0/0/1.0*,tagged,trunk
    ge-0/0/2.0*,tagged,trunk
    ge-0/0/5.0*,untagged,access
    ge-0/0/9.0*,untagged,access
Number of interfaces: Tagged 3    , Untagged 2
Total MAC count: 2
```

Configurations for Task Completion:

© 2020 Juniper Networks, Inc. All rights reserved.

EX1:

```
set vlans guest vlan-id 666
set vlans phone vlan-id 201
set vlans phone l3-interface irb.201
set vlans printer vlan-id 200
set vlans printer l3-interface irb.200
set vlans server vlan-id 101
set vlans server l3-interface irb.101
set vlans user vlan-id 100
set vlans user l3-interface irb.100
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members server
set interfaces ge-0/0/7 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/7 unit 0 family ethernet-switching vlan members user
set interfaces ge-0/0/9 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members printer
set chassis aggregated-devices ethernet device-count 11
set interfaces ge-0/0/11 gigether-options 802.3ad ae10
set interfaces ge-0/0/12 gigether-options 802.3ad ae10
set interfaces ae10 aggregated-ether-options minimum-links 2
set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 unit 0
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members all
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members all
set interfaces ae10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae10 unit 0 family ethernet-switching vlan members all
set vlans default vlan-id 1
set interfaces ge-0/0/1 native-vlan-id 1
set interfaces ge-0/0/2 native-vlan-id 1
set interfaces ae10 native-vlan-id 1
set vlans user switch-options interface ge-0/0/7.0
set switch-options voip interface ge-0/0/7.0 vlan phone
```

EX2:

```
set vlans phone vlan-id 201
set vlans printer vlan-id 200
set vlans server vlan-id 101
set vlans user vlan-id 100
set vlans user l3-interface irb.100
set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members server
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members user
set interfaces ge-0/0/9 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members user
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set chassis aggregated-devices ethernet device-count 11
set interfaces ge-0/0/11 gigether-options 802.3ad ae10
set interfaces ge-0/0/12 gigether-options 802.3ad ae10
set interfaces ae10 aggregated-ether-options minimum-links 2
set interfaces ae10 aggregated-ether-options lacp passive
set interfaces ae10 unit 0
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members all
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members all
set interfaces ae10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae10 unit 0 family ethernet-switching vlan members all
set vlans default vlan-id 1
set interfaces ge-0/0/1 native-vlan-id 1
set interfaces ge-0/0/2 native-vlan-id 1
set interfaces ae10 native-vlan-id 1
set vlans user switch-options interface ge-0/0/9.0
set switch-options voip interface ge-0/0/9.0 vlan phone
```

vQFX1:

```
set vlans server vlan-id 101
set vlans user vlan-id 100
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members server
set interfaces xe-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching vlan members user

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members all
```

vQFX2:

```
set vlans server vlan-id 101
set vlans user vlan-id 100
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members user
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members all
```

vMX7:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-0/0/0 gigether-options 802.3ad ae0
set interfaces ge-0/0/1 gigether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list [100 101]
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list [100 101]
```

vMX8:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-0/0/0 gigether-options 802.3ad ae1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list [100 101]
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list [100 101]
```

Task 2: Spanning Tree Configuration

In this task you will configure the Spanning Tree protocol and various features.

1) Configure EX1 and EX2 in MSTP region jncie with two instances: Instance 1 must be bound to VLAN100-199. Instance 2 must be bound to VLAN200-299. Ensure that Instance 1 is rooted at EX1 and Instance 2 is rooted at EX2. The CIST root is at EX1.

EX1:

```
lab@ex1# delete protocols rstp

lab@ex1# show protocols mstp
configuration-name jncie;
revision-level 1;
bridge-priority 4k;
interface ge-0/0/1;
interface ge-0/0/2;
interface ae10;

msti 1 {
    bridge-priority 4k;
    vlan 100-199;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
msti 2 {
    vlan 200-299;
}
```

EX2:

```
lab@ex2# delete protocols rstp

lab@ex2# show protocols mstp
configuration-name jncie;
revision-level 1;
interface ge-0/0/1;
interface ge-0/0/2 {
interface ae10;

msti 1 {
    vlan 100-199;
}
msti 2 {
    bridge-priority 4k;
    vlan 200-299;
}
```

9) Make sure that for the CIST interface AE10 is the root port on EX2. Only if interface AE10 and interface ge-0/0/2 are down the interface ge-0/0/1 should be the root port.

EX2:

```
lab@ex2# show protocols mstp
configuration-name jncie;
revision-level 1;
interface ge-0/0/1;
interface ge-0/0/2 {
    cost 19999;
}
```

10) Make sure that for Instance 1 interface ge-0/0/2 is root port on EX2. Only if interface ge-0/0/1 and ge-0/0/2 are down should the interface AE10 become the root port.

EX2:

```
lab@ex2# show protocols mstp
configuration-name jncie;
revision-level 1;
interface ge-0/0/1;
interface ge-0/0/2 {
    cost 19999;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
interface access-ports {
    edge;
}
interface ae10;
msti 1 {
    vlan 100-199;
    interface ge-0/0/1 {
        cost 9999;
    }
    interface ge-0/0/2 {
        cost 9998;
    }
}
msti 2 {
    bridge-priority 4k;
    vlan 200-299;
}
```

11) Ensure that ports ge-0/0/5 , ge-0/0/7 and ge-0/0/9 on EX1 and ports ge-0/0/3, ge-0/0/5 and ge-0/0/9 on EX2 transition immediately to L2 forwarding state. Make sure that if BPDUs are received on any of those interfaces they will be automatically disabled. The edge port interfaces must be grouped under the MSTP configuration using the name “access-ports”.

EX1:

```
lab@ex1# show interfaces interface-range access-ports
member ge-0/0/5;
member ge-0/0/7;
member ge-0/0/9;

lab@ex1# show protocols mstp
configuration-name jncie;
revision-level 1;
bridge-priority 4k;
bpdu-block-on-edge;
interface ge-0/0/1;
interface ge-0/0/2;
interface access-ports {
    edge;
}
interface ae10;
msti 1 {
    bridge-priority 4k;
    vlan 100-199;
}
msti 2 {
    vlan 200-299;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

EX2:

```
lab@ex2# show interfaces interface-range access-ports
member ge-0/0/3;
member ge-0/0/5;
member ge-0/0/9;

lab@ex2# show protocols mstp
configuration-name jncie;
revision-level 1;
bpdu-block-on-edge;
interface ge-0/0/1;
interface ge-0/0/2 {
    cost 19999;
}
interface access-ports {
    edge;
}
interface ae10;
msti 1 {
    vlan 100-199;
    interface ge-0/0/1 {
        cost 9999;
    }
    interface ge-0/0/2 {
        cost 9998;
    }
}
msti 2 {
    bridge-priority 4k;
    vlan 200-299;
}
```

Verification:

EX1:

```
lab@ex1# run show spanning-tree bridge
STP bridge parameters
Routing instance name              : GLOBAL
Context ID                         : 0
Enabled protocol                   : MSTP

STP bridge parameters for CIST
  Root ID                          : 4096.cc:e1:94:d2:6e:a1
  CIST regional root               : 4096.cc:e1:94:d2:6e:a1
  CIST internal root cost          : 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Hello time                          : 2 seconds
  Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Number of topology changes          : 1
  Time since last topology change     : 27 seconds
  Local parameters
    Bridge ID                         : 4096.cc:e1:94:d2:6e:a1

STP bridge parameters for MSTI 1
  MSTI regional root                  : 4097.cc:e1:94:d2:6e:a1
  Hello time                          : 2 seconds
  Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Number of topology changes          : 2
  Time since last topology change     : 21 seconds
  Local parameters
    Bridge ID                         : 4097.cc:e1:94:d2:6e:a1

STP bridge parameters for MSTI 2
  MSTI regional root                  : 4098.b8:c2:53:97:98:01
  Root cost                           : 10000
  Root port                           : ae10
  Hello time                          : 2 seconds
  Maximum age                         : 20 seconds
  Forward delay                       : 15 seconds
  Hop count                           : 19
  Number of topology changes          : 2
  Time since last topology change     : 26 seconds
  Local parameters
    Bridge ID                         : 4098.cc:e1:94:d2:6e:a1

{master:0}[edit]
lab@ex1# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface                   Port ID    Designated         Designated         Port
State  Role
                                       port ID            bridge ID          Cost
ae10                         128:13      128:13   4096.cce194d26ea1        10000    FWD
DESG
ge-0/0/5                    128:493     128:493   4096.cce194d26ea1        20000    FWD
DESG
ge-0/0/7                    128:495     128:495   4096.cce194d26ea1        20000    FWD
DESG
ge-0/0/9                    128:497     128:497   4096.cce194d26ea1        20000    FWD
DESG
ge-0/0/1                    128:504     128:504   4096.cce194d26ea1        20000    FWD
DESG
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/2                    128:505      128:505   4096.cce194d26ea1        20000    FWD
DESG

Spanning tree interface parameters for instance 1

Interface                   Port ID    Designated         Designated         Port
State  Role
                                       port ID            bridge ID          Cost
ae10                         128:13      128:13   4097.cce194d26ea1        10000    FWD
DESG
ge-0/0/5                    128:493      128:493   4097.cce194d26ea1        20000    FWD
DESG
ge-0/0/7                    128:495      128:495   4097.cce194d26ea1        20000    FWD
DESG
ge-0/0/1                    128:504      128:504   4097.cce194d26ea1        20000    FWD
DESG
ge-0/0/2                    128:505      128:505   4097.cce194d26ea1        20000    FWD
DESG

Spanning tree interface parameters for instance 2

Interface                   Port ID    Designated         Designated         Port
State  Role
                                       port ID            bridge ID          Cost
ae10                         128:13      128:13   4098.b8c253979801        10000    FWD
ROOT
ge-0/0/7                    128:495      128:495   4098.cce194d26ea1        20000    FWD
DESG
ge-0/0/9                    128:497      128:497   4098.cce194d26ea1        20000    FWD
DESG
ge-0/0/1                    128:504      128:500   4098.b8c253979801        20000    BLK
ALT
ge-0/0/2                    128:505      128:501   4098.b8c253979801        20000    BLK
ALT
```

EX2:

```
{master:0}[edit]
lab@ex2# run show spanning-tree bridge
STP bridge parameters
Routing instance name               : GLOBAL
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 4096.cc:e1:94:d2:6e:a1
  Root cost                         : 0
  Root port                         : ae10
  CIST regional root                : 4096.cc:e1:94:d2:6e:a1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  CIST internal root cost              : 10000
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 19
  Message age                          : 0
  Number of topology changes           : 2
  Time since last topology change      : 176 seconds
  Local parameters
    Bridge ID                          : 32768.b8:c2:53:97:98:01

STP bridge parameters for MSTI 1
  MSTI regional root                   : 4097.cc:e1:94:d2:6e:a1
  Root cost                            : 9998
  Root port                            : ge-0/0/2
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Hop count                            : 19
  Number of topology changes           : 2
  Time since last topology change      : 173 seconds
  Local parameters
    Bridge ID                          : 32769.b8:c2:53:97:98:01

STP bridge parameters for MSTI 2
  MSTI regional root                   : 4098.b8:c2:53:97:98:01
  Hello time                           : 2 seconds
  Maximum age                          : 20 seconds
  Forward delay                        : 15 seconds
  Number of topology changes           : 2
  Time since last topology change      : 176 seconds
  Local parameters
    Bridge ID                          : 4098.b8:c2:53:97:98:01

{master:0}[edit]
lab@ex2# run show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface                   Port ID    Designated         Designated         Port
State  Role
                                       port ID            bridge ID          Cost
ae10                         128:13      128:13   4096.cce194d26ea1        10000     FWD
ROOT
ge-0/0/3                    128:492     128:492   32768.b8c253979801        20000     FWD
DESG
ge-0/0/5                    128:494     128:494   32768.b8c253979801        20000     FWD
DESG
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/9                  128:495      128:495   32768.b8c253979801        20000    FWD
DESG
ge-0/0/1                  128:500      128:504    4096.cce194d26ea1        20000    BLK
ALT
ge-0/0/2                  128:501      128:505    4096.cce194d26ea1        19999    BLK
ALT

Spanning tree interface parameters for instance 1

Interface                 Port ID    Designated         Designated         Port
State  Role
                                       port ID          bridge ID          Cost
ae10                       128:13       128:13   4097.cce194d26ea1        10000    BLK
ALT
ge-0/0/3                  128:492      128:492   32769.b8c253979801        20000    FWD
DESG
ge-0/0/5                  128:494      128:494   32769.b8c253979801        20000    FWD
DESG
ge-0/0/9                  128:495      128:495   32769.b8c253979801        20000    FWD
DESG
ge-0/0/1                  128:500      128:504    4097.cce194d26ea1         9999    BLK
ALT
ge-0/0/2                  128:501      128:505    4097.cce194d26ea1         9998    FWD
ROOT

Spanning tree interface parameters for instance 2

Interface                 Port ID    Designated         Designated         Port
State  Role
                                       port ID          bridge ID          Cost
ae10                       128:13       128:13   4098.b8c253979801        10000    FWD
DESG
ge-0/0/9                  128:495      128:495   4098.b8c253979801         20000    FWD
DESG
ge-0/0/1                  128:500      128:500   4098.b8c253979801         20000    FWD
DESG
ge-0/0/2                  128:501      128:501   4098.b8c253979801         19999    FWD
DESG
```

Configurations for Task Completion:

EX1:

```
delete protocols rstp
set protocols mstp configuration-name jncie
set protocols mstp revision-level 1
set protocols mstp bridge-priority 4k
set protocols mstp bpdu-block-on-edge
set protocols mstp interface ge-0/0/1
set protocols mstp interface ge-0/0/2
set protocols mstp interface access-ports edge
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols mstp interface ae10
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan 100-199
set protocols mstp msti 2 vlan 200-299
set interfaces interface-range access-ports member ge-0/0/5
set interfaces interface-range access-ports member ge-0/0/7
set interfaces interface-range access-ports member ge-0/0/9
```

EX2:

```
delete protocols rstp
set protocols mstp configuration-name jncie
set protocols mstp revision-level 1
set protocols mstp bpdu-block-on-edge
set protocols mstp interface ge-0/0/1
set protocols mstp interface ge-0/0/2 cost 19999
set protocols mstp interface access-ports edge
set protocols mstp interface ae10
set protocols mstp msti 1 vlan 100-199
set protocols mstp msti 1 interface ge-0/0/1 cost 9999
set protocols mstp msti 1 interface ge-0/0/2 cost 9998
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan 200-299
set interfaces interface-range access-ports member ge-0/0/3
set interfaces interface-range access-ports member ge-0/0/5
set interfaces interface-range access-ports member ge-0/0/9
```

Task 3: Layer 2 Features

Note: For this task you will not be able to verify your configurations until you have configured the network for routing. You will configure IGP routing in Chapter 4 and BGP in Chapter 5.

1) Ensure Server1 and Server 2 receive an IPv4 address from the DHCP server.

EX1:

```
lab@ex1# show forwarding-options
dhcp-relay {
    overrides;
    server-group {
        dhcp-server {
            10.10.10.1;
        }
    }
    group dhcp-server {
        active-server-group dhcp-server;
        interface irb.101;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
```

```
lab@ex1# show interfaces ge-0/0/6
description server1;
unit 0 {
    family inet {
        dhcp;
    }
    family inet6 {
        address 2001:0db8:0100:101::1/64;
    }
}
```

EX2:

```
lab@ex2# show interfaces ge-0/0/4
description server;
unit 0 {
    family inet {
        dhcp;
    }
}

lab@ex2# show interfaces irb
unit 100 {
    family inet {
        address 10.100.100.253/24;
    }
}

{master:0}[edit]
lab@ex2# show routing-options
static {
    route 10.10.10.0/24 next-hop 10.100.100.254;
}
```

Verification:

Note: You cannot verify this task at this point. You must have routing enabled throughout your network in order for dhcp requests to be forwarded to the DHCP server at 10.10.10.1. You can verify the proper behavior after completing Chapter 5. The following outputs are after completing Chapter 5 tasks.

EX2:

```
lab@ex2# run show interfaces ge-0/0/4.0 | find Local
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Destination: 10.100.101/24, Local: 10.100.101.51, Broadcast: 10.100.101.255
```

Note: You might need to deactivate and activate the ge-0/0/4 interface to trigger DHCP requests.

2) Printer1 must be authenticated dynamically, but EX1 must not attempt 802.1x authentication. If authentication fails or the server is unreachable, the printer must be given guest access.

Note: You can use the RADIUS server at 10.10.10.1 using the secret 'juniper123'. You should source the requests from the IRB interface that is assigned to the printer VLAN.

EX1:

```
lab@ex1# show access
radius-server {
    10.10.10.1 {
        secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"; ## SECRET-DATA
        source-address 10.100.200.254;
    }
}
profile jncie {
    authentication-order radius;
    radius {
        authentication-server 10.10.10.1;
    }
}

lab@ex1# show protocols dot1x
authenticator {
    authentication-profile-name jncie;
    interface {
        ge-0/0/9.0 {
            mac-radius {
                restrict;
            }
            guest-vlan guest;
            server-fail vlan-name guest;
        }
    }
}
```

Verification:

Note: You cannot verify this task at this point. You must have routing enabled throughout your network in order for traffic to reach the RADIUS server at 10.10.10.1. You can verify

© 2020 Juniper Networks, Inc. All rights reserved.

the proper behavior after completing Chapter 5. The following outputs are after completing Chapter 5 tasks.

Log in to the printer VM to initiate traffic to EX1's IRB interface in the printer VLAN.

Printer:

```
-bash-3.2$ ping 10.100.200.254
PING 10.100.200.254 (10.100.200.254) 56(84) bytes of data.
64 bytes from 10.100.200.254: icmp_seq=1 ttl=64 time=1003 ms
64 bytes from 10.100.200.254: icmp_seq=2 ttl=64 time=3.02 ms
64 bytes from 10.100.200.254: icmp_seq=3 ttl=64 time=2.35 ms

--- 10.100.200.254 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 2.351/336.355/1003.687/471.875 ms, pipe 2
-bash-3.2$
```

This will initiate traffic which should cause EX1 to authenticate the MAC address with the RADIUS server.

EX1:

```
{master:0}[edit]
lab@ex1# run show dot1x interface
802.1X Information:
Interface     Role           State          MAC address          User
ge-0/0/9.0    Authenticator  Authenticated  00:50:56:00:00:04    005056000004
```

3) Configure interface ge-0/0/7 on EX1 and interface ge-0/0/9 on EX2 for POE and LLDP-MED support. Block location-id TLVs. No more than 20 WATT should be allocated the POE ports.

EX1:

```
lab@ex1# show poe
interface ge-0/0/7 {
    maximum-power 20;
}

{master:0}[edit]
lab@ex1# show protocols lldp-med
interface ge-0/0/7 {
    tlv-filter location-id;
}
```

EX2:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@ex2# show poe
interface ge-0/0/9 {
    maximum-power 20;
}

{master:0}[edit]
lab@ex2# show protocols lldp-med
interface ge-0/0/9 {
    tlv-filter location-id;
}
```

4) If more than 10 MAC addresses are learned on EX2 port ge-0/0/9 it should drop traffic.

EX2:

```
lab@ex2# show switch-options interface ge-0/0/9.0
interface-mac-limit {
    10;
    packet-action drop;
}
```

5) Configure OAM LFM on the ge-0/0/1 and ge-0/0/2 interfaces between EX1 and EX2. Make sure EX1 and EX2 initiate the discovery process. Ensure OAM PDUs are sent every 600ms. If 5 PDUs are missed, the peer should be declared lost/down.

EX1 & EX2:

```
lab@ex1# show protocols oam
ethernet {
    link-fault-management {
        interface ge-0/0/1 {
            pdu-interval 600;
            link-discovery active;
            pdu-threshold 5;
        }
        interface ge-0/0/2 {
            pdu-interval 600;
            link-discovery active;
            pdu-threshold 5;
        }
    }
}
```

6) Configure EX1 to shutdown a port if Unknown Unicast or Broadcast traffic consumes more than 5% of the interface bandwidth. Multicast traffic should not be included in your solution.

© 2020 Juniper Networks, Inc. All rights reserved.

EX1:

```
lab@ex1# show forwarding-options
storm-control-profiles default {
    all {
        bandwidth-percentage 5;
        no-multicast;
    }
    action-shutdown;
}
```

Configurations for Task Completion:

EX1:

```
set forwarding-options dhcp-relay overrides
set forwarding-options dhcp-relay server-group dhcp-server 10.10.10.1
set forwarding-options dhcp-relay group dhcp-server active-server-group dhcp-server
set forwarding-options dhcp-relay group dhcp-server interface irb.101
set interfaces ge-0/0/6 unit 0 family inet dhcp
set access radius-server 10.10.10.1 secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"
set access radius-server 10.10.10.1 source-address 10.100.200.254
set access profile jncie authentication-order radius
set access profile jncie radius authentication-server 10.10.10.1
set protocols dot1x authenticator authentication-profile-name jncie
set protocols dot1x authenticator interface ge-0/0/9.0 mac-radius restrict
set protocols dot1x authenticator interface ge-0/0/9.0 guest-vlan guest
set protocols dot1x authenticator interface ge-0/0/9.0 server-fail vlan-name guest
set poe interface ge-0/0/7 maximum-power 20
set protocols lldp-med interface ge-0/0/7 tlv-filter location-id
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-interval 600
set protocols oam ethernet link-fault-management interface ge-0/0/1 link-discovery
active
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-threshold 5
set protocols oam ethernet link-fault-management interface ge-0/0/2 pdu-interval 600
set protocols oam ethernet link-fault-management interface ge-0/0/2 link-discovery
active
set protocols oam ethernet link-fault-management interface ge-0/0/2 pdu-threshold 5
set forwarding-options storm-control-profiles default all bandwidth-percentage 5
set forwarding-options storm-control-profiles default all no-multicast
set forwarding-options storm-control-profiles default action-shutdown
```

EX2:

```
set interfaces ge-0/0/4 unit 0 family inet dhcp
set interfaces irb unit 100 family inet address 10.100.100.253/24
set routing-options static route 10.10.10.0/24 next-hop 10.100.100.254
set poe interface ge-0/0/9 maximum-power 20
set protocols lldp-med interface ge-0/0/9 tlv-filter location-id
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set switch-options interface ge-0/0/9.0 interface-mac-limit 10 packet-action drop
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-interval 600
set protocols oam ethernet link-fault-management interface ge-0/0/1 link-discovery
active
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-threshold 5
set protocols oam ethernet link-fault-management interface ge-0/0/2 pdu-interval 600
set protocols oam ethernet link-fault-management interface ge-0/0/2 link-discovery
active
set protocols oam ethernet link-fault-management interface ge-0/0/2 pdu-threshold 5
```

Appendix – Chapter 3: Protocol Independent Features

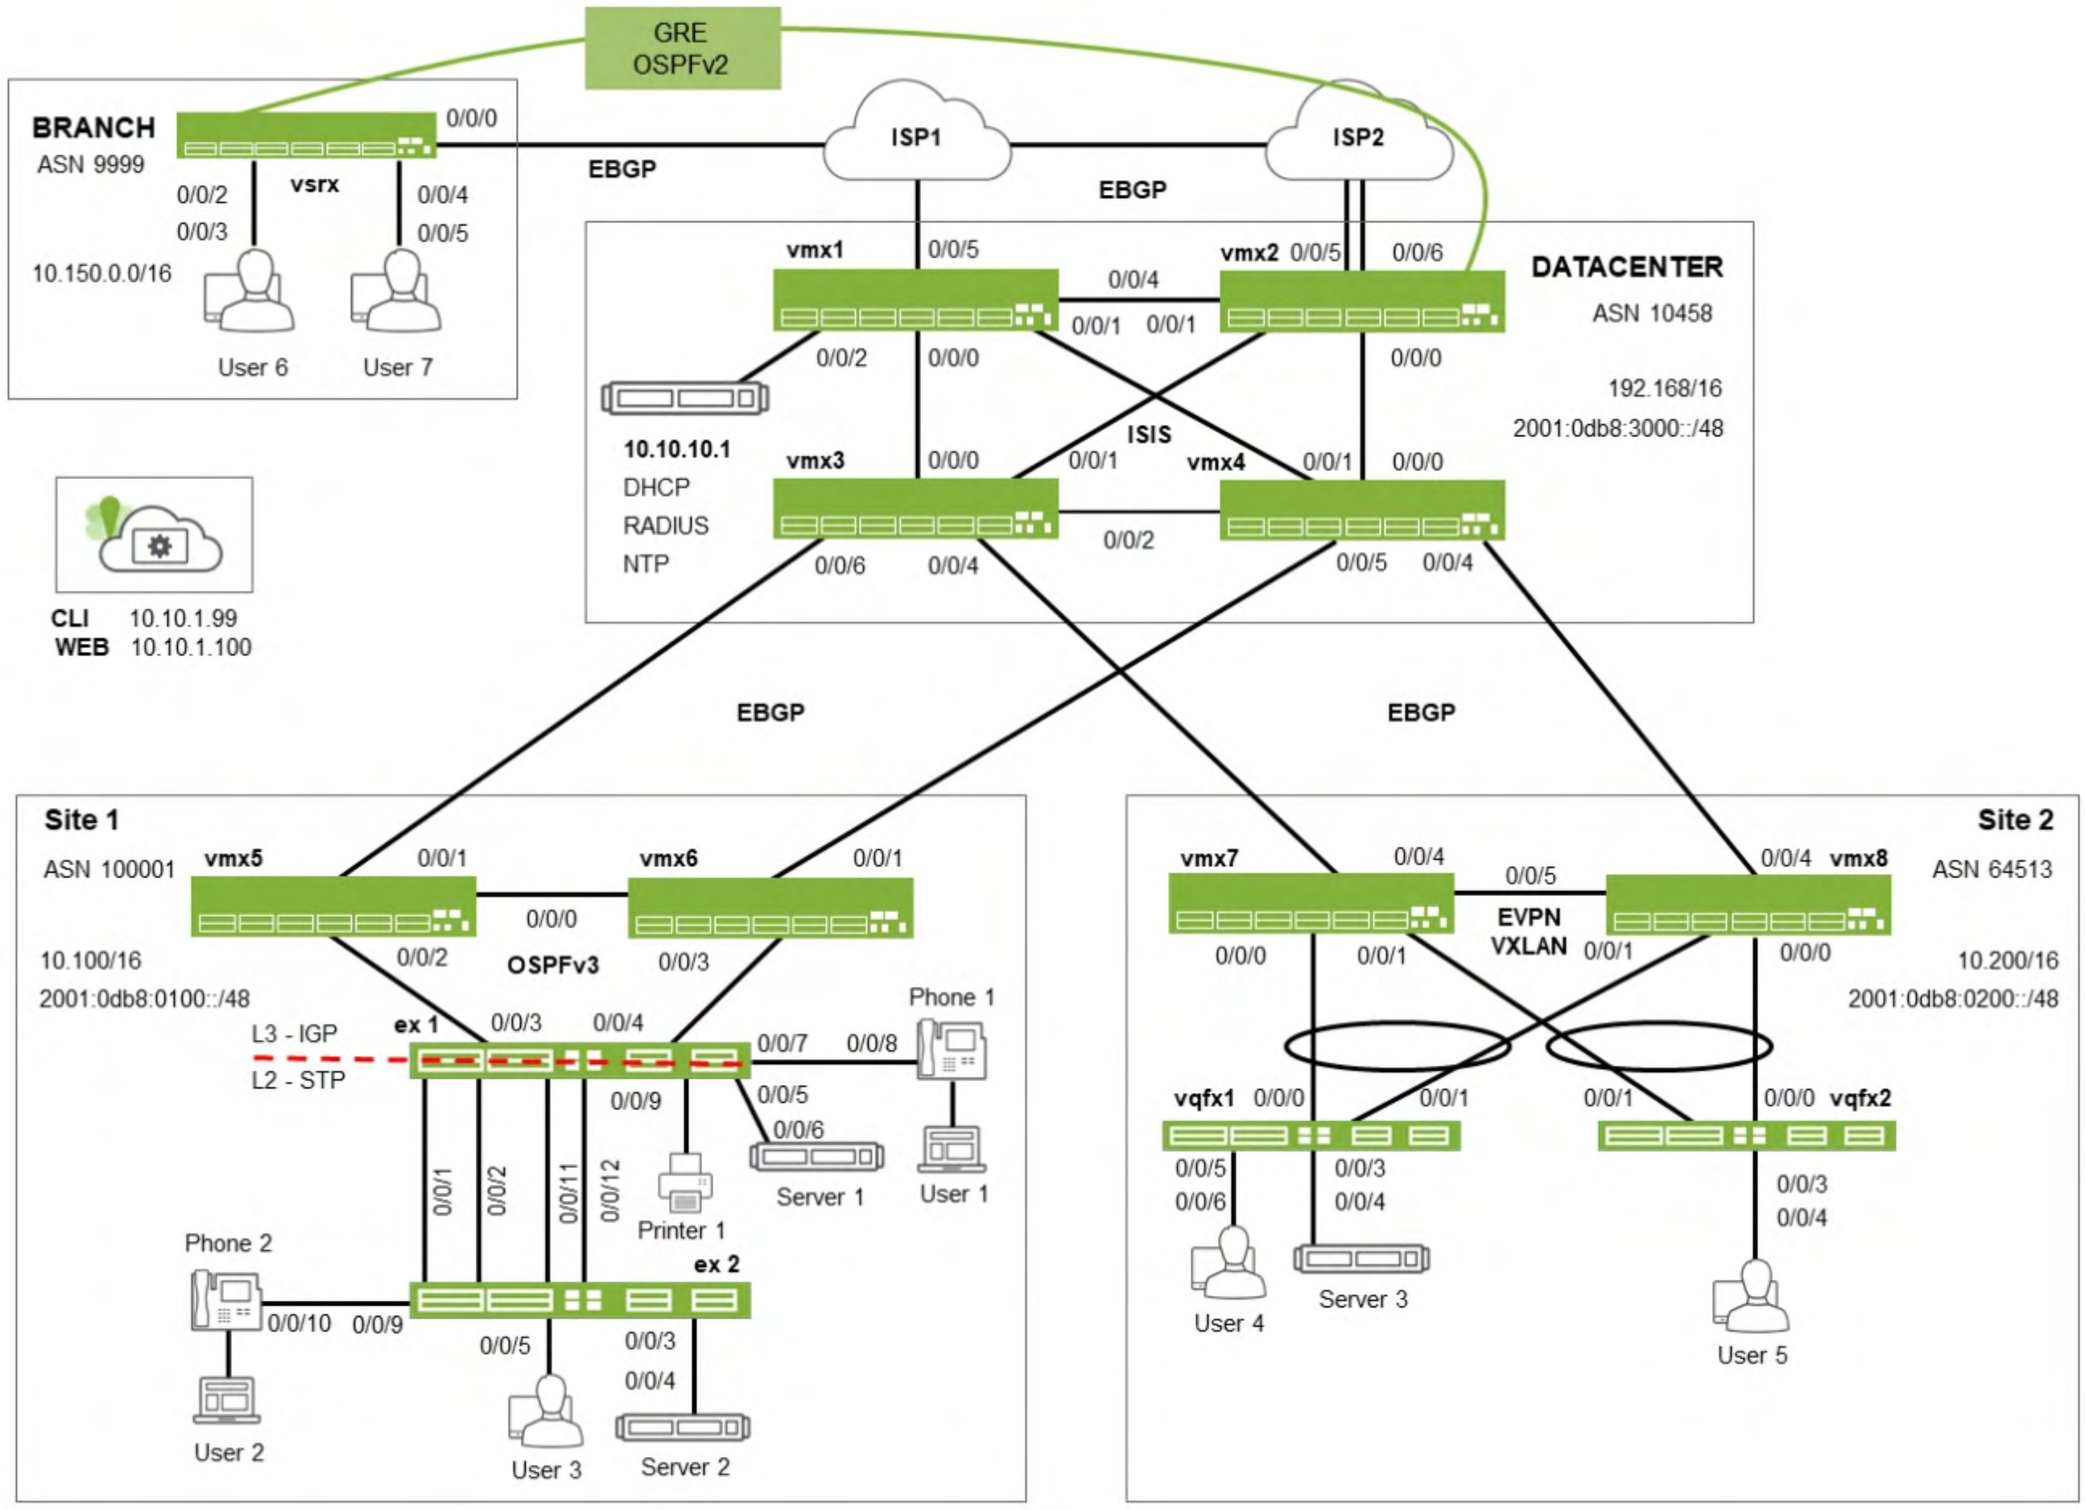

Task 1: GRE Tunnel

In this task you will configure a GRE tunnel between the data center and the Branch office.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure a GRE tunnel between the branch office and vMX2 in the data center. Use the IP addressing as shown in the Lab setup Information chapter. Use the loopback interface as GRE tunnel source address.

Note: For this task you will not be able to verify your tunnel is established until you have configured the BGP routing from vMX2 through the internet. You will configure this functionality in Chapter 5.

vMX2:

```
lab@vmx2# show chassis
fpc 0 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
    }
    lite-mode;
}

lab@vmx2# run show interfaces terse | match gr
gr-0/0/10               up    up

lab@vmx2# show interfaces gr-0/0/10
no-traps;
unit 0 {
    tunnel {
        source 192.168.0.2;
        destination 10.150.0.99;
    }
    family inet {
        address 173.17.99.2/24;
    }
}
```

vSRX:

```
lab@vsrx# show interfaces gr-0/0/0
unit 0 {
    tunnel {
        source 10.150.0.99;
        destination 192.168.0.2;
    }
    family inet {
        address 173.17.99.99/24;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vsrx# show routing-options
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
```

Verification:

Note: You cannot verify this task at this point. You must have routing enabled through your ISP networks. You can verify the proper behavior after completing Chapter 5. The following outputs are after completing Chapter 5 tasks.

vMX2:

```
lab@vmx2# run show interfaces terse | match gr-
gr-0/0/10               up    up
gr-0/0/10.0             up    up   inet     173.17.99.2/24

[edit]
lab@vmx2# run ping 173.17.99.99
PING 173.17.99.99 (173.17.99.99): 56 data bytes
64 bytes from 173.17.99.99: icmp_seq=0 ttl=64 time=200.979 ms
64 bytes from 173.17.99.99: icmp_seq=1 ttl=64 time=1.359 ms
64 bytes from 173.17.99.99: icmp_seq=2 ttl=64 time=1.301 ms
^C
--- 173.17.99.99 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.301/67.880/200.979/94.115 ms
```

Configurations for Task Completion:

vMX2:

```
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set interfaces gr-0/0/10 no-traps
set interfaces gr-0/0/10 unit 0 tunnel source 192.168.0.2
set interfaces gr-0/0/10 unit 0 tunnel destination 10.150.0.99
set interfaces gr-0/0/10 unit 0 family inet address 173.17.99.2/24
```

vSRX:

```
set interfaces gr-0/0/0 unit 0 tunnel source 10.150.0.99
set interfaces gr-0/0/0 unit 0 tunnel destination 192.168.0.2
set interfaces gr-0/0/0 unit 0 family inet address 173.17.99.99/24
set routing-options static route 192.168.0.2/32 next-hop 10.1.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Filter-Based Forwarding

1) Ensure that vMX2 forwards IPv4 web traffic received from Site1 destinated to 6.83.176.1 using the ge-0/0/6 interface. If interface ge-0/0/6 is unavailable, then interface ge-0/0/5 must be used. Other traffic should follow the routing requirements as mentioned in the other tasks.

vMX2:

```
lab@vmx2# show routing-options
interface-routes {
    rib-group inet interfaces;
}
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 ];
    }
}

lab@vmx2# show firewall filter fbf-v4
apply-flags omit;
term 1 {
    from {
        source-address {
            10.100.0.0/16;
        }
        destination-address {
            6.83.176.1/32;
        }
        destination-port [ 80 443 ];
    }
    then {
        routing-instance fbf;
    }
}
term 2 {
    then accept;
}


lab@vmx2# show routing-instances
fbf {
    instance-type forwarding;
    routing-options {
        static {
            route 0.0.0.0/0 {
                next-hop 172.16.3.9;
                qualified-next-hop 172.16.2.9 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    preference 6;
                }
            }
        }
    }
}


lab@vmx2# show interfaces ge-0/0/0
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.24.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0024::2/64;
    }
}

[edit]
lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.23.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
}
```

```
[edit]
lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

Verification:

```
lab@vmx2# run show route table fbf.inet

fbf.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:14:41
                    > to 172.16.3.9 via ge-0/0/6.0
                    [Static/6] 00:14:41
                    > to 172.16.2.9 via ge-0/0/5.0
...
```

Configurations for Task Completion:

vMX2:

```
set routing-options interface-routes rib-group inet interfaces
set routing-options rib-groups interfaces import-rib [ inet.0 fbf.inet.0 ]
set firewall filter fbf-v4 apply-flags omit
set firewall filter fbf-v4 term 1 from source-address 10.100.0.0/16
set firewall filter fbf-v4 term 1 from destination-address 6.83.176.1/32
set firewall filter fbf-v4 term 1 from destination-port [80 443]
set firewall filter fbf-v4 term 1 then routing-instance fbf
set firewall filter fbf-v4 term 2 then accept
set routing-instances fbf instance-type forwarding
set routing-instances fbf routing-options static route 0.0.0.0/0 next-hop 172.16.3.9
set routing-instances fbf routing-options static route 0.0.0.0/0 qualified-next-hop
172.16.2.9 preference 6
set interfaces ge-0/0/0 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/1 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/4 unit 0 family inet filter input fbf-v4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Chapter 4: IGP

This chapter focuses on the implementation of various IGPs like IS-IS, OSPFv2 and OSPv3.

OSPF Overview

OSPFv2 is a link-state protocol that can be used to distribute IPv4 routing-information. Link State Advertisements (LSAs) are used to reliable flood routing information across the AS. These LSAs are stored in an LSDB. The SPF algorithm is run against the records in this database and the outcome is placed into the routing-table.

OSPFv2 uses the following packets:

1) **Hello** – used to dynamically discover, establish and maintain neighbor relationships. Send to 224.0.0.5 (or all DRs 224.0.0.6) on a regular interval.
2) **Database Description** – used during neighbor relationship establishment on multiaccess links. Packets transmit LSA headers between OSPF routers.
3) **Link-State Request** – used to request missing LSAs or LSAs that are not up-to-date.
4) **Link-State Update** – send in response to LSR or link state change. Carries a collection of LSAs.
5) **Link- State Acknowledgement** – every flooded LSA is acknowledged with an LSAck, providing for a reliable flooding.

Inside the LSDB, we find different types of LSAs. These types are the following:

Type 1: Router-LSA.

Description of the state and costs of the routers' link inside an area. Originated for each area that a router belongs to and flooded throughout the entire area.

Type 2: Network-LSA.

Originated by the DR, describing the networks connected to a multi-access segment. Flooded throughout the area.

Type 3: Network-Summary LSA.

Describes inter-area routes to networks in other areas, originated by the ABR. These LSAs are flooded throughout the area.

© 2020 Juniper Networks, Inc. All rights reserved.

Type 4: ASBR-Summary LSA.

Originated by an ABR attached to an area that contains an ASBR. Describes inter-area routes to AS boundary routers. Basically, the ABR telling other areas, or the rest of the OSPF domain, where they can find a particular ASBR.

Type 5: External LSA.

Originated by ASBRs, describes the routes towards external destinations that have been redistributed into OSPF.

During the formation of an OSPF neighbor relationship, the router will go through various states. The description of the different neighbor states (or conversation states) is as follows;

- **Down**: nothing happened. OSPF is awaiting a start event. When there is a start event, a Hello is sent towards the neighbor.
- **Init**: Hello is seen but it did not contain the router's own router-id.
- **2Way**: Bidirectional communication is established. Both routers see their own router-id listed in the 'neighbor' field in the received hello packets.
- **ExStart**: The routers determine who is slave and who is master during the following exchange state.
- **Exchange**: The master and slave exchange LSA headers, describing their databases to each other. LSR's and LSU's may also be sent during this state in case a router is missing some LSA's.
- **Loading**: Link state information is exchanged in the form of LSR's and LSU's.
- **Full**: When the LSR list is empty, the state is changed to Full.

OSPFv3 was developed to offer support for both IPv6 and IPv4. To quote RFC 5340: 'The fundamental mechanisms of OSPF (flooding, Designated Router (DR) election, area support, (Shortest Path First) SPF calculations, etc.) remain unchanged.'

In the [**protocols ospf3**] hierarchy, you can configure the IPv6 realm. Under the [**protocols ospf3 realm ipv4-unicast**], you can configure the IPv4 realm. In the IPv4 realm, you will find that some features (like plain or MD5 authentication, virtual link and traffic engineering) are not supported.

Apart from the address family (or realm) support, there are some additional differences between OSPFv2 and OSPFv3:

- In OSPFv3, authentication can only be done using IPsec.

© 2020 Juniper Networks, Inc. All rights reserved.

- OSPF3 has a different packet format and it introduces new LSA types (some of which have a new flooding scope):

OSPFv3 LSA type - function code:	OSPFv3 LSA name:	OSPFv2 LSA type:	OSPFv2 LSA name:
1 (Area scope)	Router-LSA	1	Router-LSA
2 (Area scope)	Network-LSA	2	Network-LSA
3 (Area scope)	Inter-Area-Prefix-LSA	3	Summary-LSA
4 (Area scope)	Inter-Area-Router-LSA	4	ASBR-Summary-LSA
5 (AS scope)	AS-External-LSA	5	AS-External-LSA
6 (Area scope)	Deprecated		
7 (Area scope)	NSSA-LSA	7	NSSA External LSA
8 (link-local)	Link-LSA		
9 (Area scope)	Intra-Area-Prefix-LSA		

- OSPFv3 does protocol processing per-link instead of per subnet
- There is explicit support for multiple instances of OSPF3 on a single link
- The router-ID is still a 32-bit number, but it must be configured manually

Very often, you'll find OSPFv2 configured for IPv4 and OSPFv3 configured for IPv6.

Multi-Area OSPF Implementation

To be able to optimize and scale a network running OSPF, it can be divided into smaller groups called areas. An area is a collection of networks that are administratively grouped together. A separate LSDB is maintained for every single area and the LSDB describing an area should be identical on all the routers inside that area.

Through the configuration of these different areas, LSA flooding is restrained. Furthermore, by dividing the network into different areas, you will be able to summarize prefixes. Summarization can have two positive effects on the network. It can reduce the size of the LSDB, and it provide for extra stability. The extra stability is achieved by prevent routing-instability in one area, or a network external to OSPF, from affecting another area's LSDB.

A multi-area OSPF network consists of several different areas that must be directly connected to the backbone area. Routers connected to multiple areas are called Area Border Routers. These ABRs can transmit network information between different OSPF areas. Routers injecting information from outside of the OSPF domain (i.e. static, RIP, etc.) are called Autonomous System Boundary Routers. By default, an area is considered a 'regular' or 'normal' area and supports the flooding of LSA type 1, 2, 3, 4 and 5.

© 2020 Juniper Networks, Inc. All rights reserved.

Apart from the 'regular' areas, there are a number of other stub area types to consider:

Stub Area: the regular stub area is void of any LSA type 4 and 5. Instead of having the external routing-information, stub routers rely on a default route from the ABRs inside the stub area. The stub area cannot contain an ASBR and only LSA type 1, 2 and 3 are flooded throughout the area.

Totally Stubby Area: totally stubby areas have their LSDBs further reduced by relying solely on the default route generated by the ABR.

Not So Stubby Area (NSSA): it is possible for routers inside the NSSA to be an ASBR and inject external routing-information into the OSPF domain. To redistribute external routing information into OSPF, the NSSA ASBR injects type 7 LSAs. These type 7 LSAs are translated into type 5 LSAs by the ABR that borders the non-NSSA areas. When there are multiple ABRs, the ABR with the highest router-id does the translation. External routing information from other areas is not send into the NSSA. Inside the NSSA area, we can find type 1, 2 3 and 7 LSAs.

Totally NSSA: the ABR for a totally NSSA will not inject type 3 LSAs, making the routers inside the area rely on the default route for all internal and external routes that are not generated from within the NSSA.

Converting a stub or NSSA area into a totally stub or NSSA area can be done by using the **no-summaries** command. When you inject a route into a totally stubby area or totally NSSA, you can have its metric and its metric type adjusted. When you play around with the metric values, always keep in mind that a type 1 route will always be preferred over a type 2 route.

Also note that because the different stub areas do not contain all routing information, a virtual-link cannot transit these types of areas.

Summarization and Route Filtering

We can hide network instability and increase the network scalability through summarization and filtering. OSPF imposes several restrictions on where you can perform summarization and route filtering. Because the LSDB has to be consistent on all routers within an area, you cannot summarize within an area. The general rule is that you can summarize where LSA and route conversion takes place (ABR/ASBR).

Within Junos, performing OSPF summarization and route-filtering can be done using the following options:

© 2020 Juniper Networks, Inc. All rights reserved.

- Inter-area internal LSA summarization and filtering on the ABR, using the **area-range** configuration command. The **restrict** option can be used to filter out the routes contained within the summary.
- Inter-area internal LSA filtering on the ABR using the **network-summary-import** and **network-summary-export** configuration command.
- External route summarization and filtering on the ASBR through the use of aggregate routes and export policies.
- Inter-area NSSA external route summarization and filtering on the NSSA ABR using the **nssa area-range** command.
- Route-table filtering of external routes using import policies.

Keep in mind that only type 1, 2 and 7 LSAs can be summarized or filtered. Type 5 LSAs cannot be summarized or filtered.

Redistributing Routes into OSPF

When you configure route-redistribution on an OSPF router, you turn the router into an ASBR. By default, the external routes redistributed into OSPF are external type 2 routes. The metric for these routes is never incremented. Through the use of a routing policy, you are able to change the route type. This is done using the **'then external type 1**' action. This will redistribute the matched routing information as external type 1. The external type 1 routes will have their metrics incremented when they are flooded throughout the network.

By using external type 1 and external type 2 routes when redistributing routing-information on multiple routers, you can easily create a primary and a backup path.

© 2020 Juniper Networks, Inc. All rights reserved.

Virtual link and Multi Area Adjacencies

All of the OSPF areas have to be connected to area 0. When an OSPF router does not have a physical connection to the backbone area, we can still connect the router to area 0 by means of a virtual link.

The virtual-link is configured in area 0 using the following configuration command:

[edit protocols ospf area 0.0.0.0]

virtual-link neighbor-id *router-id* transit-area *area-id*

The verification of the virtual-link can be done by issuing the **show ospf neighbor** and **show ospf interface** commands. The virtual-link can be configured with additional parameters such as authentication, altered hello interval, etc.

As stated previously, keep in mind that a virtual link cannot be established through a stub area.

To configure a multi-area adjacency, we need to add the **secondary** option to the interface in the OSPF configuration. The state of the secondary interface is always point-to-point. For more information about multi-area adjacency , please study RFC 5185.

OSPF Options and Considerations.

- To divert traffic away from a router, we can use the **overload** command. This will set the metric for all of the routers' transit links (except links into stub areas) to 0xFFFF (the highest value).
- Type 1 and 2 LSAs are summarized using the **area-range** command under **[edit protocols ospf area X].** The type 7 LSAs summarization is configured under the **[edit protocols ospf area X nssa]** hierarchy.
- By default, every link has a cost of 1. We can change this manually, link by link, using the **metric** command. We can also change the way OSPF calculates the metric by issuing the **reference-bandwidth** command. Using this command, we change the metric calculation to *cost=Ref bandwidth/interface bandwidth*.
- To avoid a DR and BDR election, we can configure the link to be point-to-point using the '**set protocols ospf area x interface x interface-type p2p'** configuration command. Avoiding a DR and BDR election speeds up the neighbor relationship formation. It also avoids the network from advertising type 2 LSAs.
- To make sure a router is ineligible to become either the DR or the BDR on a link, you can configure it with a priority 0.
- To inject a prefix that is configured on an interface without forming an OSPF neighbor relationship on that interface, configure it using the **passive** option.

© 2020 Juniper Networks, Inc. All rights reserved.

- The **no-nssa-abr** command disables exporting type 7 LSAs on ABR and ASBR alike.

Monitoring and troubleshooting OSPF

There are many reasons that can prevent an OSPF neighbor relationship to form between two routers:

1) Timer mismatch
2) Area number mismatch
3) Authentication mismatch (type of authentication, used passphrase, key-id in case of MD5)
4) Area type mismatch (stub area on one end versus 'regular' area on the other end)
5) Subnet mask mismatch
6) Duplicate router-id
7) MTU mismatch
8) A DR priority of 0 configured on both routers
9) Duplicate interface IP address
10) Interface type mismatch (p2p and p2mp for example)

To investigate the problems related to an OSPF neighbor relationship establishment, consider one of the following methods or commands:

- **monitor traffic interface *interface-name***

The monitor traffic interface command is used to display packets send to and from the routing-engine. Adding the keyword **detail** or **extensive** will give you additional information from the fields contained within the packet header. You can also narrow down the packets that you want to display by specifying a matching condition. For example, to see detailed OSPF information send and received across interface ge-0/0/4.123, issue the following command: **monitor traffic interface ge-0/0/4.123 matching "proto ospf" extensive**

- **traceoptions**

Traceoptions for OSPF are activated under the [**protocols ospf traceoptions**] or [**protocols ospf3 traceoptions**] hierarchy. To start a trace for general OSPF information, and to have the output stored in a file called 'ospf.log', you can configure the following:

```
set protocols ospf traceoptions file ospf.log
set protocols ospf traceoptions flag normal
```


© 2020 Juniper Networks, Inc. All rights reserved.

You can use the '**show log ospf.log | last**' command to see the latest messages logged to the file. To continuously follow any messages added to the log file, you can issue the '**monitor start ospf.log**' command.

- **show ospf(3) neighbor**

This command displays what neighbor relationships are formed across the interfaces running OSPF. It will also display the priorities for the interfaces on the other side, the router-ID of the neighboring router and the countdown timer until the neighbor is declared dead. Using the extensive option will also tell you how long the adjacency has been up.

- **show ospf(3) interface**

This command will output the options that the interface is configured for. It will list the interface type, the area that the interface is placed in, the DR and BDR for the link (0.0.0.0 or none for p2p links) and the number of neighbors active on the link.

Using the 'extensive' option will also inform you of the configured Hello and Dead interval, the retransmit interval, area number, area-type, authentication type and subnet-mask.

- **show interface terse**

Used to check the subnet that is configured on the interface and to verify whether or not the interface is up.

- **show interfaces ge-0/0/x.x | match MTU**

To verify the MTU that is applicable to the address family active on the interface.

IS-IS Overview

IS-IS is another Interior Gateway Protocol that uses link state information to build forwarding paths through the network. Like OSPF it uses the Shortest Path First (SPF) algorithm. IS-IS was originally developed for OSI routing and later adopted for routing IP. In fact, IS-IS is designed in a way that it is easily extendible. So, for example it's much easier to add new address families or protocol extensions then with OSPF.

As with OSPF a single IS-IS domain can be divided into smaller areas. Routing between areas is organized in a hierarchical fashion. Routers in IS-IS can be configured as Level 1, Level 2 or both. Level 1 routers route inside the area whereas Level 2 routers connect different areas. Level 1/Level 2 routers have interfaces participating in Level 1 and others in Level 2. Unlike OSPF, a router can be part of a single area only. L1/L2 routers send PDUs

© 2020 Juniper Networks, Inc. All rights reserved.

with the attached bit, announcing that they are a gateway to other areas. Areas affect the adjacency formation. Levels control the flooding of LSPs.

IS-IS uses hello PDUs (IIH PDUs). They are used to discover other IS-IS routers on the link. Once the adjacency is established IIH are used as keepalives between IS-IS neighbors. There are two types of IIH PDUs: LAN and point-to-point. Different LAN IIHs are used for Level 1 and Level 2 but are very similar. By default, hello PDUs are send every 9 seconds, except on broadcast networks where the DIS sends hellos every 3 seconds. Adjacency establishment is much simpler compared to OSPF. Neighbors start in the "Down" state, when the first message is received from a neighbor, they reach the "Initializing" state. When a bidirectional conversation is verified the neighbors move to state 'UP' state.

The following adjacency formation rules are used by IS-IS:

- Level 1 routers never form adjacencies with Level 2 routers.
- Level 2 routers never form adjacencies with Level 1 routers.
- The area ID must match for Level 1 adjacencies
- For Level 2 adjacencies Area IDs might be different.

As mentioned, on broadcast multi-access networks IS-IS elects a Designated Intermediate Systems (DIS). For DIS election the tie breaking rule is the highest MAC address in case the DIS priorities area equal. DIS priorities range from 0 to 127. The device with the highest priority wins the election. The default is 64 on all interfaces. A separate DIS is elected for Level 1 and Level 2. The DIS represents the broadcast segment. All routers on the segment synchronize their link state database with the DIS. There is no backup DIS elected as with OSPF. The other difference with OSPF is that a full mesh of adjacencies is formed on broadcast networks. The DIS election process starts all over again if new router is connected to the segment. If a new router has a higher priority, it becomes the new DIS.

IS-IS Link State PDUs (LSP) are exchanged and are used to build the link-state database. They are sent when a new adjacency is established, network changes or in response to sequence number PDUs. Based on the LSP in the database the SPF algorithm is executed. Separate LSPs are exchanged for Level 1 and Level 2.

Complete sequence number PDUs (CSNPs) contain a full description of the IS-IS link state database. They are periodically sent to all interfaces. IS-IS routers use them to update and synchronize their databases. On broadcast segment only the DIS sends CSNPs. The DIS also uses them for the acknowledgment of the received LSPs.

Partial sequence number PDUs (PSNPs) are used to request missing or outdated LSPs, and for the acknowledgment on point-to-point links.

© 2020 Juniper Networks, Inc. All rights reserved.

IS-IS supports two types of metrics (narrow and wide). Narrow metrics are in the range of 1 to 63 but are insufficient for modern networks. Wide metrics can range up to 16 million. To force the use of wide metrics you have configure the router to use them. It is recommended that all routers in the network have a consistent metric type. Like OSPF, the JUNOS IS-IS implementation supports a reference bandwidth. The configured value is divided by the bandwidth of the interface in bits per second. By default, all interfaces are assigned an IS-IS metric of 10. You can configure a manual metric on each IS-IS interface.

When enabling IS-IS on a router a network entity title (NET) address must be configured, which is a special ISO Network service access point (NSAP) address. The NET is a combination of the Area ID and the System ID (SysID). The NET is represented in hexadecimal. The NET format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier and a selector. The simplest format omits the domain ID and is 10 bytes long. As an example, the NET address 49.0001.1921.6800.1002.00 consists of the following parts:

49—AFI

0001—Area ID

1921.6800.1002—System identifier

00—Selector

Please note that the system identifier must be unique within the network. For IP networks, it's recommended to use the loopback IP address of the router. Configure a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first part of the address is the area number, which is a variable number from 1 to 13 bytes. The first byte of the area number, 49, is the authority and format indicator (AFI). The next bytes are the assigned area identifier and can be 0 to 12 bytes. In the examples, 0001 is the area identifier.

The NET address is configured under *family iso* and usually under the Loopback interface.

To include interfaces for IS-IS routing, you have to configure *family iso* together with adding it under the *protocol isis* hierarchy. As with OSPF, to configure IS-IS you have to specify the interfaces on which IS-IS adjacencies should be established. By default, both Level 1 and Level 2 adjacency formation is enabled. You can deactivate a particular Level if needed.

IS-IS supports simple and MD5 authentication. MD5 authentication uses an encrypted checksum for the PDUs. IS-IS authentication can be configured at different places in the

© 2020 Juniper Networks, Inc. All rights reserved.

configuration. IS-IS authentication is per Level. Adding authentication under the Level, authenticates every PDUs (Hellos, LSPs, CSNPs and PSNPs) generated by the router. Different types of authentication can be configured for Level 1 and Level 2. Adding authentication at the interface level under the **protocols isis** hierarchy, authenticates only the Hello messages transmitted on that interface.

Multilevel IS-IS networks have certain rules for the distribution of routing information between Levels. Flooding of LSPs is limited within the Level. L1/L2 routers control the distribution based on the following rules:

Internal routes:

- Level 1 routers flood LSPs to other L1 routers and L1/L2 routers.
- L1/L2 routers flood the L1 information to other L2 routers.
- L1/L2 routers do not flood L2 information to L1 routers.
- L2 routers send LSPs with the Attached bit set towards L1. L1 routers install a default route pointing to the L1/L2 routers.

External routes:

- L1 routers redistribute external information using policies. All L1 and L1/L2 routers receive the external information.
- L1/L2 routers are not distributing external routes outside L1. Except in one case, which is described later.
- L1/L2 routers do not distribute external routes to L1 routers. L1 routers use a default route to reach destinations outside of the L1 domain.

The default routing policy for IS-IS is similar to OSPF. As in OSPF, the information in the database should be identical on all routers in the same level, so import policies are not supported. There is difference with OSPF regarding the use of export policies. The default export policy permits local IS-IS information. This information is retrieved from the routing table.

JUNOS has different default route preferences for L1 internal, L1 external routes and for L2 internal and external routes.

- *IS-IS L1 internal route – 15*
- *IS-IS L2 internal route – 18*
- *IS-IS L1 external route – 160*
- *IS-IS L2 external route – 165*

© 2020 Juniper Networks, Inc. All rights reserved.

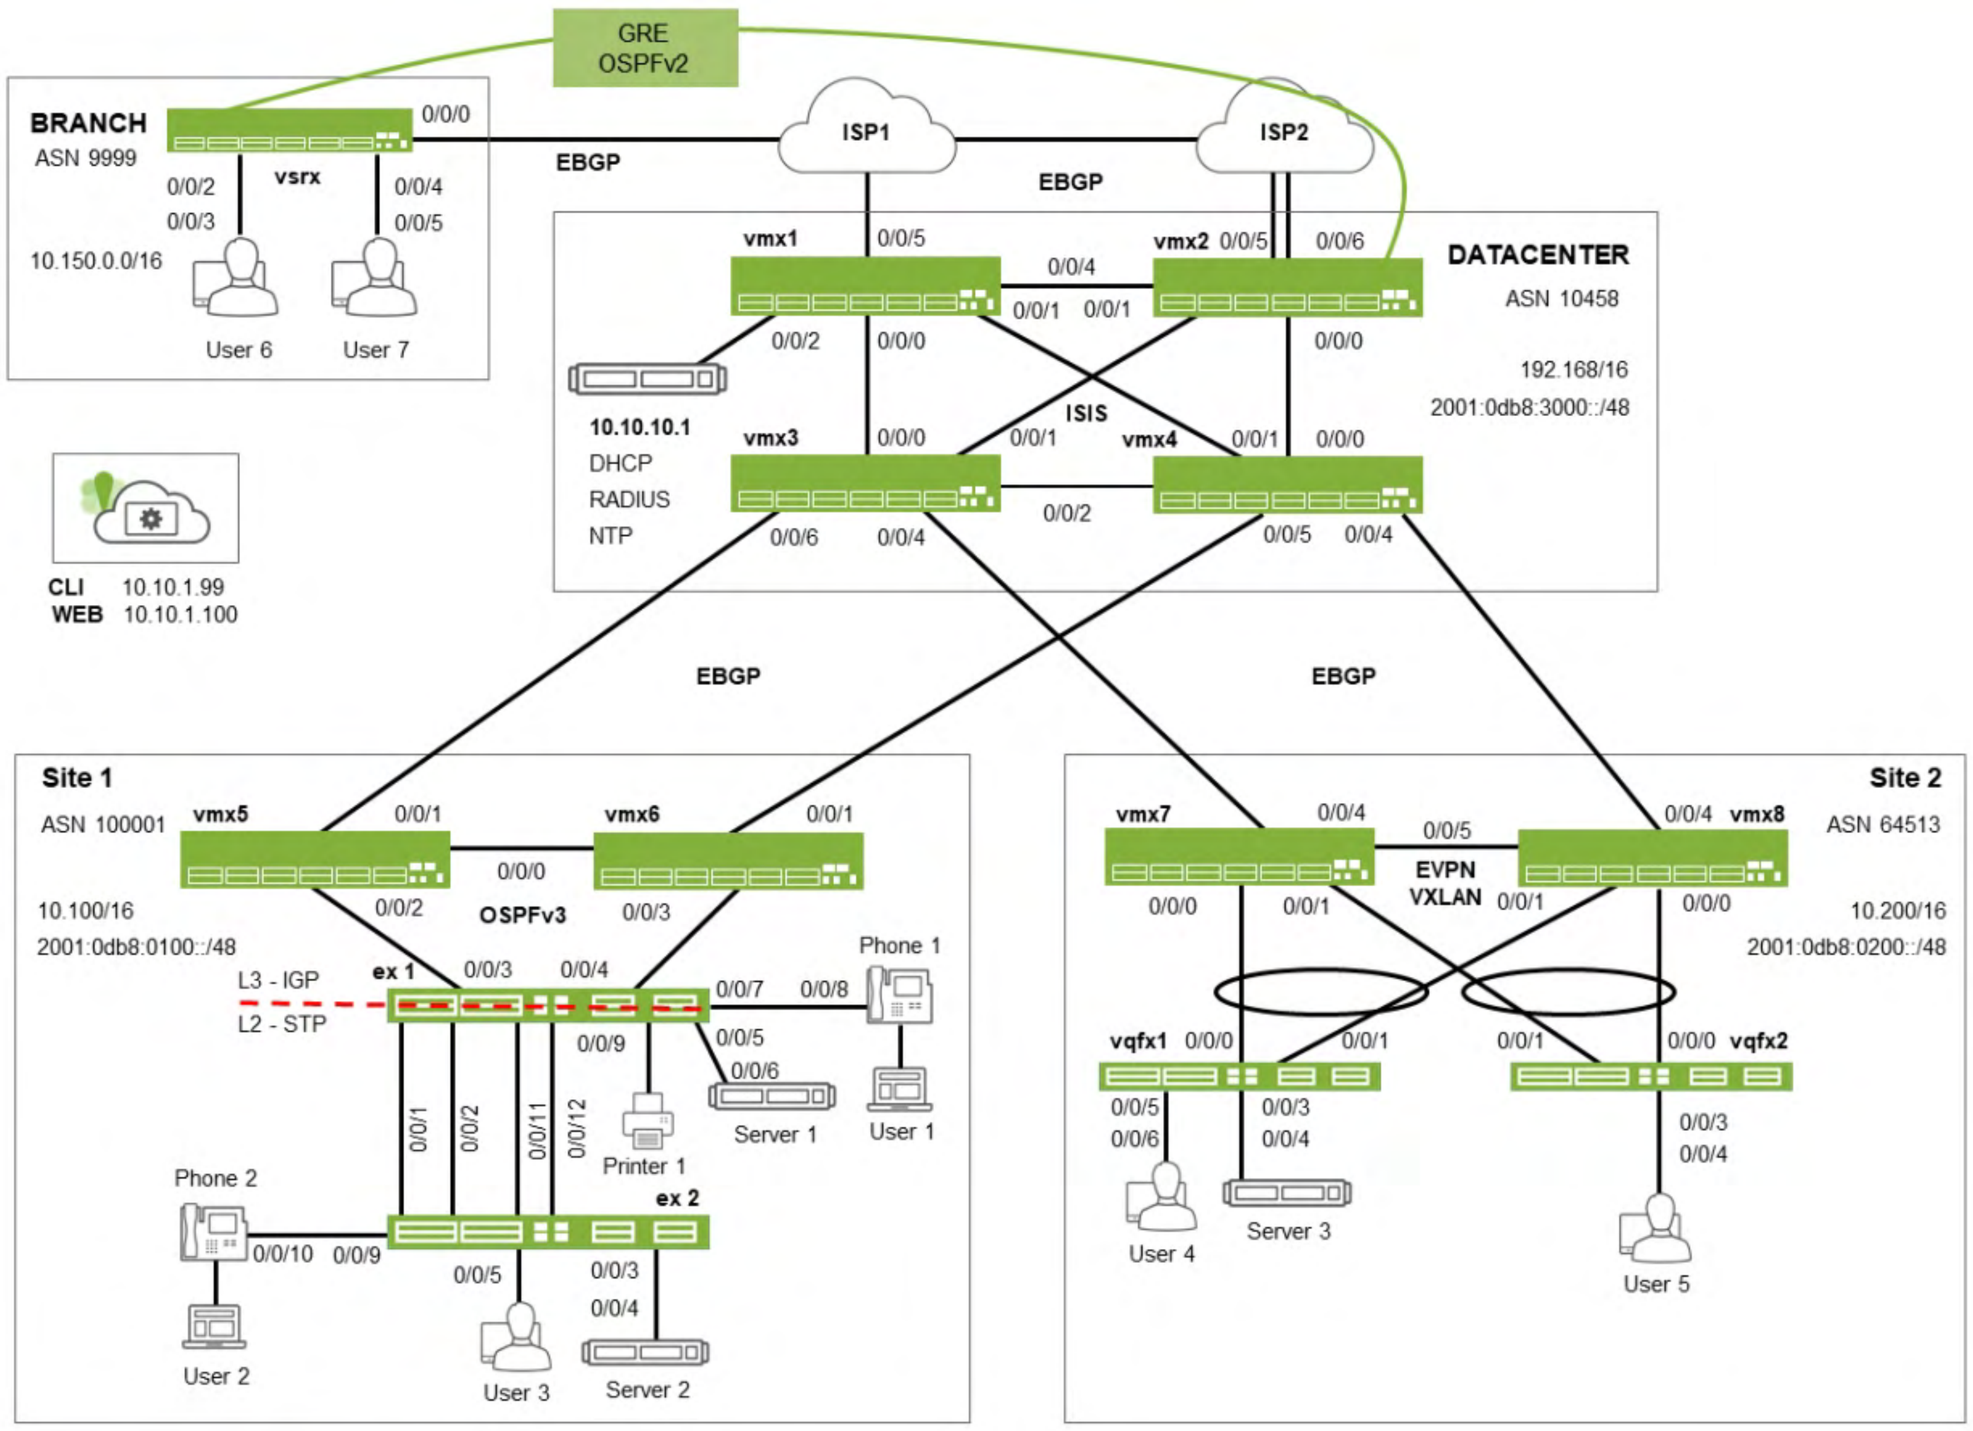

Task 1: OSPFv3

In this task you will configure OSPFv3 in Site 1 and Site 2.

Site 1:

Device	Interface	area
vMX5	lo0	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/2	0.0.0.1
vMX6	lo0	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/3	0.0.0.1
EX1	lo0	0.0.0.1
	ge-0/0/3	0.0.0.1

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/4	0.0.0.1

Table 32 OSPF information site 1

1) Configure OSPFv3 in Site 1 using the information as shown in the tables above.
2) Ensure type 0x2002 LSAs are not found in the OSPFv3 databases in Site 1.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four seconds.

vMX5:

```
lab@vmx5# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/0.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
    area 0.0.0.1 {
        interface ge-0/0/2.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}
```

vMX6:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx6# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/0.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
    area 0.0.0.1 {
        interface ge-0/0/3.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/3.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}
```

EX1:

```
lab@ex1# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.1 {
        interface ge-0/0/3.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
        interface ge-0/0/4.0 {
            interface-type p2p;
            hello-interval 1;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            dead-interval 4;
        }
        interface lo0.0 {
            passive;
        }
    }
}
area 0.0.0.1 {
    interface ge-0/0/3.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
    interface ge-0/0/4.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
    interface lo0.0 {
        passive;
    }
}
```

Verification:

EX1:

```
lab@ex1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
10.100.0.6      ge-0/0/4.0            Exchange  128      3
  Neighbor-address fe80::250:56ff:fea9:5c18

{master:0}[edit]
lab@ex1# run show ospf3 neighbor realm ipv4-unicast
ID              Interface             State     Pri   Dead
10.100.0.6      ge-0/0/4.0            Exchange  128      2
  Neighbor-address fe80::250:56ff:fea9:5c18
```

The output above shows that EX1 is unable to form an adjacency with vMX5 and vMX6. There is no neighborship state shown with vMX5. The neighborship state with vMX6 is in the “Exchange” state. This usually indicates an MTU issue.

vMX5:

```
lab@vmx5# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
10.100.0.6      ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4873
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx5# run show ospf3 neighbor realm ipv4-unicast
ID              Interface             State     Pri   Dead
10.100.0.6      ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4873
```

It seems there is no adjacency formed between vMX5 and EX1. The adjacency between vMX5 and vMX6 is in the “Full” state.

vMX6:

```
lab@vmx6# run show ospf3 neighbor
run show ospf3 neighbor realm ipv4-unicast ID              Interface             State
Pri   Dead
10.100.0.5      ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4cfd
10.100.0.5      ge-0/0/3.0            ExStart   128      3
  Neighbor-address fe80::cee1:94ff:fed2:6e47

[edit]
lab@vmx6# run show ospf3 neighbor realm ipv4-unicast
ID              Interface             State     Pri   Dead
10.100.0.5      ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4cfd
10.100.0.5      ge-0/0/3.0            ExStart   128      3
  Neighbor-address fe80::cee1:94ff:fed2:6e47
```

The output of vMX6 shows that indeed the adjacency between vMX5 and vMX6 is in the “Full” state”, however the adjacency with EX1 in in the “ExStart” state. What is also strange is the fact that the neighbor ID on the ge-0/0/3 link is the same as vMX5. It is clear that router ID 10.100.0.5 should belong to vMX5, so it looks like there is a configuration issue on EX1.

EX1:

```
lab@ex1# show routing-options
router-id 10.100.0.5;
```

Correct the router ID on EX1

EX1:

```
lab@ex1# set routing-options router-id 10.100.0.1

{master:0}[edit]
lab@ex1# show routing-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
router-id 10.100.0.1;

{master:0}[edit]
lab@ex1# commit
```

Verify again OSPF neighborship on EX1.

EX1:

```
lab@ex1# run show ospf3 neighbor
ID                  Interface             State     Pri   Dead
10.100.0.5          ge-0/0/3.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:740c
10.100.0.6          ge-0/0/4.0            Exchange  128      3
  Neighbor-address fe80::250:56ff:fea9:5c18

lab@ex1# run show ospf3 neighbor realm ipv4-unicast
ID                  Interface             State     Pri   Dead
10.100.0.5          ge-0/0/3.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:740c
10.100.0.6          ge-0/0/4.0            Exchange  128      3
  Neighbor-address fe80::250:56ff:fea9:5c18
```

It seems that correcting the router-id on EX1 fixed the issue with the neighborship between EX1 and vMX5. The neighborship between EX1 and vMX6 is still down.

```
lab@ex1# show interfaces ge-0/0/4
mtu 9216;
unit 0 {
    family inet {
        address 10.100.16.1/24;
    }
    family inet6 {
        address 2001:0db8:0100:016::1/64;
    }
}


lab@ex1# delete interfaces ge-0/0/4 mtu

{master:0}[edit]
lab@ex1# commit
configuration check succeeds
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Correcting the MTU on either EX1 or vMX6 so that they match would have fixed the adjacency issue between EX1 and vMX6. In this case we chose to correct the MTU on EX1 to make the MTU consistent on all links in the labs.

```
{master:0}[edit]
lab@ex1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
10.100.0.5      ge-0/0/3.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:740c
10.100.0.6      ge-0/0/4.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:5c18

lab@ex1# run show ospf3 neighbor realm ipv4-unicast
ID              Interface             State     Pri   Dead
10.100.0.5      ge-0/0/3.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:740c
10.100.0.6      ge-0/0/4.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:5c18
```

The output above shows that now the adjacencies with vMX5 and vMX6 are in the “Full” state.

4) Area 1 should contain Intra Area Prefixes LSAs for the IPv4 IRB addresses of EX1.

EX1:

```
lab@ex1# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.1 {
        interface ge-0/0/3.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
        interface ge-0/0/4.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
        interface lo0.0 {
            passive;
        }
        interface irb.100 {
            passive;
        }
        interface irb.101 {
            passive;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        interface irb.200 {
            passive;
        }
        interface irb.201 {
            passive;
        }
    }
}
```

Verification:

```
lab@ex1# run show ospf3 database realm ipv4-unicast intra-area-prefix detail
    OSPF3 database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
IntraArPfx *0.0.0.1          10.100.0.1       0x80000003   991  0x9129  88
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.100.0.1
  Prefix-count 7
  Prefix 10.100.15.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.16.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.100.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.101.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.200.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.201.0/24
    Prefix-options 0x0, Metric 1
  Prefix 10.100.0.1/32
    Prefix-options 0x2, Metric 0
...
```

5) Area 1 should contain External LSAs for the IPv6 IRB addresses of EX1.

EX1:

```
lab@ex1# show policy-options
policy-statement 2ospf-v6 {
    term 1 {
        from {
            rib inet6.0;
            interface [ irb.100 irb.101 irb.200 irb.201 ];
        }
        then accept;
    }
    term 2 {
        then reject;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@ex1# show protocols ospf3
export 2ospf-v6;
```

Verification:

```
lab@ex1# run show ospf3 database external detail
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern     *0.0.0.1          10.100.0.1       0x80000002   956  0x492d  36
  Prefix 2001:db8:100:100::/64
  Prefix-options 0x0, Metric 0, Type 2,
Extern     *0.0.0.2          10.100.0.1       0x80000002   600  0x5123  36
  Prefix 2001:db8:100:101::/64
  Prefix-options 0x0, Metric 0, Type 2,
Extern     *0.0.0.3          10.100.0.1       0x80000002   244  0x462d  36
  Prefix 2001:db8:100:200::/64
  Prefix-options 0x0, Metric 0, Type 2,
Extern     *0.0.0.4          10.100.0.1       0x80000001  1463  0x5022  36
  Prefix 2001:db8:100:201::/64
  Prefix-options 0x0, Metric 0, Type 2,
```

6) Area 0 should contain an InterArea Prefix LSA that encompasses the IPv4 IRB addresses of EX1. The individual IPv4 networks of EX1 IRB interface should not exist in Area 0.

vMX5 & vMX6:

```
lab@vmx5# show protocols ospf3 realm ipv4-unicast area 1
area-range 10.100.0.0/16;
interface ge-0/0/2.0 {
    interface-type p2p;
    hello-interval 1;
    dead-interval 4;
}
```

Verification:

```
lab@vmx5# run show ospf3 database realm ipv4-unicast area 0 detail inter-area-prefix

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
InterArPfx *0.0.0.8          10.100.0.5       0x80000001   705  0x63db  32
  Prefix 10.100.0.0/16
  Prefix-options 0x0, Metric 2
InterArPfx  0.0.0.8          10.100.0.6       0x80000001   337  0x5de0  32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix 10.100.0.0/16
  Prefix-options 0x0, Metric 2
```

7) No more than 1000 non self-generated LSAs are allowed for IPV4. Only generate a warning if this limit is exceeded.

vMX5, vMX6, and EX1:

```
lab@vmx5# show protocols ospf3 realm ipv4-unicast
database-protection {
    maximum-lsa 1000;
    warning-only;
}
```

8) The time to wait before running an SFP calculation should be 50 ms for IPV4 and 200 ms for IPV6.

vMX5, vMX6, and EX1:

```
lab@vmx5# show protocols ospf3 realm ipv4-unicast
spf-options {
    delay 50;
}
```

Verification:

vMX5:

```
lab@vmx5# run show ospf3 overview realm ipv4-unicast | match "SPF delay"
    SPF delay: 0.050000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

lab@vmx5# run show ospf3 overview  | match "SPF delay"
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

vMX6:

```
lab@vmx6# run show ospf3 overview realm ipv4-unicast | match "SPF delay"
    SPF delay: 0.050000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

lab@vmx6# run show ospf3 overview| match "SPF delay"
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

EX1:

```
lab@ex1# run show ospf3 overview| match "SPF delay"
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

lab@ex1# run show ospf3 overview realm ipv4-unicast | match "SPF delay"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
SPF delay: 0.050000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

9) Ensure that self-generated LSAs exchanged between the vMX devices in Site1 are only refreshed if the content of the LSA changes.

vMX5 & vMX6:

```
lab@vmx5# show protocols ospf3 area 0 interface ge-0/0/0
interface-type p2p;
hello-interval 1;
dead-interval 4;
flood-reduction;
```

```
lab@vmx5# show protocols ospf3 realm ipv4-unicast area 0 interface ge-0/0/0
interface-type p2p;
hello-interval 1;
dead-interval 4;
flood-reduction;
```

Verification:

vMX5:

```
lab@vmx5# run show ospf3 interface ge-0/0/0.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Address fe80::250:56ff:fea9:4cfd, Prefix-length 64
  OSPF3-Intf-index 1, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0, Flood Reduction
  Hello 1, Dead 4, ReXmit 5, Not Stub
  Protection type: None

lab@vmx5# run show ospf3 interface ge-0/0/0.0 detail realm ipv4-unicast
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Address fe80::250:56ff:fea9:4cfd, Prefix-length 64
  OSPF3-Intf-index 1, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0, Flood Reduction
  Hello 1, Dead 4, ReXmit 5, Not Stub
  Protection type: None
```

vMX6:

```
lab@vmx6# run show ospf3 interface ge-0/0/0.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Address fe80::250:56ff:fea9:4873, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1500, Cost 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Adj count: 1, Router LSA ID: 0, Flood Reduction
  Hello 1, Dead 4, ReXmit 5, Not Stub
  Protection type: None

lab@vmx6# run show ospf3 interface ge-0/0/0.0 detail realm ipv4-unicast
Interface           State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Address fe80::250:56ff:fea9:4873, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0, Flood Reduction
  Hello 1, Dead 4, ReXmit 5, Not Stub
  Protection type: None
```

Configurations for Task Completion in Site 1:

vMX5:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 dead-interval 4
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/2.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/2.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/2.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 dead-interval 4
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 area-range 10.100.0.0/16
set protocols ospf3 realm ipv4-unicast database-protection maximum-lsa 1000
set protocols ospf3 realm ipv4-unicast database-protection warning-only
set protocols ospf3 realm ipv4-unicast spf-options delay 50
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
```

vMX6:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 dead-interval 4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 dead-interval 4
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 area-range 10.100.0.0/16
set protocols ospf3 realm ipv4-unicast database-protection maximum-lsa 1000
set protocols ospf3 realm ipv4-unicast database-protection warning-only
set protocols ospf3 realm ipv4-unicast spf-options delay 50
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
```

EX1:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/3.0 dead-interval 4
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/4.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/4.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface ge-0/0/4.0 dead-interval 4
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/3.0 dead-interval 4
set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.0 dead-interval 4
set protocols ospf3 area 0.0.0.1 interface lo0.0 passive
set routing-options router-id 10.100.0.1
delete interfaces ge-0/0/4 mtu
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface irb.100 passive
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface irb.101 passive
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface irb.200 passive
set protocols ospf3 realm ipv4-unicast area 0.0.0.1 interface irb.201 passive
set policy-options policy-statement 2ospf-v6 term 1 from rib inet6.0 interface [ irb.100
irb.101 irb.200 irb.201 ]
set policy-options policy-statement 2ospf-v6 term 1 then accept
set policy-options policy-statement 2ospf-v6 term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 export 2ospf-v6
set protocols ospf3 realm ipv4-unicast database-protection maximum-lsa 1000
set protocols ospf3 realm ipv4-unicast database-protection warning-only
set protocols ospf3 realm ipv4-unicast spf-options delay 50
```

Site 2:

Device	Interface	area
vMX7	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0
vMX8	lo0	0.0.0.0
	ge-0/0/5	0.0.0.0

Table 33 OSPF information site 2

1) Configure OSPFv3 in Site 2 using the information as shown in the tables above.
2) Ensure no LSA type 0x2002 are found in the OSPFv3 databases in Site 1.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four seconds.

vMX7:

```
lab@vmx7# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-0/0/5.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
    interface ge-0/0/5.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}
```

vMX8:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx8# show protocols ospf3
realm ipv4-unicast {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-0/0/5.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
    interface ge-0/0/5.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}
```

Verification:

```
lab@vmx8# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
10.200.0.7      ge-0/0/5.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:60f5

lab@vmx8# run show ospf3 neighbor realm ipv4-unicast
ID              Interface             State     Pri   Dead
10.200.0.7      ge-0/0/5.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:60f5
```

Configurations for Task Completion in Site 2:

vMX7:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 hello-interval 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 dead-interval 4
```

vMX8:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 interface-type
p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 hello-interval
1
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface ge-0/0/5.0 dead-interval 4
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/5.0 dead-interval 4
```

Task 2: OSPFv2

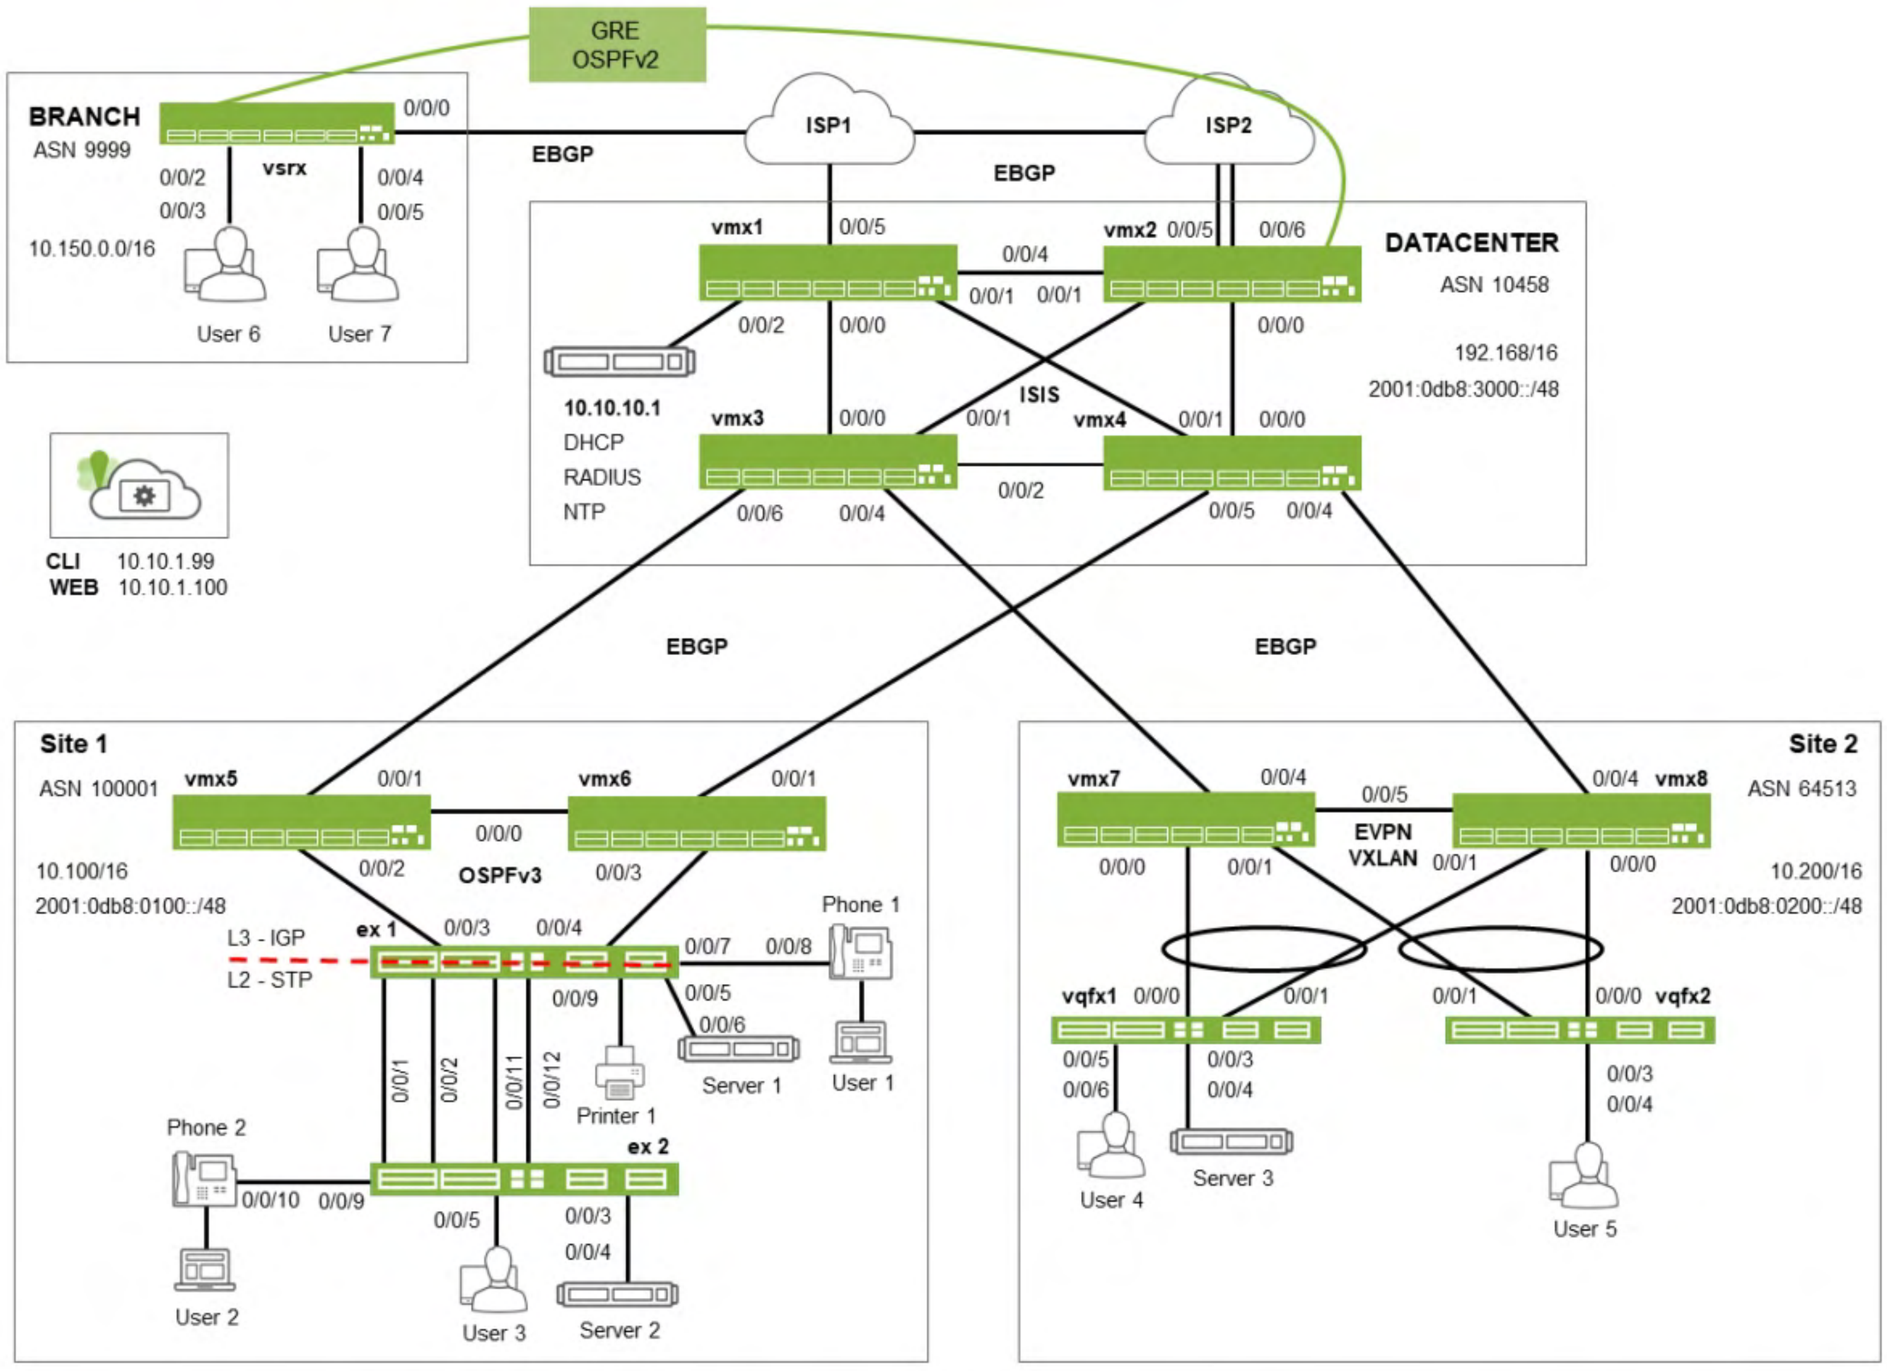

In this task you will configure an OSPFv2 over the GRE tunnel between the data center and the Branch office.

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Interface	area
vMX2	GRE interface	0.0.0.0
vSRX	GRE interface	0.0.0.0

Table 34 OSPF information for GRE tunnel

1) Enable OSPFv2 over the GRE interface as shown in the table above.
2) Ensure the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.
3) vSRX should receive a default route and the data center IPv4 aggregate route from vMX2. Both routes should exist in the OSPFv2 database.
4) Configure OSPF in such a way that vSRX only accepts the default route from vMX2.

Note: For this task, you will not be able to verify your OSPF neighborships are established until your GRE tunnel is operational and this requires BGP routing from vMX2 through the internet. You will configure this functionality in Chapter 5.

vMX2:

```
lab@vmx2# show policy-options policy-statement 2ibgpandospfv2
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol ospf;
        route-filter 10.150.99.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}

lab@vmx2# show protocols ospf
export 2ibgpandospfv2;
area 0.0.0.0 {
    interface gr-0/0/10.0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        hello-interval 1;
    }
}
```

vSRX:

```
lab@vsrx# show policy-options
policy-statement 2ospf {
    term 1 {
        from interface irb.99;
        then {
            external {
                type 1;
            }
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
policy-statement ospfimport {
    term 1 {
        from {
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}

lab@vsrx# show protocols ospf
export 2ospf;
import ospfimport;
area 0.0.0.0 {
    interface gr-0/0/0.0 {
        hello-interval 1;
    }
}
```

Verification:

Note: You cannot verify this task at this point. You must have routing enabled through your ISP networks and your GRE tunnel must be established. You can verify the proper behavior after completing Chapter 5. The following outputs are after completing Chapter 5 tasks.

© 2020 Juniper Networks, Inc. All rights reserved.

vSRX:

```
lab@vsrx# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
173.17.99.2      gr-0/0/0.0             Full      192.168.0.2      128     3

lab@vsrx# run show route protocol ospf

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:51:02, metric 0, tag 0
                    >  via gr-0/0/0.0
173.17.99.0/24      [OSPF/10] 1d 18:11:26, metric 1
                    >  via gr-0/0/0.0
224.0.0.5/32       *[OSPF/10] 1d 18:11:31, metric 1
                       MultiRecv
...
```

Configurations for Task Completion:

vMX2:

```
set protocols ospf export 2ibgpandospfv2
set protocols ospf area 0.0.0.0 interface gr-0/0/10.0 hello-interval 1
set policy-options policy-statement 2ibgpandospfv2 term 1 from protocol aggregate
set policy-options policy-statement 2ibgpandospfv2 term 1 from route-filter
192.168.0.0/16 exact
set policy-options policy-statement 2ibgpandospfv2 term 1 then accept
set policy-options policy-statement 2ibgpandospfv2 term 2 from protocol ospf
set policy-options policy-statement 2ibgpandospfv2 term 2 from route-filter
10.150.99.0/24 exact
set policy-options policy-statement 2ibgpandospfv2 term 2 then accept
set policy-options policy-statement 2ibgpandospfv2 term 3 from protocol aggregate
set policy-options policy-statement 2ibgpandospfv2 term 3 from route-filter 0.0.0.0/0
exact
set policy-options policy-statement 2ibgpandospfv2 term 3 then accept
```

vSRX:

```
set protocols ospf export 2ospf
set protocols ospf import ospfimport
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 hello-interval 1
set policy-options policy-statement 2ospf term 1 from interface irb.99
set policy-options policy-statement 2ospf term 1 then external type 1
set policy-options policy-statement 2ospf term 1 then accept
set policy-options policy-statement 2ospf term 2 then reject
set policy-options policy-statement ospfimport term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement ospfimport term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement ospfimport term 2 then reject
```

Task 3: IS-IS

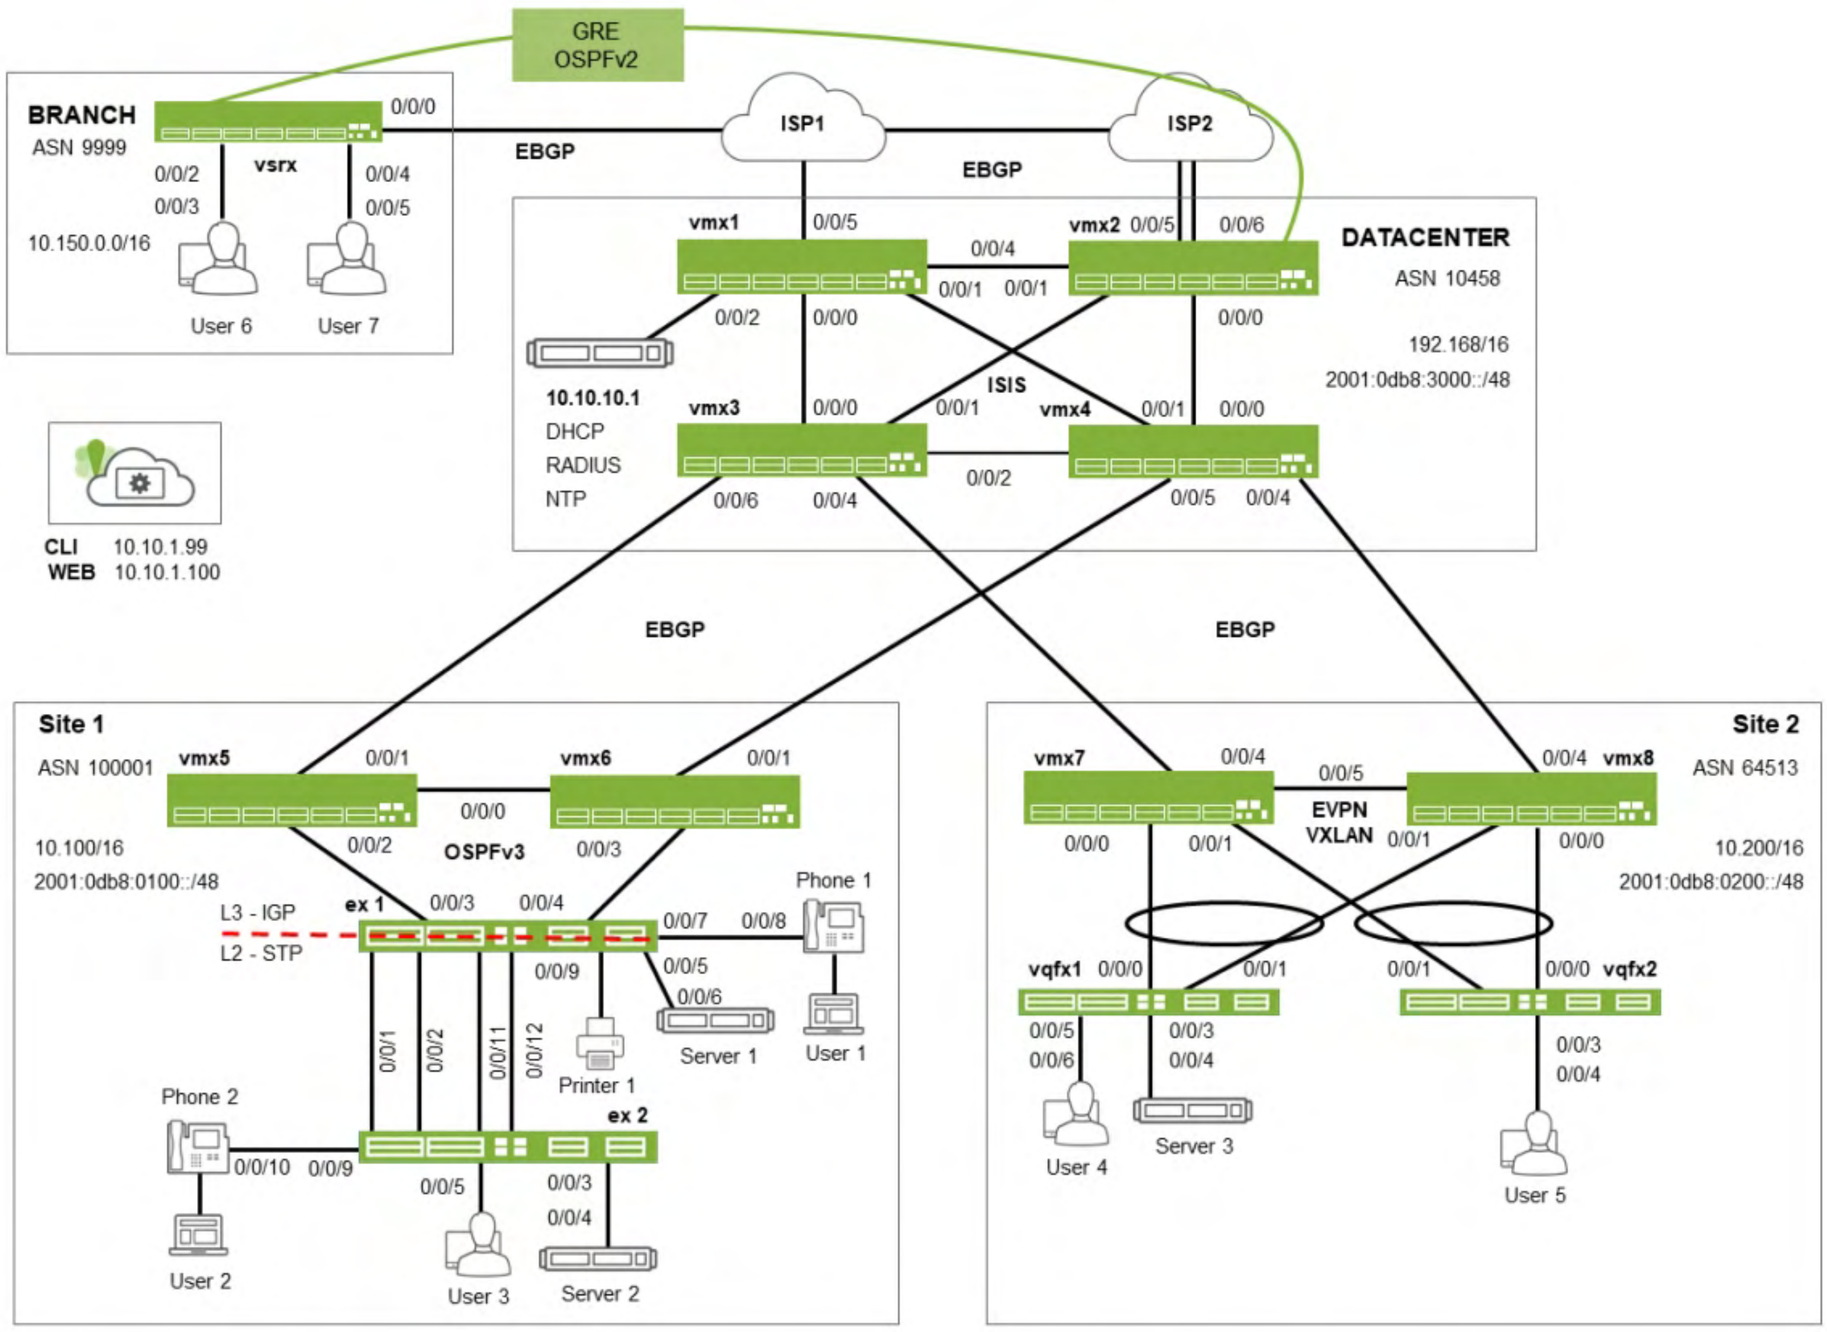

In this task you will configure IS-IS as the IGP in the Data Center.

1) Configure IS-IS in the data center using the information as shown in the tables below.

Device	Level	Interfaces
vMX1	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
vMX2	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4

© 2020 Juniper Networks, Inc. All rights reserved.

vMX3	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2
vMX4	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2

Table 35 IS-IS information

Device	Required IS-IS Configuration
vMX1	Sysid: 0168.0000.0001 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, Ipv6 is enabled
vMX2	Sysid: 0168.0000.0002 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled
vMX3	Sysid: 0168.0000.0003 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled

© 2020 Juniper Networks, Inc. All rights reserved.

vMX4	Sysid: 0168.0000.0004 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is enabled

2) Do not enable Level 1 routing.
3) Ensure that no DIS is elected.
4) Ensure interface ge-0/0/2 on vMX1 participates in IS-IS, but no adjacency can be formed.
5) The maximum number of routes that can be redistributed in IS-IS should be 100.
6) Enable MD5 authentication for level 2. Use “juniper123” as the key.
7) Configure the hello and hold interval to 1 and 4 for the IS-IS level 2 interface between vMX1 and vMX3.

vMX1:

```
lab@vmx1# show protocols isis
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"; ## SECRET-DATA
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
    level 2 {
        hello-interval 1;
        hold-time 4;
    }
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    passive;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
interface ge-0/0/4.0 {
    point-to-point;
}
interface lo0.0;

lab@vmx1# show interfaces lo0.0 family iso
address 49.0001.0192.0168.0000.0001.00;

lab@vmx1# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.13.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::1/64;
    }
}

lab@vmx1# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::1/64;
    }
}

lab@vmx1# show interfaces ge-0/0/4
unit 0 {
    family inet {
        address 192.168.12.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::1/64;
    }
}
```

vMX2:

```
lab@vmx2# show protocols isis
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
level 2 {
    authentication-key "$9$jMkmT69pRhrz3hrev7Nik.Pz3/CtOIE"; ## SECRET-DATA
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    passive;
}
interface ge-0/0/4.0 {
    point-to-point;
}
interface lo0.0;

lab@vmx2# show interfaces lo0.0 family iso
address 49.0001.0192.0168.0000.0002.00;


lab@vmx2# show interfaces ge-0/0/0
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.24.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0024::2/64;
    }
}

[edit]
lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.23.2/24;
    }
    family iso;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
}

[edit]
lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
}
```

vMX3:

```
lab@vmx3# show protocols isis
lsp-lifetime 600;
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$Yy4JD.PQ9A0ikA0BIrl24aZikmfT3/C"; ## SECRET-DATA
    authentication-type md5;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
    level 2 {
        hello-interval 1;
        hold-time 4;
    }
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    point-to-point;
}
interface lo0.0;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx3# show interfaces lo0.0 family iso
address 49.0001.0192.0168.0000.0003.00;


lab@vmx3# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.13.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::3/64;
    }
}


lab@vmx3# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.23.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0023::3/64;
    }
}


lab@vmx3# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::3/64;
    }
}
```

vMX4:

```
lab@vmx4# show protocols isis
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$NKVs4UjqQF/aZF/CtIR-VwYaZDikPTz"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    authentication-type md5;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    point-to-point;
}
interface lo0.0;
```

```
lab@vmx4# show interfaces lo0.0 family iso
address 49.0001.0192.0168.0000.0004.00;

lab@vmx4# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.24.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0024::4/64;
    }
}

lab@vmx4# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::4/64;
    }
}


lab@vmx4# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::4/64;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

Verification:

vMX1:

```
lab@vmx1# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0001
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled, IPv6 is enabled

[edit]
lab@vmx1# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx3           2  Up                  3
ge-0/0/1.0            vmx4           2  Up                 22
ge-0/0/4.0            vmx2           2  Up                 26

[edit]
lab@vmx1# run show route protocol isis

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

192.168.0.2/32     *[IS-IS/18] 00:23:53, metric 10
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.0.3/32     *[IS-IS/18] 00:22:40, metric 10
                    > to 192.168.13.3 via ge-0/0/0.0
192.168.0.4/32     *[IS-IS/18] 00:20:03, metric 10
                    > to 192.168.14.4 via ge-0/0/1.0
192.168.23.0/24    *[IS-IS/18] 00:22:40, metric 20
                      to 192.168.13.3 via ge-0/0/0.0
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.24.0/24    *[IS-IS/18] 00:20:03, metric 20
                    > to 192.168.14.4 via ge-0/0/1.0
                      to 192.168.12.2 via ge-0/0/4.0
192.168.34.0/24    *[IS-IS/18] 00:20:03, metric 20
                      to 192.168.13.3 via ge-0/0/0.0
                    > to 192.168.14.4 via ge-0/0/1.0

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

inet6.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
2001:db8:3000::2/128
                   *[IS-IS/18] 00:23:53, metric 10
                    > to fe80::250:56ff:fea9:3b9f via ge-0/0/4.0
2001:db8:3000::3/128
                   *[IS-IS/18] 00:22:40, metric 10
                    > to fe80::250:56ff:fea9:3aaa via ge-0/0/0.0
2001:db8:3000::4/128
                   *[IS-IS/18] 00:20:03, metric 10
                    > to fe80::250:56ff:fea9:12b5 via ge-0/0/1.0
2001:db8:3000:23::/64
                   *[IS-IS/18] 00:22:40, metric 20
                    > to fe80::250:56ff:fea9:3aaa via ge-0/0/0.0
                      to fe80::250:56ff:fea9:3b9f via ge-0/0/4.0
2001:db8:3000:24::/64
                   *[IS-IS/18] 00:20:03, metric 20
                    > to fe80::250:56ff:fea9:12b5 via ge-0/0/1.0
                      to fe80::250:56ff:fea9:3b9f via ge-0/0/4.0
2001:db8:3000:34::/64
                   *[IS-IS/18] 00:20:03, metric 20
                    > to fe80::250:56ff:fea9:3aaa via ge-0/0/0.0
                      to fe80::250:56ff:fea9:12b5 via ge-0/0/1.0
```

vMX2:

```
lab@vmx2# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0002
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled, IPv6 is enabled

[edit]
lab@vmx2# run show isis adjacency
run show route protocol isis
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx4           2  Up                   23
ge-0/0/1.0            vmx3           2  Up                   26
ge-0/0/4.0            vmx1           2  Up                   25

[edit]
lab@vmx2# run show route protocol isis

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


© 2020 Juniper Networks, Inc. All rights reserved.

```
10.10.10.0/24      *[IS-IS/18] 00:25:00, metric 20
                    > to 192.168.12.1 via ge-0/0/4.0
192.168.0.1/32     *[IS-IS/18] 00:25:00, metric 10
                    > to 192.168.12.1 via ge-0/0/4.0
192.168.0.3/32     *[IS-IS/18] 00:23:18, metric 10
                    > to 192.168.23.3 via ge-0/0/1.0
192.168.0.4/32     *[IS-IS/18] 00:21:35, metric 10
                    > to 192.168.24.4 via ge-0/0/0.0
192.168.13.0/24    *[IS-IS/18] 00:23:18, metric 20
                    > to 192.168.23.3 via ge-0/0/1.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.14.0/24    *[IS-IS/18] 00:21:35, metric 20
                    > to 192.168.24.4 via ge-0/0/0.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.34.0/24    *[IS-IS/18] 00:21:35, metric 20
                      to 192.168.24.4 via ge-0/0/0.0
                    > to 192.168.23.3 via ge-0/0/1.0

fbf.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)

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

inet6.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:3000::1/128
                   *[IS-IS/18] 00:25:00, metric 10
                    > to fe80::250:56ff:fea9:4466 via ge-0/0/4.0
2001:db8:3000::3/128
                   *[IS-IS/18] 00:23:18, metric 10
                    > to fe80::250:56ff:fea9:32cf via ge-0/0/1.0
2001:db8:3000::4/128
                   *[IS-IS/18] 00:21:35, metric 10
                    > to fe80::250:56ff:fea9:51c6 via ge-0/0/0.0
2001:db8:3000:13::/64
                   *[IS-IS/18] 00:23:18, metric 20
                    > to fe80::250:56ff:fea9:32cf via ge-0/0/1.0
                      to fe80::250:56ff:fea9:4466 via ge-0/0/4.0
2001:db8:3000:14::/64
                   *[IS-IS/18] 00:21:35, metric 20
                    > to fe80::250:56ff:fea9:51c6 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:4466 via ge-0/0/4.0
2001:db8:3000:34::/64
                   *[IS-IS/18] 00:21:35, metric 20
                    > to fe80::250:56ff:fea9:51c6 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:32cf via ge-0/0/1.0
```

vMX3:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx3# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0003
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 600
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled, IPv6 is enabled

[edit]
lab@vmx3# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx1           2  Up                  3
ge-0/0/1.0            vmx2           2  Up                 24
ge-0/0/2.0            vmx4           2  Up                 26

[edit]
lab@vmx3# run show route protocol isis

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

10.10.10.0/24      *[IS-IS/18] 00:24:23, metric 20
                    > to 192.168.13.1 via ge-0/0/0.0
192.168.0.1/32     *[IS-IS/18] 00:24:23, metric 10
                    > to 192.168.13.1 via ge-0/0/0.0
192.168.0.2/32     *[IS-IS/18] 00:23:53, metric 10
                    > to 192.168.23.2 via ge-0/0/1.0
192.168.0.4/32     *[IS-IS/18] 00:21:45, metric 10
                    > to 192.168.34.4 via ge-0/0/2.0
192.168.12.0/24    *[IS-IS/18] 00:23:53, metric 20
                    > to 192.168.13.1 via ge-0/0/0.0
                      to 192.168.23.2 via ge-0/0/1.0
192.168.14.0/24    *[IS-IS/18] 00:21:45, metric 20
                    > to 192.168.13.1 via ge-0/0/0.0
                      to 192.168.34.4 via ge-0/0/2.0
192.168.24.0/24    *[IS-IS/18] 00:21:45, metric 20
                    > to 192.168.23.2 via ge-0/0/1.0
                      to 192.168.34.4 via ge-0/0/2.0

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

inet6.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:3000::1/128
                   *[IS-IS/18] 00:24:23, metric 10
                    > to fe80::250:56ff:fea9:3ed3 via ge-0/0/0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
2001:db8:3000::2/128
                   *[IS-IS/18] 00:23:53, metric 10
                    > to fe80::250:56ff:fea9:35aa via ge-0/0/1.0
2001:db8:3000::4/128
                   *[IS-IS/18] 00:21:45, metric 10
                    > to fe80::250:56ff:fea9:36 via ge-0/0/2.0
2001:db8:3000:12::/64
                   *[IS-IS/18] 00:23:53, metric 20
                    > to fe80::250:56ff:fea9:3ed3 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:35aa via ge-0/0/1.0
2001:db8:3000:14::/64
                   *[IS-IS/18] 00:21:45, metric 20
                    > to fe80::250:56ff:fea9:3ed3 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:36 via ge-0/0/2.0
2001:db8:3000:24::/64
                   *[IS-IS/18] 00:21:45, metric 20
                    > to fe80::250:56ff:fea9:35aa via ge-0/0/1.0
                      to fe80::250:56ff:fea9:36 via ge-0/0/2.0
```

vMX4:

```
lab@vmx4# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0004
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled, IPv6 is enabled

[edit]
lab@vmx4# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx2           2  Up                    21
ge-0/0/1.0            vmx1           2  Up                    23
ge-0/0/2.0            vmx3           2  Up                    23

[edit]
lab@vmx4# run show route protocol isis

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

10.10.10.0/24      *[IS-IS/18] 00:22:54, metric 20
                    > to 192.168.14.1 via ge-0/0/1.0
192.168.0.1/32     *[IS-IS/18] 00:22:54, metric 10
                    > to 192.168.14.1 via ge-0/0/1.0
192.168.0.2/32     *[IS-IS/18] 00:23:18, metric 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 192.168.24.2 via ge-0/0/0.0
192.168.0.3/32     *[IS-IS/18] 00:22:54, metric 10
                    > to 192.168.34.3 via ge-0/0/2.0
192.168.12.0/24    *[IS-IS/18] 00:22:54, metric 20
                    > to 192.168.24.2 via ge-0/0/0.0
                      to 192.168.14.1 via ge-0/0/1.0
192.168.13.0/24    *[IS-IS/18] 00:22:54, metric 20
                    > to 192.168.14.1 via ge-0/0/1.0
                      to 192.168.34.3 via ge-0/0/2.0
192.168.23.0/24    *[IS-IS/18] 00:22:54, metric 20
                      to 192.168.24.2 via ge-0/0/0.0
                    > to 192.168.34.3 via ge-0/0/2.0

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

inet6.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:3000::1/128
                   *[IS-IS/18] 00:22:54, metric 10
                    > to fe80::250:56ff:fea9:427f via ge-0/0/1.0
2001:db8:3000::2/128
                   *[IS-IS/18] 00:23:18, metric 10
                    > to fe80::250:56ff:fea9:3944 via ge-0/0/0.0
2001:db8:3000::3/128
                   *[IS-IS/18] 00:22:54, metric 10
                    > to fe80::250:56ff:fea9:6e2a via ge-0/0/2.0
2001:db8:3000:12::/64
                   *[IS-IS/18] 00:22:54, metric 20
                    > to fe80::250:56ff:fea9:3944 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:427f via ge-0/0/1.0
2001:db8:3000:13::/64
                   *[IS-IS/18] 00:22:54, metric 20
                    > to fe80::250:56ff:fea9:427f via ge-0/0/1.0
                      to fe80::250:56ff:fea9:6e2a via ge-0/0/2.0
2001:db8:3000:23::/64
                   *[IS-IS/18] 00:22:54, metric 20
                    > to fe80::250:56ff:fea9:3944 via ge-0/0/0.0
                      to fe80::250:56ff:fea9:6e2a via ge-0/0/2.0
```

Configurations for Task Completion:

vMX1:

```
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"
set protocols isis level 2 authentication-type md5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/0.0 level 2 hello-interval 1
set protocols isis interface ge-0/0/0.0 level 2 hold-time 4
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 passive
set protocols isis interface ge-0/0/4.0 point-to-point
set protocols isis interface lo0.0
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0001.00
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/4 unit 0 family iso
```

vMX2:

```
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$jMkmT69pRhrz3hrev7Nik.Pz3/CtOIE"
set protocols isis level 2 authentication-type md5
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 passive
set protocols isis interface ge-0/0/4.0 point-to-point
set protocols isis interface lo0.0
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0002.00
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/4 unit 0 family iso
```

vMX3:

```
set protocols isis lsp-lifetime 600
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$Yy4JD.PQ9A0ikA0BIrl24aZikmfT3/C"
set protocols isis level 2 authentication-type md5
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/0.0 level 2 hello-interval 1
set protocols isis interface ge-0/0/0.0 level 2 hold-time 4
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 point-to-point
set protocols isis interface lo0.0
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0003.00
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family iso
```

vMX4:

```
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$NKVs4UjqQF/aZF/CtIR-VwYaZDikPTz"
set protocols isis level 2 authentication-type md5
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 point-to-point
set protocols isis interface lo0.0
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0004.00
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family iso
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Chapter 5: BGP

This chapter focuses on the implementation of BGP for routing within the data center and within Site 2. BGP is also used for routing between the data center, ISP1, ISP2, and between the data center, Site 1, and Site 2. At various points in the topology you need to perform route redistribution between and IGP and BGP for end-to-end connectivity.

BGP Overview

BGP is an Exterior Gateway Protocol used to exchange routing information between Autonomous Systems (AS). An AS is a set of devices under a common administration and each AS represents a different domain on the Internet.

BGP is a path-vector protocol and uses the AS path to evaluate or gauge the distance to a destination prefix. The AS path is also the method BGP uses to prevent routing-loops. If a BGP speaking router encounters its own path in the AS path, the route is not eligible for use.

BGP uses TCP (port 179) to provide for reliable communication. BGP neighbors are explicitly configured and as soon as a neighbor relationship is set up, the BGP peers start exchanging Network Layer Reachability Information (NLRI).

There are two different types of BGP neighbor relationships. The peering session between routers that belong to a different AS is called an External BGP (EBGP) peering session. The peering session between routers belonging to the same AS peer is called an Internal BGP (IBGP) peering session. The rules for prefix advertisements differ for EBGP and IBGP peers:

- IBGP peers advertise routes received from EBGP peers to other IBGP peers.
- EBGP peers advertise routes learned from IBGP or EBGP peers to other EBGP peers.
- IBGP peers do not advertise routes learned from IBGP peers to other IBGP peers.

That last rule is needed to prevent loops inside an AS. This is because the AS path is only updated across EBGP peering sessions. Because IBGP peers do not re-advertise IBGP learned routes, the IBGP neighbors should be fully meshed. Scaling methods to circumvent this requirement include route-reflection and confederation (neither are a topic on the exam).

The AS path is not the only attribute used by BGP. The attributes are grouped into four main classes:

© 2020 Juniper Networks, Inc. All rights reserved.

- **Well-known mandatory**: must be supported by all BGP neighbors and must be included in all BGP updates.
- **Well-known discretionary**: must be supported by all BGP neighbors, do not have to be included in all updates.
- **Optional transitive**: support is not mandatory. When the attribute is encountered and it is not recognized, it must be passed to other neighbors.
- **Optional non-transitive**: support is not mandatory. When the attribute is encountered and it is not recognized, it can be ignored.

The following attributes are some of the most used attributes in BGP Enterprise environments:

- **Next-Hop**: this a well-known mandatory attribute. It represents the IP address of the advertising router. For a router to install a BGP route into the routing-table, this IP address must be reachable. EBGP peers update this attribute automatically to the IP address of their outgoing interface. IBGP peers do not update this attribute automatically, but they can be configured to do so.
- **AS Path**: well-known mandatory attribute that represents the path to the source of the prefix. The AS-path is a list of autonomous-systems. BGP also uses the AS-path to prevent loops. The AS-path is updated automatically across EBGP sessions. The router that sends the routing-update prepends its own AS to the list before updating its neighbor.
- **Origin**: well-known mandatory attribute. The origin describes the way in which the route was originally injected into BGP. The values can be IGP (I), EGP (E) and Incomplete (?). By default, JUNOS will set this value to IGP.
- **Local Preference**: Well-known discretionary attribute that is used to define the most preferable way out of the local AS. The route with the highest Local Preference wins. The local-preference attribute is only sent across IBGP sessions. the default value is 100.
- **Multi-Exit Discriminator**: optional non-transitive attribute, used to influence the preferred path into the AS when there are multiple links between the two ASs. After the MED is send across the EBGP session into the neighboring AS, that AS will not advertise it further across its EBGP sessions. When the MED is absent, the Junos OS treats it as though the value was set to 0.
- **Community**: optional transitive attribute that is used to tag a prefix or a group of prefixes. After tagging them, they can be easily identified later on and they can help simplify policies applied to routes. Even between different ASs.

BGP configuration and implementation

IBGP sessions are usually sourced from the loopback interface, which is specified using the **local-address** configuration command. In case of IBGP, we can either configure the

© 2020 Juniper Networks, Inc. All rights reserved.

group/neighbor to be **internal**, or we can configure the group/neighbor with the **peer-as** statement. In addition to this, we also need to specify the AS of the system itself:

set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.0.0.1
set protocols bgp group ibgp neighbor 10.0.0.2
set protocols bgp group ibgp neighbor 10.0.0.3
set protocols bgp group ibgp neighbor 10.0.0.4

set routing-options autonomous-system 65000

After setting the AS in the IBGP configuration, adding a standard EBGP peer would require the following configuration:

set protocols bgp group ebgp neighbor 192.168.1.2 peer-as 65001

The EBGP peer will set the next-hop attribute to the IP address that it used to source the EBGP session. To be able to use the routes that are advertised by this peer, the rest of the network needs to be able to resolve this next-hop IP address. This can be achieved in one of several ways:

- Use a next-hop self policy towards the IBGP neighbors
- Export the subnet used for the EBGP peering into the IGP through the use of a policy
- Configure the interface used for EBGP as a passive interface under the IGP
- By configuring a static route to the EBGP next-hop on each IBGP router
- Configuring an IGP adjacency with the EBGP peer

The solution that is most-often used (and recommended) is to apply a next-hop self policy on the IBGP sessions. IBGP sessions do not change the next-hop IP address by default. Through the use of a policy, you can change this behavior. When the IBGP session are sourced from the loopback interface, this IP address will (probably) be redistributed into the IGP already. This way, the other IBGP neighbors will not have any problems resolving the next-hop. To implement this solution, you can use the following policy and export it to the IBGP:

set policy-options policy-statement *nhs* term 1 then next-hop self

set protocols bgp group ibgp export *nhs*

Load sharing across multiple BGP sessions can be achieved using the **multipath** configuration command. In case both EBGP speakers are using the loopback IP address to source the EBGP session, don't forget the **multihop** feature to overcome BGP's default

© 2020 Juniper Networks, Inc. All rights reserved.

TTL of 1. IP connectivity towards the loopback of the other BGP speaker can be provided through the configuration of a static route or by creating an IGP adjacency between the routers and injecting the BGP speakers' loopback IP address. The **local-address** command should be used to source the local EBGP session using the loopback IP.

Both methods of load sharing, multiple sessions or a single multihop session, require the configuration of a policy that is applied to the forwarding table:

set policy-options policy-statement *load-balancing* then load-balance per-packet

set routing-options forwarding-table export *load-balancing*

In some situations, especially when we use a link to share IPv6 routes over an IPv4 BGP session, we need to accept a remote next hop that does not belong to the same subnet. In such cases, we can use the **accept-remote-nexthop** configuration knob.

Note that we cannot configure **multihop** and **accept-remote-nexthop** at the same time.

Some common configuration items related to BGP you should know of include the following:

set protocols bgp log-updown

Configure the system to generate a logging message for every BGP state transition. This is also configurable per-group.

set protocols bgp group xx authentication-key workbook

Configures all peering session within the group with an MD5 authentication key.

set protocols bgp group xx hold-time 100

Specifies the advertised hold-time value. The default value is 90. The hold-time is three times the interval at which keepalives will be send. The lowest hold-time of the two peers will be used.

set protocols bgp group xx family inet unicast

Configures the BGP peers for the inet unicast address-family. This address family is active by default but has to be specified as soon as another address family is configured.

set protocols bgp group xx family inet6 unicast

Configures the BGP peers for the inet6 unicast address family.

© 2020 Juniper Networks, Inc. All rights reserved.

BGP Route Selection

BGP determines the best path based on the following comprehensive route-selection procedure:

1) lowest **route preference**
2) highest **local-preference**
3) shortest **AS-PATH**
4) lowest **origin value** (IGP over EGP over incomplete)
5) lowest **multiple exit discriminator** (no MED is treated as 0)
    a) checked only for routes from same AS, changed with 'always-compare-med'
6) **EBGP** over **IBGP**. When there are only EBGP routes, skip the next step.
7) lowest **IGP** cost towards BGP next-hop
8) Differs for internal and external routes:
    a) For internal routes: prefer the route with the lowest router-id
    b) For external routes: select the oldest (or currently) active route
9) (when route-reflection is used) shortest cluster list
10) lowest peer IP address

4 Byte AS Number

AS numbers are assigned by the Regional Internet Registers. Initially, they were 16-bits (2-bytes). In anticipation of the 2-byte AS numbers' depletion, 32-bits (or 4-bytes) AS numbers were introduced. These 4-byte AS number are made backwards compatible with the old 2-byte AS numbers. The way in which BGP provides for this backwards compatibility is described in RFC 4893.

In that RFC, two new optional transitive attributes are defined, the AS4_PATH and the AS4_AGGREGATOR. Both these attributes were developed to make sure that 4-byte AS information (AS path and aggregates) could be propagated across networks containing 'old' BGP speakers that do not offer support for a 4-byte AS. The RFC also stated that a 2-octet AS number (AS 23456) is to be reserved.

The way in which this all works is when an old and a new BGP speaker peer with each other, the new BGP speaker notices that the old BGP speaker did not send a 4-byte AS number capability along with its open message. The new BGP speaker will present itself as AS 23456 to the old BGP speaker and it will encode the AS-path information into the new communities.

© 2020 Juniper Networks, Inc. All rights reserved.

When the old BGP speaker advertises the path information to another new BGP speaker, that BGP speaker can reconstruct the AS path information from the added attributes. This way, the information was effectively tunneled across the old BGP speaker.

Notice that there are three formats to display 4-byte AS numbers:

- **asplain**: decimal representation of the AS, ranging from 0 – 4294967295.
- **asdot+**: the 32-bit number is broken up into a lower-order half and a high-order half. If you start incrementing the AS numbers at a low value, you can count from '0.1' up until '0.65535'. If you were to increment that AS number by 1, you would change the high-order bit to 1 and the low order bits to 0. You'd end up with '1.0' From then on, you can increment the low-order value until you reach 65535 again ('1.65535').
- **asdot**: mixture of asplain and asdot+. Values in the range of asplain are written in asplain. The values that exceed this range are written in asdot+.

Note: Junos OS allows for the configuration of BGP peers in asplain as well as asdot. Junos OS can also be configured to have the output of the BGP show commands display the AS in asdot.

Route Policy

Before applying any routing-policies, it is important to know that the BGP protocol uses three different tables: RIB-IN, RIB-LOCAL and RIB-OUT.

The RIB-IN table is created for every established peer. The routes received from peers are stored in this table. The import policies are applied after the routes are stored.

The RIB-LOCAL is the route-table of the device. This information is used to populate the forwarding table.

The RIB-OUT is created for every peer and stores the routes that are advertised to the peer. Routes in this table have already passed the export policy.

The BGP protocol is policy driven protocol. We can use a routing-policy to determine how we want traffic to flow. The following BGP attributes can be manipulated using the **then** action in a routing policy:

- **next hop - then next-hop**
  - we can alter the next-hop across IBGP sessions to make sure other routers can resolve the advertised routes
  - the next-hop can be set to a discard interface, in case we want to black hole traffic

© 2020 Juniper Networks, Inc. All rights reserved.

  - in case an IPv4 BGP session is used to advertise information for both IPv4 and IPv6, we can alter the next-hop for IPv6 routing information.

- **local preference - then local-preference**
  - the highest local preference is preferred. Used to determine the exit point for outbound traffic inside the AS.
- **origin - then origin**
  - lowest origin value will win (IGP (I) over EGP (E) over incomplete(?))
- **MED - then metric**
  - used to manipulate inbound traffic into the AS in case multiple EBPG sessions exist with another AS.
- **AS path - then as-path-prepend**
  - In an import policy, it can be used to manipulate outbound traffic.
  - In an export policy, it can be used to manipulate inbound traffic.
- **Community - then community**
  - add: adds a community to an empty or existing list of communities.
  - set: overwrite the list of existing communities (removing previously applied communities, if any).
  - delete: deletes the reference community.
  - well-known communities can be configured by specifying the name:
    - **no-export**: used to prevent route-advertisements to other AS.
    - **no-advertise:** used to prevent route-advertisement to other BGP peers.

The following attributes can be manipulated through BGP peering configuration without the use of a policy:

- **as-override**
  - used when an AS number is used on different sites that are connected to the same network.
  - replace the peer AS with your own AS before advertising the route to a peer.
- **remove-private**
  - strips all private AS numbers from the AS path. The router stops stripping private AS numbers from left to right until it encounters a public AS.
- **metric-out**
  - explained earlier.
- **local-preference**
  - explained earlier.

© 2020 Juniper Networks, Inc. All rights reserved.

AS-path and community regular expressions

You can use AS-path regular expressions to locate or identify a set of routes in a policy or when issuing a 'show route aspath-regex' command. AS-path regular expressions are configured under the [**policy-options as-path**] hierarchy and can be referenced as a **from** condition inside a routing-policy.

Commonly used AS-path Regex:

.	Any one AS (not a single character).
*	0 or more repetitions of the AS.
{m,n}	Match at least m repetitions and at most n repetitions of the term.
{m}	Matches exactly m repetitions of the term.
{m,}	Matches m or more repetitions of the term.
+	Matches one or more instances of the term.
?	Zero or 1 repetitions of the term.
|	Matches the term on either side of this sign, performing an OR.
-	Between a starting and an ending range.
()	matches a null-path.

© 2020 Juniper Networks, Inc. All rights reserved.

Note:

^	Added implicitly at the beginning of the regex, matches the start of the string.
$	Added implicitly at the end of the regex, matches the end of the string.

Some examples:

1234*	matches 1234, 1234 1234, null AS.
.*456.*	matches any length path that contains the list 4,5,6.
64{2,6} 65?	matches 2 to 6 instances of AS 64 followed by zero or 1 instance of AS 65.
.* 62 63 64 .*	matches a path that contains the list '62 63 64' anywhere.
"()"	matches null AS (can be used to advertise only the routes that originated from your own AS).

You can configure and apply communities to tag, or group together, certain prefixes. Additionally, you can use the configured communities to gather output for the 'show route community-name' command. When you are configuring communities, you can choose to use regex.

Commonly used community Regex:

.	=	Any single digit
*****	=	Any single community
.*	=	Everything in complex form

© 2020 Juniper Networks, Inc. All rights reserved.

Examples:

"* : *"	matches all communities (in case you want to strip routes from any communities)
600:*	match any community which first part is 600
111:50.0	matches 111:5000, 111:5010, 111:5020, …
"^56: (2.*)$"	matches 56:2111, 56:26765, ….

Please refer to Juniper web site for more information about regular expressions.

Monitoring and troubleshooting

Common BGP issues related to the BGP sessions include issues with:

- Connectivity (no IP connectivity between the hosts)
- BGP misconfiguration (wrong AS, wrong source address, TTL of 1)
- Authentication mismatch (wrong password, no password)
- Firewall filters (BGP peer address no allowed in the filter used to protect the control plane)

In some cases, your BGP session will be established, but you will find that there is still no communication between the autonomous systems. In these cases, the issue could be related to the following:

- Policy misconfiguration (too much or the wrong prefixes are filtered)
- BGP next hop reachability (the next-hop IP cannot be resolved)

Troubleshooting BGP usually starts with the **show bgp summary** command. Here we can quickly spot which peer is not in the established state. We should then check for connectivity issues. If there is IP connectivity between the IP addresses used to setup the peering session, we can see if **show log message** outputs anything related to the issue.

If the log leaves us with nothing, we can either use **traceoptions** for BGP or monitor **traffic interfaces *<interface name>* detail no-resolve**

Be aware of the fact that hidden routes can be in that state for the following reasons:

- next-hop is not reachable
- policy rejects the route

© 2020 Juniper Networks, Inc. All rights reserved.

The following commands are helpful troubleshooting and verifying BGP:

show bgp neighbor
show route advertising-protocol bgp
show route receive-protocol bgp
show route protocol bgp
show route hidden
show route resolution unresolved

Device	Interface	Ipv4 Address	Ipv6 Address	Description
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128
		ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64
		ge-0/0/6	172.16.3.2/24	Link local
vMX3	vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128
		ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64
		ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64
vMX4	vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128
		ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64
		ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table 36 BGP peering information

© 2020 Juniper Networks, Inc. All rights reserved.

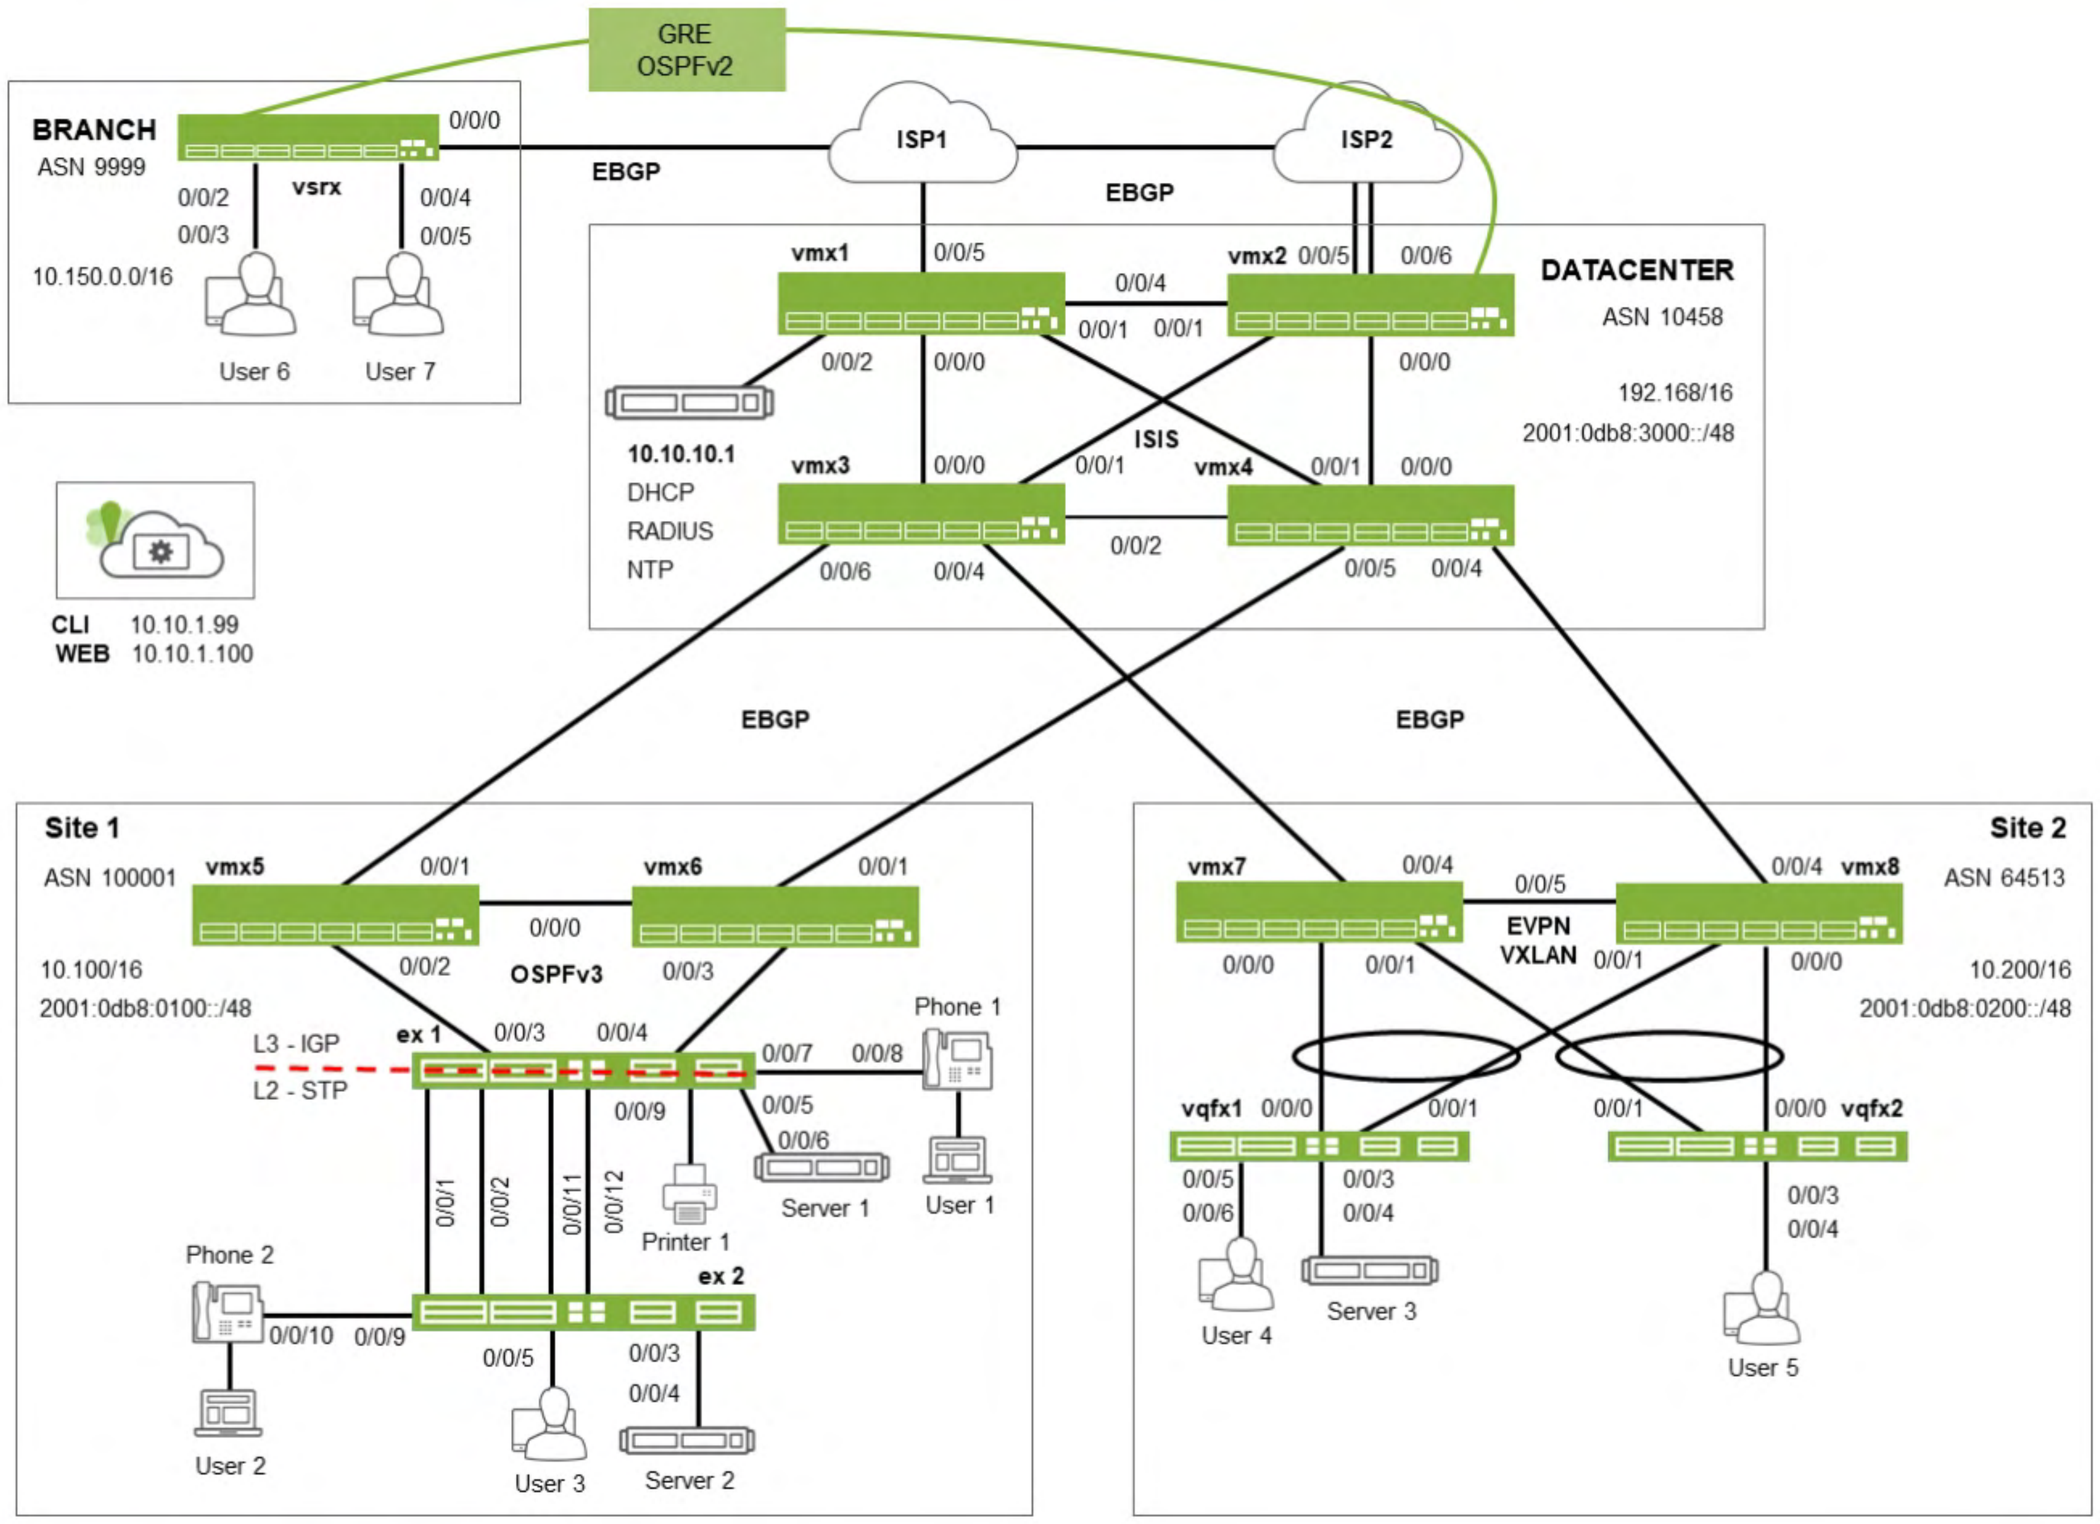

Task 1: IBGP Data Center

1) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have a IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and also accept routes received from each other.
2) Do not configure neighbor statements for IBGP sessions on vMX1 and vMX2.
3) Enable BFD with authentication for IBGP sessions in the data center. Allow multihop sessions. Do not configure BFD between the Route Reflectors.

vMX1:

```
lab@vmx1# show routing-options autonomous-system
10458;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.1;
    cluster 192.168.0.1;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.1;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::2;
}

lab@vmx1# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

```
lab@vmx2# show routing-options autonomous-system
10458;

lab@vmx2# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.2;
    cluster 192.168.0.2;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.2;
    neighbor 192.168.0.1;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.2;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::1;
}

lab@vmx2# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

vMX3:

```
lab@vmx3# show routing-options autonomous-system
10458;

lab@vmx3# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.3;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}

lab@vmx3# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

vMX4:

```
lab@vmx4# show routing-options autonomous-system
10458;

lab@vmx4# show protocols bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
group ibgp {
    type internal;
    local-address 192.168.0.4;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}

lab@vmx4# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

Verification:

```
lab@vmx1# run show bgp summary | match 10458
192.168.0.2           10458       1075       1514       0       0     2:01:25 0/2/2/0
0/0/0/0
192.168.0.3           10458        268        271       0       0     2:00:11 0/0/0/0
0/0/0/0
192.168.0.4           10458        266        269       0       0     1:59:23 0/0/0/0
0/0/0/0
2001:db8:3000::2       10458        418        573       0       0      2:01:17 Establ
2001:db8:3000::3       10458        266        266       0       0      1:59:55 Establ
2001:db8:3000::4       10458        265        265       0       0      1:59:07 Establ
```

vMX2, vMX3, and vMX4 should provide similar outputs

© 2020 Juniper Networks, Inc. All rights reserved.

4) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS and the Branch Office.
5) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down
   vMX2 should not advertise the default route if the BGP peering with ISP2 goes down

vMX1:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}
```

```
lab@vmx1# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.1.9;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx1# show policy-options policy-statement 2ibgp
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}
term 4 {
    then reject;
}
```

```
lab@vmx1# show protocols bgp group ibgp export
export 2ibgp;
```

vMX2:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.200.176;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx2# show policy-options policy-statement 2ibgpandospfv2
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol ospf;
        route-filter 10.150.99.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        accept;
    }
}
```

```
lab@vmx2# show protocols bgp
group ibgp {
    export 2ibgpandospfv2;
}
```

Verification:

Note: The EBGP sessions with ISP1 and ISP2 must be configured before you can test this solution. These outputs are taken after completing Step 1 and 2 for Task 3 in this chapter.

```
lab@vmx1# run show route
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 245 destinations, 470 routes (245 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 00:00:08
                    > to 172.16.1.9 via ge-0/0/5.0

lab@vmx1# run show bgp summary | match 172.16.1.9
172.16.1.9             1111        417        216       0       0     1:36:47
222/222/222/0           0/0/0/0

lab@vmx1# run show route advertising-protocol bgp 192.168.0.3

inet.0: 245 destinations, 470 routes (245 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               172.16.1.9                   100        I

lab@vmx1# set interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete


lab@vmx1# run show bgp summary | match 172.16.1.9
172.16.1.9             1111        419        218       0       1          30 Idle

lab@vmx1# run show route 0/0 exact
```

Since the EBGP peering with ISP1 is down, the generate route is not installed in the routing-table due to the applied policy.

```
lab@vmx1# delete interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete
```

6) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the AS.

vMX1 & vMX2:

```
lab@vmx1# show routing-options
rib inet6.0 {
    static {
        route 0::/0 {
            discard;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            no-install;
        }
    }
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        accept;
    }
}
term 3 {
    then reject;
}

lab@vmx1# show protocols bgp group ibgp-v6 export
export 2ibgp-v6;
```

Verification:

```
lab@vmx1# run show route advertising-protocol bgp 2001:db8:3000::3

inet6.0: 102 destinations, 182 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                         150        I
```

Configurations for Task Completion:

vMX1:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.1
set protocols bgp group ibgp cluster 192.168.0.1
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.1
set protocols bgp group rr neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.1
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::1
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.1.9
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set policy-options policy-statement 2ibgp term 3 from protocol aggregate
set policy-options policy-statement 2ibgp term 3 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2ibgp term 3 then accept
set policy-options policy-statement 2ibgp term 4 then reject
set protocols bgp group ibgp export 2ibgp
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
```

vMX2:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.2
set protocols bgp group ibgp cluster 192.168.0.2
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.2
set protocols bgp group rr neighbor 192.168.0.1
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.2
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::1
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.200.176
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set protocols bgp group ibgp export 2ibgpandospfv2
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
```

vMX3:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.3
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::3
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
```

vMX4:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.4
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::4
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
```

Task 2: IBGP Site 1 and Site 2

1) Configure IBGP in Site1 between vMX5 and vMX6 and in Site 2 between vMX7 and vMX8. Use the information as shown in the table above.

Site 1:

vMX5:

```
lab@vmx5# show routing-options
autonomous-system 100001;

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
    log-updown;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    neighbor 2001:0db8:0100:0000::6;
}
```

vMX6:

```
lab@vmx6# show routing-options
autonomous-system 100001;

lab@vmx6# show protocols bgp
group ibgp {
    type internal;
    neighbor 10.100.0.5 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    neighbor 2001:0db8:0100:0000::5;
}
```

Site 2:

vMX7:

```
lab@vmx7# show routing-options
autonomous-system 64513;

lab@vmx7# show protocols bgp
group ibgp {
    type internal;
    local-address 10.200.0.7;
    family inet {
        unicast;
    }
    neighbor 10.200.0.8;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::7;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::8;
}
```

vMX8:

```
lab@vmx8# show routing-options
autonomous-system 64513;

lab@vmx8# show protocols bgp
group ibgp {
    type internal;
    local-address 10.200.0.8;
    family inet {
        unicast;
    }
    neighbor 10.200.0.7;
}
group ibgp-v6 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    type internal;
    local-address 2001:0db8:3000:0000::8;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::7;
}
```

Verification:

vMX5:

```
lab@vmx5# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                       2          1          0          0          0          0
inet.0
                       8          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.100.0.6           100001       2064       2064       0       0    15:31:15 0/0/0/0
0/0/0/0
2001:db8:100::6      100001       2052       2051       0       0    15:31:15 Establ
  inet6.0: 0/0/0/0
```

vMX7:

```
lab@vmx7# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       8          4          0          0          0          0
bgp.evpn.0
                       0          0          0          0          0          0
inet6.0
                       6          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.8           64513       2003       2004       0       0    15:03:36 Establ
  inet.0: 0/0/0/0
  bgp.evpn.0: 0/0/0/0
2001:db8:3000::8       64513       1993       1993       0       0    15:03:39 Establ
  inet6.0: 0/0/0/0
```

Configurations for Task Completion:

Site 1:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX5:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp log-updown
set protocols bgp group ibgp neighbor 10.100.0.6 local-address 10.100.0.5
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::5
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::6
```

vMX6:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp log-updown
set protocols bgp group ibgp neighbor 10.100.0.5 local-address 10.100.0.6
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::6
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::5
```

Site 2:

vMX7:

```
set routing-options autonomous-system 64513
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.200.0.7
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp neighbor 10.200.0.8
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::7
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::8
```

vMX8:

```
set routing-options autonomous-system 64513
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.200.0.8
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp neighbor 10.200.0.7
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::8
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::7
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: EBGP Data Center

1) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
2) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.

vMX1:

```
lab@vmx1# show protocols bgp
group isp1 {
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
group isp1-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0db9:4000:0001::9 {
        peer-as 1111;
   }
```

Verification:

```
lab@vmx1# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Unconfigured peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                      80         79          0          0          0          0
inet.0
                     226        224          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.9             1111        253         50       0       2       20:45
222/222/222/0        0/0/0/0
2001:db9:4000:1::9        1111        120         47       0       2       20:31 Establ
  inet6.0: 78/78/78/0
```

vMX2:

The task states to establish one IPv4 peering with ISP2 and we need to provide load balancing using both links towards ISP2. As ISP2 is outside of our control it must be preconfigured. This means that if we can debug the BGP packets received from ISP2 we

© 2020 Juniper Networks, Inc. All rights reserved.

might be able to find out valuable information to configure the BGP peering. Note that the BGP packets received from ISP2 can come in on interface ge-0/0/5 or interface ge-0/0/6. Keep this in mind when debugging the BGP packets.

In this example we are going to use the “monitor traffic” command to verify incoming BGP packets destinated to the control plane of vMX2.

```
lab@vmx2# run monitor traffic interface ge-0/0/5 extensive matching tcp
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on ge-0/0/5, capture size 1514 bytes
...
        PFE proto 2 (ipv4): (tos 0xc0, ttl  64, id 53971, offset 0, flags [none], proto:
TCP (6), length: 64) 172.16.200.176.55517 > 192.168.0.2.bgp: S 3936139201:3936139201(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 3639125731 0,sackOK,eol>
23:50:28.618250 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
```

The output above shows that we received an incoming tcp packet from 172.16.200.176 destined to the loopback address of vMX2 (192.168.0.2), port BGP.

We now know the source and destination address for the IPv4 BGP peering with ISP2. Due to the load balancing requirement we need to make sure that IP 172.16.200.176 can be reached using both interfaces towards ISP2. For this we need to configure a static route

```
lab@vmx2# show routing-options
static {
    route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ];
}
```

Now configure the BGP peering with ISP2. Keep in mind that the “multihop” statement is needed as EBGP sessions by default are single-hop sessions (TTL checking).

```
lab@vmx2# show protocols bgp
group isp2 {
    neighbor 172.16.200.176 {
        multihop;
        peer-as 2222;
    }
}
```

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called per-packet load balancing. The naming may be counter-intuitive. However, Junos per-packet load balancing

© 2020 Juniper Networks, Inc. All rights reserved.

is functionally equivalent to what other vendors may term per-flow load balancing. Configure load balancing to spread traffic across both links towards ISP2.

```
lab@vmx2# show policy-options
policy-statement lbpp {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}
...
```

```
lab@vmx2# show routing-options forwarding-table
export lbpp;
```

Verification:

```
lab@vmx2# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.200.176         2222        238         36       0       0       14:39
222/222/222/0         0/0/0/0

lab@vmx2# run show route receive-protocol bgp 172.16.200.176

inet.0: 250 destinations, 255 routes (250 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 1.84.160.0/20           172.16.200.176                          2222 1111 33112 I
...

[edit]
lab@vmx2# run show route forwarding-table destination 1.84.160.1
...
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
1.84.160.0/20      user     0                    indr  1048575   225
                                                 ulst  1048574     4

                              172.16.2.9         ucst      660     3 ge-0/0/5.0
                              172.16.3.9         ucst      659     4 ge-0/0/6.0
...
```

3) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use

© 2020 Juniper Networks, Inc. All rights reserved.

troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.

First configure the IPv6 peering with ISP2 for the ge-0/0/5 link.

vMX2:

```
lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
}
```

The IPv6 peering over the ge-0/0/6 interface must be established using the IPv6 link local address. We need to find out the peering addresses using the "`monitor traffic interface`" command.

```
lab@vmx2# run monitor traffic interface ge-0/0/6 no-resolve matching tcp
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on ge-0/0/6, capture size 96 bytes

10:51:52.548899 Out IP6 fe80::250:56ff:fea9:43ac.179 > fe80::250:56ff:fea9:7fec.61600: R
1293636104:1293636104(0) win 0
10:51:52.548993  In IP6 fe80::250:56ff:fea9:7fec.61600 > fe80::250:56ff:fea9:43ac.179: F
91:91(0) ack 52 win 17085 <nop,nop,timestamp 2230277161[|tcp]>
^C

lab@vmx2# run show interfaces ge-0/0/6 | match fe80
        Destination: fe80::/64, Local: fe80::250:56ff:fea9:43ac
```

```
lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
    neighbor fe80::250:56ff:fea9:7fec {
        local-interface ge-0/0/6.0;
        peer-as 2222;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

As link-local addresses are only unique per interface, it is required to specify the "local-interface" command for the peering so that Junos known on which interface the peering must be established.

Verification:

```
lab@vmx2# run show bgp summary
Groups: 6 Peers: 9 Down peers: 0
Unconfigured peers: 4
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                       234         78          0          0          0          0
inet.0
                       356        222          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
...
2001:dba:4000:2::9         2222          84           9       0       0        3:38 Establ
  inet6.0: 78/78/78/0
fe80::250:56ff:fea9:7fec         2222          76           2       0       0           2
Establ
  inet6.0: 0/78/78/0
```

4) Advertise the IPv4 summary routes received from Site 1, Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs.

vMX1:

```
lab@vmx1# show routing-options aggregate
route 192.168.0.0/16;

lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        accept;
    }
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp
group isp1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    export 2ebgp;
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
```

Verification:

Note: The following verification output is taken after everything through step 10 has been completed. You must have BGP established between the data center and Site 1 and Site 2.

```
lab@vmx1# run show route advertising-protocol bgp 172.16.1.9

inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.0.0/16           Self                                    100001 I
* 10.200.0.0/16           Self                                    I
* 192.168.0.0/16          Self                                    I
```

vMX2:

```
lab@vmx2# show routing-options aggregate
route 192.168.0.0/16;

[edit]
lab@vmx2# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vmx2# show protocols bgp group isp2
export 2ebgp;
neighbor 172.16.200.176 {
    multihop;
    peer-as 2222;
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

Note: The following verification output is taken after everything through step 10 has been completed. You must have BGP established between the data center and Site 1 and Site 2.

```
lab@vmx2# run show route advertising-protocol bgp 172.16.200.176

inet.0: 251 destinations, 255 routes (251 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.0.0/16           Self                                    100001 I
* 10.200.0.0/16           Self                                    I
* 192.168.0.0/16          Self                                    I
```

5) Advertise a single /32 route to both ISPs that encompasses the IPv6 network of the data center, Site 1, and Site 2.

vMX1:

```
lab@vmx1# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;

lab@vmx1# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp group isp1-v6
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0db9:4000:0001::9 {
    peer-as 1111;
}
```

vMX2:

```
lab@vmx2# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;

lab@vmx2# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx2# show protocols bgp group isp2-v6
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0dba:4000:0002::9 {
    peer-as 2222;
}
neighbor fe80::250:56ff:fea9:7fec {
    local-interface ge-0/0/6.0;
    peer-as 2222;
}
```

Verification:

vMX1:

```
lab@vmx1# run show route advertising-protocol bgp 2001:0db9:4000:0001::9

inet6.0: 102 destinations, 181 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I
```

vMX2:

```
lab@vmx2# run show route advertising-protocol bgp 2001:0dba:4000:0002::9

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I

[edit]
lab@vmx2# run show route advertising-protocol bgp fe80::250:56ff:fea9:7fec

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I
```


© 2020 Juniper Networks, Inc. All rights reserved.

6) Ensure that ISP2 is the primary path to all IPv4 address blocks announced by your AS. Both ISPs apply route policy based on received communities. Use the information in the table below to meet the task requirement.

ISP	COMMUNITY	ISP ACTION
ISP1	1111:1	Prepend 1x
	1111:2	Prepend 2x
	1111:3	Prepend 3x
ISP2	2222:1	Prepend 1x
	2222:2	Prepend 2x
	2222:3	Prepend 1x

vMX1:

```
lab@vmx1# show policy-options community prepend
members 1111:2;

lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        community add prepend;
        accept;
    }
}
term 2 {
    then reject;
}
```

Verification:

Note: You can review the summary route on the vr-device using the *lab/lab123* credentials.

VR-Device:

```
lab@vr-device> show route 192.168.0.0

ISP1.inet.0: 232 destinations, 235 routes (232 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 01:07:01, localpref 100
                      AS path: 2222 10458 I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    >  to 10.20.20.2 via lt-0/0/0.1
                    [BGP/170] 02:36:50, localpref 100
                      AS path: 1111 1111 10458 I, validation-state: unverified
                    >  to 172.16.1.1 via ge-0/0/2.0
...
```

7) Ensure that ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.

vMX1:

```
lab@vmx1# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}
lab@vmx1# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
```

vMX2:

```
lab@vmx2# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}

lab@vmx2# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        accept;
    }
}

lab@vmx2# show protocols bgp group isp2 import
import fromisp2;
```

Verification:

```
lab@vmx1# run show route protocol bgp

inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.84.160.0/20      *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 1111 33112 I, validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 33112 I, validation-state: unverified
                    > to 172.16.1.9 via ge-0/0/5.0
1.176.0.0/12       *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 172.16.1.9 via ge-0/0/5.0
3.189.0.0/16       *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                    > to 172.16.1.9 via ge-0/0/5.0
...
```

8) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2. Prefixes from Site1 that are advertised to both ISPs should contain the ASN of Site 1.

vMX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show protocols bgp | display set | match remove
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
```

vMX2:

```
lab@vmx2# show protocols bgp | display set | match remove-private
set protocols bgp group isp2 remove-private
set protocols bgp group isp2-v6 remove-private
```

9) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Site 1 and Site 2.

vMX1:

```
lab@vmx1# show policy-options community no-export
members no-export;

lab@vmx1# show policy-options policy-statement fromisp1
term 1 {
    then {
        community set no-export;
        accept;
    }
}

lab@vmx1# show protocols bgp group isp1 import
import fromisp1;
```

vMX2:

```
lab@vmx2# show policy-options community no-export
members no-export;

lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        community set no-export;
        accept;
    }
}
```

Verification:

vMX3:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx3# run show route protocol bgp detail

inet.0: 247 destinations, 249 routes (247 active, 0 holddown, 0 hidden)
...

1.84.160.0/20 (1 entry, 1 announced)
        *BGP    Preference: 170/-151
                Next hop type: Indirect, Next hop index: 0
                Address: 0xce3dd50
                Next-hop reference count: 450
                Source: 192.168.0.2
                Next hop type: Router, Next hop index: 643
                Next hop: 192.168.23.2 via ge-0/0/1.0, selected
                Session Id: 0x238
                Protocol next hop: 192.168.0.2
                Indirect next hop: 0xba3f580 1048574 INH Session ID: 0x256
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 1:18:42    Metric2: 10
                Validation State: unverified
                ORR Generation-ID: 0
                Task: BGP_10458.192.168.0.2
                Announcement bits (2): 0-KRT 4-Resolve tree 4
                AS path: 2222 1111 33112 I
                Communities: no-export
                Accepted
                Localpref: 150
                Router ID: 192.168.0.2
...
```

vMX4:

Note: This sample verification output from vMX4 was taken after the BGP peering is setup between the data center and Site 1. This step is next, you can verify this output after completing step 10.

```
lab@vmx4# run show route advertising-protocol bgp 192.168.46.6

inet.0: 247 destinations, 250 routes (247 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    I
  10.10.10.0/24           Self                                    I
* 10.200.0.0/16           Self                                    64513 I
* 192.168.0.0/16          Self                                    I
```

10) Site 1 and Site 2 should learn the summarized prefixes of the other site, data center, and branch office. Site 1 and Site 2 should also learn the IPv4 and IPv6

© 2020 Juniper Networks, Inc. All rights reserved.

default route and the management network 10.10.10/24 that originated in the data center.

vMX1:

```
lab@vmx1# show policy-options policy-statement 2ibgp
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol direct;
        route-filter 10.10.10.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}
term 4 {
    then reject;
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then accept;
}
term 3 {
    then reject;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

```
lab@vmx2# show policy-options policy-statement 2ibgpandospfv2
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol ospf;
        route-filter 10.150.99.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        local-preference 150;
        accept;
    }
}

[edit]
lab@vmx2# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        local-preference 150;
        accept;
    }
}
term 3 {
    then reject;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

vMX3:

```
[edit]
lab@vmx3# show policy-options
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx3# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.3;
    export nhs;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site1 {
    neighbor 192.168.35.5 {
        peer-as 100001;
    }
}
group site2 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    neighbor 192.168.37.7 {
        peer-as 64513;
    }
}
group site1-v6 {
    neighbor 2001:0db8:3000:0035::5 {
        family inet6 {
            unicast;
        }
        peer-as 100001;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0037::7 {
        family inet6 {
            unicast;
        }
        peer-as 64513;
    }
}
```

vMX4:

```
lab@vmx4# show policy-options
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}
```

```
lab@vmx4# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.4;
    export nhs;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site1 {
    neighbor 192.168.46.6 {
        peer-as 100001;
    }
}
group site2 {
    neighbor 192.168.48.8 {
        peer-as 64513;
    }
}
group site1-v6 {
    neighbor 2001:0db8:3000:0046::6 {
        family inet6 {
            unicast;
        }
        peer-as 100001;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0048::8 {
        family inet6 {
            unicast;
        }
        peer-as 64513;
    }
}
```

vMX5:

```
lab@vmx5# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0100::/48;
    }
}

lab@vmx5# show policy-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
policy-statement 2bgp {
    term 1 {
        from {
            protocol ospf3;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0100::/48 exact;
        }
        then accept;
    }
}
policy-statement 2ospf {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then accept;
    }
}

lab@vmx5# show protocols ospf3
realm ipv4-unicast {
    spf-options {
        delay 50;
    }
    database-protection {
        maximum-lsa 1000;
        warning-only;
    }
    export 2ospf;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/0.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
            flood-reduction;
        }
    }
    area 0.0.0.1 {
        area-range 10.100.0.0/16;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        interface ge-0/0/2.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
export 2ospf;
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
        flood-reduction;
    }
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
    log-updown;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    neighbor 2001:0db8:0100:0000::6;
}
group dc {
    export 2bgp;
    neighbor 192.168.35.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0035::3 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vMX6:

```
lab@vmx6 show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0100::/48;
    }
}

lab@vmx6# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol ospf3;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0100::/48 exact;
        }
        then accept;
    }
}
policy-statement 2ospf {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then accept;
    }
}

lab@vmx6# show protocols ospf3
realm ipv4-unicast {
    spf-options {
        delay 50;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    database-protection {
        maximum-lsa 1000;
        warning-only;
    }
    export 2ospf;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/0.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
            flood-reduction;
        }
    }
    area 0.0.0.1 {
        area-range 10.100.0.0/16;
        interface ge-0/0/3.0 {
            interface-type p2p;
            hello-interval 1;
            dead-interval 4;
        }
    }
}
export 2ospf;
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
        flood-reduction;
    }
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/3.0 {
        interface-type p2p;
        hello-interval 1;
        dead-interval 4;
    }
}

lab@vmx6# show protocols bgp
group ibgp {
    type internal;
    log-updown;
    neighbor 10.100.0.5 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    neighbor 2001:0db8:0100:0000::5;
}
group dc {
    export 2bgp;
    neighbor 192.168.46.4 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0046::4 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vSRX1:

```
lab@vsrx# show routing-options
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
aggregate {
    route 10.150.0.0/16;
}
```

```
lab@vsrx# show policy-options policy-statement 2bgp
term 1 {
    from {
        route-filter 10.150.0.0/16 exact;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vsrx# show policy-options policy-statement frombgp
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vsrx# show protocols bgp
group ebgp {
    import frombgp;
    export 2bgp;
    neighbor 10.1.1.1 {
        peer-as 1111;
        local-as 9999;
    }
}
```

Verification Branch Office:

vSRX:

```
lab@vsrx# 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
                       224          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.1               1111        207          3       0       0        1:14 1/224/1/0
0/0/0/0

[edit]
lab@vsrx# run show route protocol bgp

inet.0: 237 destinations, 238 routes (14 active, 0 holddown, 223 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 00:05:27, localpref 100
                      AS path: 1111 2222 10458 I, validation-state: unverified
                    >  to 10.1.1.1 via ge-0/0/0.0
...

lab@vsrx# run show route advertising-protocol bgp 10.1.1.1

inet.0: 237 destinations, 238 routes (14 active, 0 holddown, 223 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.150.0.0/16           Self                                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification Site 1:

vMX6:

```
lab@vmx6# run show route protocol bgp

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

0.0.0.0/0           [BGP/170] 00:12:07, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 00:16:36, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
                    > to 10.100.56.5 via ge-0/0/0.0
10.10.10.0/24       [BGP/170] 00:12:07, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 00:16:36, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
                    > to 10.100.56.5 via ge-0/0/0.0
10.150.99.0/24     *[BGP/170] 00:02:23, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
192.168.0.0/16      [BGP/170] 00:12:07, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 00:16:36, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
                    > to 10.100.56.5 via ge-0/0/0.0

inet6.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                [BGP/170] 00:12:16, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:46::4 via ge-0/0/1.0
                    [BGP/170] 00:16:32, localpref 100, from 2001:db8:100::5
                      AS path: 10458 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4cfd via ge-0/0/0.0
```

It seems we are missing the 2001:db8::/32 aggregate network from the data center. Verify if vMX4 announces the route to Site1.

vMX4:

```
lab@vmx4# run show route advertising-protocol bgp 2001:db8:3000:46::6
inet6.0: 27 destinations, 29 routes (27 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    I
* 2001:db8::/32           Self                                    {100001} I
```

The IPv6 aggregate route is announced by vMX4 to Site1, but not accepted by Site1. It seems that vMX1 and vMX2 are aggregating the 2001:db8::/32 network using AS 100001 contributing routes. This is the reason that Site1 does not accept the route. To fix this we need to ensure that vMX1 and vMX2 create the aggregate route 2001:db8::/32 based on local contributing routes. We can do this with route policy as shown below.

vMX1:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}


lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
}
```

vMX2:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}


lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

Verification:

vMX4:

```
lab@vmx4# run show route advertising-protocol bgp 2001:db8:3000:46::6

inet6.0: 27 destinations, 29 routes (27 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    I
* 2001:db8::/32           Self                                    I
```

vMX6:

```
lab@vmx6# run show route receive-protocol bgp 2001:db8:3000:46::4

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

inet6.0: 23 destinations, 27 routes (23 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  ::/0                    2001:db8:3000:46::4                     10458 I
  2001:db8::/32           2001:db8:3000:46::4                     10458 I
```

vMX7:

```
lab@vmx7# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}
```

```
lab@vmx7# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx7# show protocols bgp
group ibgp {
    type internal;
    local-address 10.200.0.7;
    family inet {
        unicast;
    }
    export nhs;
    neighbor 10.200.0.8;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::7;
    family inet6 {
        unicast;
    }
    export nhs;
    neighbor 2001:0db8:3000:0000::8;
}
group dc {
    export 2bgp;
    neighbor 192.168.37.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0037::3 {
        family inet6 {
            unicast;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        peer-as 10458;
    }
}
```

vMX8:

```
lab@vmx8# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}

lab@vmx8# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx8# show protocols bgp
group ibgp {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        type internal;
        local-address 10.200.0.8;
        family inet {
            unicast;
        }
        export nhs;
        neighbor 10.200.0.7;
    }
    group ibgp-v6 {
        type internal;
        local-address 2001:0db8:3000:0000::8;
        family inet6 {
            unicast;
        }
        export nhs;
        neighbor 2001:0db8:3000:0000::7;
    }
    group dc {
        export 2bgp;
        neighbor 192.168.48.4 {
            peer-as 10458;
        }
    }
    group dc-v6 {
        export 2bgp-v6;
        neighbor 2001:0db8:3000:0048::4 {
            family inet6 {
                unicast;
            }
            peer-as 10458;
        }
    }
```

Verification Site 2:

```
lab@vmx7# run show route protocol bgp

inet.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:05:21, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 00:00:48, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.10.10.0/24      *[BGP/170] 00:05:21, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    [BGP/170] 00:00:48, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.100.0.0/16       *[BGP/170] 00:05:21, localpref 100
                      AS path: 10458 100001 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 00:00:48, localpref 100, from 10.200.0.8
                      AS path: 10458 100001 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.150.99.0/24      *[BGP/170] 00:05:21, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 00:00:48, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
192.168.0.0/16      *[BGP/170] 00:05:21, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 00:00:48, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0

inet6.0: 15 destinations, 18 routes (14 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                *[BGP/170] 00:05:17, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
                    [BGP/170] 00:00:44, localpref 100, from 2001:db8:3000::8
                      AS path: 10458 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
2001:db8::/32       *[BGP/170] 00:05:17, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
                    [BGP/170] 00:00:44, localpref 100, from 2001:db8:3000::8
                      AS path: 10458 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
2001:db8:100::/48   *[BGP/170] 00:05:17, localpref 100
                      AS path: 10458 100001 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
                    [BGP/170] 00:00:44, localpref 100, from 2001:db8:3000::8
                      AS path: 10458 100001 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
```

Configurations for Task Completion:

vMX1:

```
set protocols bgp group isp1 neighbor 172.16.1.9 peer-as 1111
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group isp1-v6 family inet6 unicast
set protocols bgp group isp1-v6 neighbor 2001:0db9:4000:0001::9 peer-as 1111
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp1 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp1-v6 export 2ebgp-v6
set policy-options community prepend members 1111:2
set policy-options policy-statement 2ebgp term 1 then community add prepend
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
set policy-options community no-export members no-export
set policy-options policy-statement fromisp1 term 1 then community set no-export
set policy-options policy-statement fromisp1 term 1 then accept
set protocols bgp group isp1 import fromisp1
set policy-options policy-statement 2ibgp term 1 from protocol aggregate
set policy-options policy-statement 2ibgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ibgp term 1 then accept
set policy-options policy-statement 2ibgp term 2 from protocol direct
set policy-options policy-statement 2ibgp term 2 from route-filter 10.10.10.0/24 exact
set policy-options policy-statement 2ibgp term 2 then accept
insert policy-options policy-statement 2ibgp term 1 before term 3
insert policy-options policy-statement 2ibgp term 2 before term 3
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::/48 orlonger
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX2:

```
set routing-options static route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ]
set protocols bgp group isp2 neighbor 172.16.200.176 multihop
set protocols bgp group isp2 neighbor 172.16.200.176 peer-as 2222
set policy-options policy-statement lbpp term 1 then load-balance per-packet
set routing-options forwarding-table export lbpp
set protocols bgp group isp2-v6 family inet6 unicast
set protocols bgp group isp2-v6 neighbor 2001:0dba:4000:0002::9 peer-as 2222
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec local-interface ge-
0/0/6.0
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec peer-as 2222
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp2 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp2-v6 export 2ebgp-v6
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set policy-options policy-statement fromisp2 term 1 then local-preference 150
set policy-options policy-statement fromisp2 term 1 then accept
set protocols bgp group isp2 import fromisp2
set protocols bgp group isp2 remove-private
set protocols bgp group isp2-v6 remove-private
set policy-options community no-export members no-export
set policy-options policy-statement fromisp2 term 1 then community set no-export
set policy-options policy-statement 2ibgpandospfv2 term 3 then local-preference 150
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ibgp-v6 term 2 then local-preference 150
set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::2/128 exact
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX3:

```
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
set protocols bgp group ibgp export nhs
set protocols bgp group site1 neighbor 192.168.35.5 peer-as 100001
set protocols bgp group site2 neighbor 192.168.37.7 peer-as 64513
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0035::5 family inet6 unicast
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0035::5 peer-as 100001
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 peer-as 64513
```

vMX4:

```
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
set protocols bgp group ibgp export nhs
set protocols bgp group site1 neighbor 192.168.46.6 peer-as 100001
set protocols bgp group site2 neighbor 192.168.48.8 peer-as 64513
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0046::6 family inet6 unicast
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0046::6 peer-as 100001
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 peer-as 64513
```

vMX5:

```
set routing-options rib inet6.0 aggregate route 2001:0db8:0100::/48
set policy-options policy-statement 2bgp term 1 from protocol ospf3
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement 2ospf term 1 from protocol bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ospf term 1 from external
set policy-options policy-statement 2ospf term 1 then accept
set protocols ospf3 realm ipv4-unicast export 2ospf
set protocols ospf3 export 2ospf
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.35.3 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 peer-as 10458
```

vMX6:

```
set routing-options rib inet6.0 aggregate route 2001:0db8:0100::/48
set policy-options policy-statement 2bgp term 1 from protocol ospf3
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement 2ospf term 1 from protocol bgp
set policy-options policy-statement 2ospf term 1 from external
set policy-options policy-statement 2ospf term 1 then accept
set protocols ospf3 realm ipv4-unicast export 2ospf
set protocols ospf3 export 2ospf
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.46.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 peer-as 10458
```

vMX7:

```
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.37.3 peer-as 10458
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 peer-as 10458
```

vMX8:

```
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.48.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 peer-as 10458
```

vSRX:

```
set routing-options aggregate route 10.150.0.0/16
set policy-options policy-statement 2bgp term 1 from route-filter 10.150.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp term 2 then reject
set policy-options policy-statement frombgp term 1 from route-filter 192.168.0.0/16
exact
set policy-options policy-statement frombgp term 1 then accept
set policy-options policy-statement frombgp term 2 then reject
set protocols bgp group ebgp import frombgp
set protocols bgp group ebgp export 2bgp
set protocols bgp group ebgp neighbor 10.1.1.1 peer-as 1111
set protocols bgp group ebgp neighbor 10.1.1.1 local-as 9999
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Chapter 6: EVPN/VXLAN

EVPN Overview

Ethernet VPN (EVPN) is a standards-based technology that provides virtual multipoint bridged connectivity between different Layer 2 domains over an IP or IP/MPLS backbone network. Like other VPN technologies, such as IP VPN and virtual private LAN service (VPLS), EVPN instances are configured on provider edge (PE) routers to maintain logical service separation between customers. The PE routers connect to customer edge (CE) devices, which can be routers, switches, or hosts. The PE routers then exchange reachability information.

The EVPN is used as a Layer 2 overlay solution to provide Layer 2 connection over an IP underlay for the endpoints within a virtual network whenever Layer 2 connectivity is required by an end station such as bare-metal server (BMS). Otherwise, Layer 3 routing is used through VRF tables between Contrail vRouters and MX Series routers. EVPN technology offers multitenancy, flexible services that can be extended on demand, frequently using compute resources of different physical data centers for a single service (Layer 2 extension).

EVPN's MP-BGP control plane enables you to dynamically move live virtual machines from one data center to another, also known as virtual machine (VM) motion. After you move a VM to a destination server or hypervisor, it transmits a gratuitous ARP, which updates the Layer 2 forwarding table of the PE device at the destination data center. The PE device then transmits a MAC route update to all remote PE devices which in turn update their forwarding tables. An EVPN tracks the movement of the VM, which is also known as MAC mobility.

EVPN also has mechanisms that detect and stop MAC flapping, and prevent the looping of broadcast, unknown unicast, and multicast (BUM) traffic in an all-active multi-homed topology.

The EVPN technology, similar to Layer 3 MPLS VPN, includes the concept of routing MAC addresses using IP/MPLS core. EVPN provides the following benefits:

1) Ability to have an active multihomed edge device
2) Aliasing
3) Fast convergence
4) Load balancing across dual-active links

© 2020 Juniper Networks, Inc. All rights reserved.

5) MAC address mobility
6) Multitenancy

In addition, EVPN uses these techniques:

- *Multihoming* provides redundancy in the event that an access link or one of the PE routing devices fails. In either case, traffic flows from the CE device towards the PE router, using the remaining active links. For traffic in the other direction, the remote PE router updates its forwarding table to send traffic to the remaining active PE routers connected to the multihomed Ethernet segment. EVPN provides a fast convergence mechanism, which reduces traffic restoration time so that the time it takes to make this adjustment is independent of the number of media access control (MAC) addresses learned by the PE router. All-active multihoming enables a CE device to connect to two or more PE routers such that traffic is forwarded using all of the links between the devices. This multihoming enables the CE device to load-balance traffic to multiple PE routers. More importantly, multihoming enables a remote PE router to load-balance traffic to the multihomed PE routers across the core network. This load balancing of traffic flows between data centers is known as *aliasing*, which causes different signals to become indistinguishable—they become aliases of one another. Aliasing is used with digital audio and digital images.
- *Split horizon* prevents the looping of broadcast, unknown unicast, and multicast (BUM) traffic in a network. The split horizon basic principle is simple: Information about the routing for a particular packet is never sent back in the direction from which it was received.
- *Local link bias* conserves bandwidth by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a link aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.
- EVPN with VXLAN encapsulation is used for Layer 2 connectivity between virtual machines and a top-of-rack (TOR) switch, for example, a QFX5100 switch, within a Layer 2 domain.

You can use Contrail to provision an MX Series router as a Layer 2 or Layer 3 VXLAN gateway. MX Series routers implement the NETCONF XML management protocol, which is an XML-based protocol that client applications use to request and change configuration information on routing, switching, and security devices. The NETCONF XML management protocol uses an XML based data encoding for the configuration data and remote procedure calls. The NETCONF protocol defines basic operations that are equivalent to configuration mode commands in the command-line interface (CLI). Applications use the

© 2020 Juniper Networks, Inc. All rights reserved.

protocol operations to display, edit, and commit configuration statements similarly to how administrators use CLI configuration mode commands to perform those same operations.

VXLAN Overview

Virtual extensible LANs (VXLANs) introduced an overlay scheme that expands the Layer 2 network address space from 4K to 16 million, largely solving the scaling issues seen in VLAN-based environments.

Network overlays are created by encapsulating traffic and tunneling the traffic over a physical network. You can use a number of tunneling protocols in the data center to create network overlays—the most common protocol is VXLAN. VXLAN tunneling protocol encapsulates Layer 2 Ethernet frames in Layer 3 UDP packets. This encapsulation enables you to create virtual Layer 2 subnets or segments that can span physical Layer 3 networks.

In a VXLAN overlay network, a VXLAN network identifier (VNI) uniquely identifies each Layer 2 subnet or segment. A VNI segments traffic the same way that an IEEE 802.1Q VLAN ID segments traffic. As is the case with VLAN, virtual machines on the same VNI can communicate directly with each other, whereas virtual machines on different VNIs need a router to communicate with each other.

The entity that performs the encapsulation and de-encapsulation is called a VXLAN tunnel endpoint (VTEP). In the physical network, a Juniper Networks device that functions as a Layer 2 or Layer 3 VXLAN gateway can encapsulate and decapsulate data packets. This type of VTEP is known as a hardware VTEP. In the virtual network, VTEPs can reside in hypervisor hosts, such as kernel-based virtual machine (KVM) hosts. This type of VTEP is known as a software VTEP.

Each VTEP has two interfaces.

- One interface is a switching interface that faces the virtual machines in the host and provides communication between VMs on the local LAN segment.
- The other is an IP interface that faces the Layer 3 network.

Each VTEP has a unique IP address that is used for routing the UDP packets between VTEPs. For example, when VTEP1 receives an Ethernet frame from VM1 addressed to VM3, it uses the VNI and the destination MAC to look up in its forwarding table which VTEP sends the packet to. It then adds a VXLAN header that contains the VNI to the Ethernet frame and encapsulates the frame in a Layer 3 UDP packet and routes the packet to VTEP2 over the Layer 3 network. VTEP2 de-encapsulates the original Ethernet frame and forwards it to VM3. VM1 and VM3 cannot detect the VXLAN tunnel and the Layer 3 network between them.

© 2020 Juniper Networks, Inc. All rights reserved.

EVPN-VXLAN Integration Overview

VXLAN defines a tunneling scheme to overlay Layer 2 networks on top of Layer 3 networks. This tunneling scheme allows for optimal forwarding of Ethernet frames with support for multipathing of unicast and multicast traffic with the use of UDP/IP encapsulation for tunneling and is mainly used for the intra-data center site connectivity.

A unique characteristic of EVPN is that MAC address learning between PE routers occurs in the control plane. The local PE router detects a new MAC address from a CE device and then, using MP-BGP, advertises the address to all the remote PE routers. This method differs from existing Layer 2 VPN solutions such as VPLS, which learn by flooding unknown unicast in the data plane. This control plane MAC learning method is the key enabler of the many useful features that EVPN provides.

Because MAC learning is handled in the control plane, EVPN has the flexibility to support different data plane encapsulation technologies between PE routers. This flexibility is important because not every backbone network might be running MPLS, especially in enterprise networks.

EVPN addresses many of the challenges faced by network operators building data centers to offer cloud and virtualization services. The main application of EVPN is Data Center Interconnect (DCI), which refers to the ability to extend Layer 2 connectivity between different data centers that are deployed to improve the performance of delivering application traffic to end users and for disaster recovery.

Although various DCI technologies are available, EVPN has an advantage over the other MPLS technologies because of its unique features, such as active/active redundancy, aliasing, and mass MAC withdrawal. As a result, to provide a solution for DCI, VXLAN is integrated with EVPN.

Each PE node implements EVPN to distribute the client MAC addresses learned over the VXLAN tunnel into BGP. Each PE node encapsulates the VXLAN or Ethernet frames with MPLS when sending the packets over the MPLS core and with the VXLAN tunnel header when sending the packets over the VXLAN network.

© 2020 Juniper Networks, Inc. All rights reserved.

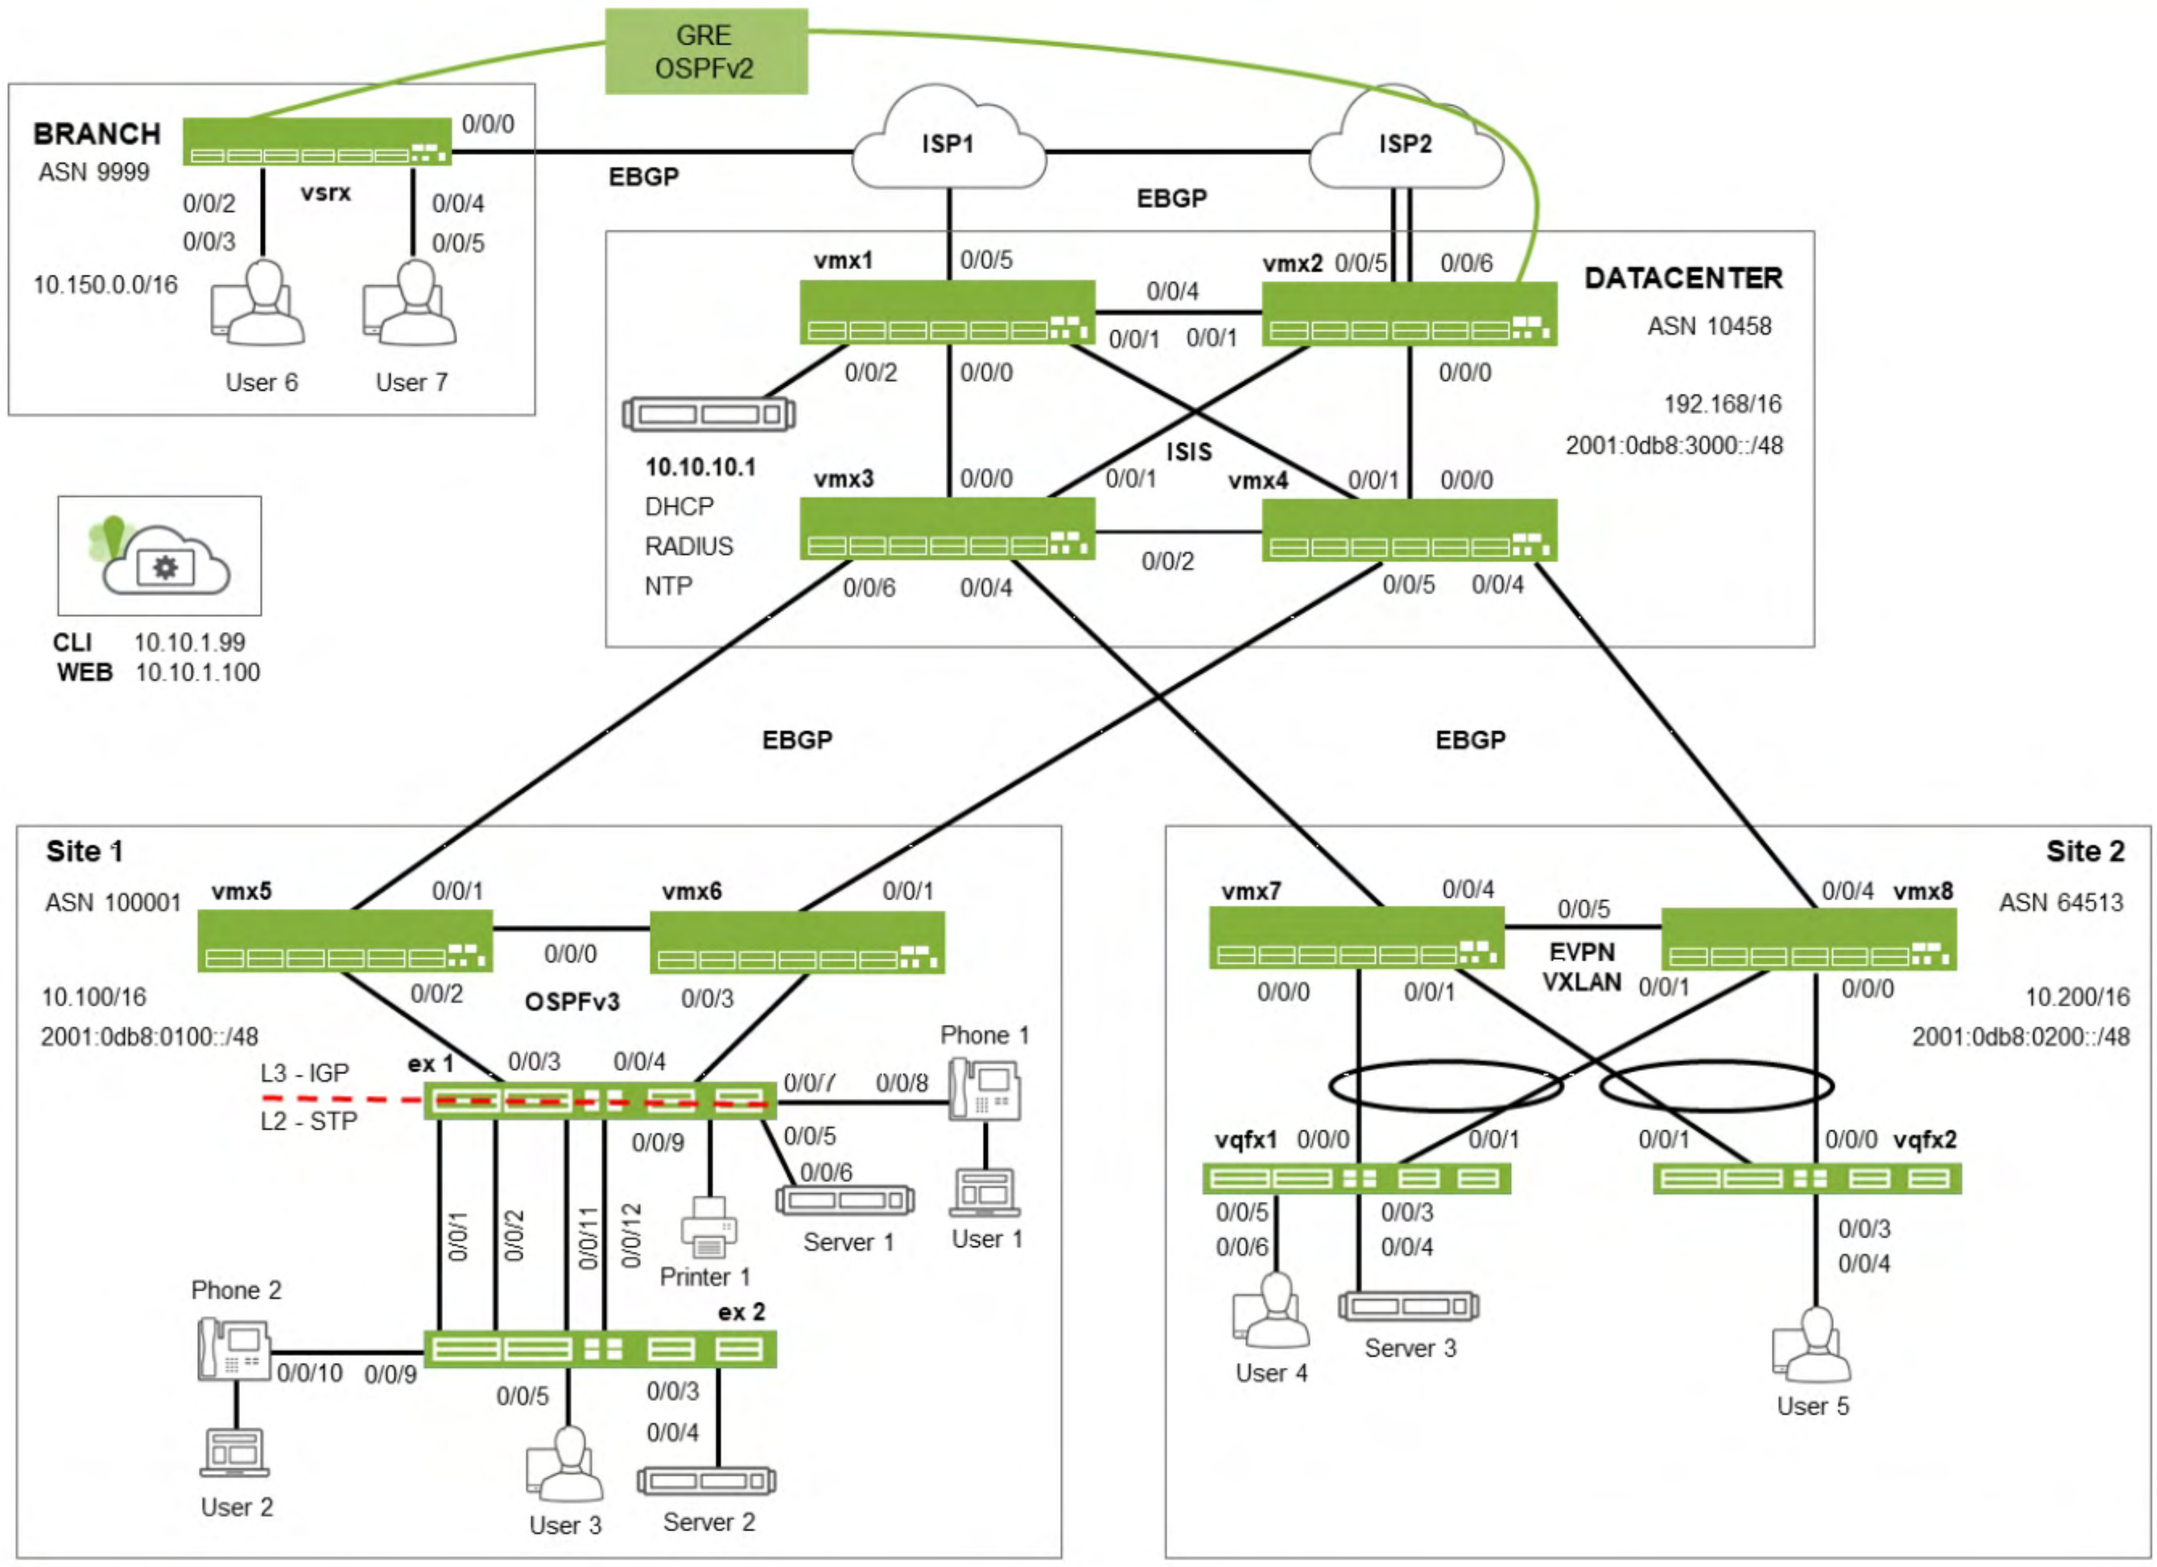

Task 1: EVPN VXLAN

In this task, you will configure the Layer 2 connectivity in Site 2 using the EVPN/VXLAN technology.

1) Configure one EVPN instance on vMX7 and VMX8 that uses VXLAN transport. The EVPN should provide Layer 2 connectivity for the VLANs in Site 2. Ensure that active multi-homing is supported for unicast traffic.
2) Ensure that route targets are automatically generated.

vMX7:

```
lab@vmx7# show routing-options router-id
router-id 10.200.0.7;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx7# show protocols bgp group ibgp
type internal;
local-address 10.200.0.7;
family inet {
    unicast;
}
family evpn {
    signaling;
}
export nhs;
neighbor 10.200.0.8;
```

```
lab@vmx7# show interfaces ae0
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:01;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:01;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}

[edit]
lab@vmx7# show interfaces ae1
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:02;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:02;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx7# show interfaces irb
unit 100 {
    family inet {
        address 10.200.100.250/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::250/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    family inet {
        address 10.200.101.250/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::250/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}

lab@vmx7# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
    interface ae0.0;
    interface ae1.0;
    route-distinguisher 10.200.0.7:1;
    vrf-target {
        target:64513:1;
        auto;
    }
    protocols {
        evpn {
            encapsulation vxlan;
            extended-vni-list all;
            default-gateway do-not-advertise;
        }
    }
    bridge-domains {
        vlan100 {
            vlan-id 100;
            routing-interface irb.100;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            vxlan {
                vni 100;
            }
        }
        vlan101 {
            vlan-id 101;
            routing-interface irb.101;
            vxlan {
                vni 101;
            }
        }
    }
}
```

vMX8:

```
lab@vmx8# show routing-options router-id
router-id 10.200.0.8;

lab@vmx8# show protocols bgp group ibgp
type internal;
local-address 10.200.0.8;
family inet {
    unicast;
}
family evpn {
    signaling;
}
export nhs;
neighbor 10.200.0.7;

lab@vmx8# show interfaces ae0
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:01;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:01;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx8# show interfaces ae1
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:02;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:78:02;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}

lab@vmx8# show interfaces irb
unit 100 {
    family inet {
        address 10.200.100.251/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::251/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    family inet {
        address 10.200.101.251/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::251/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}

lab@vmx8# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
interface ae0.0;
interface ae1.0;
route-distinguisher 10.200.0.8:1;
vrf-target {
    target:64513:1;
    auto;
}
protocols {
    evpn {
        encapsulation vxlan;
        extended-vni-list all;
        default-gateway do-not-advertise;
    }
    bridge-domains {
        vlan100 {
            vlan-id 100;
            routing-interface irb.100;
            vxlan {
                vni 100;
            }
        }
        vlan101 {
            vlan-id 101;
            routing-interface irb.101;
            vxlan {
                vni 101;
            }
        }
    }
}
```

Verification:

vMX7:

```
lab@vmx7# run show bgp summary group ibgp
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.8            64513         42         53       0       1        8:45 Establ
  inet.0: 0/5/5/0
  bgp.evpn.0: 26/26/26/0
  evpn1.evpn.0: 24/24/24/0
  __default_evpn__.evpn.0: 2/2/2/0

lab@vmx7# run show evpn database
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:fc:01:00:00:00:64:00  May 01 09:20:06
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:fc:01:00:00:00:64:00  May 01 09:20:06
2001:db8:200:100::254
     100        02:05:86:71:76:1b  00:78:78:78:78:78:78:78:78:01  May 01 09:40:58
10.200.100.4

2001:db8:200:100::4

fe80::205:86ff:fe71:761b
     100        02:05:86:71:9b:13  00:78:78:78:78:78:78:78:78:02  May 01 09:40:52
10.200.100.5

2001:db8:200:100::5

fe80::205:86ff:fe71:9b13
     100        2c:6b:f5:9c:a0:f0  irb.100                        May 01 09:20:06
10.200.100.250

2001:db8:200:100::250

fe80::2e6b:f500:649c:a0f0
     101        00:00:5e:00:01:01  05:00:00:fc:01:00:00:00:65:00  May 01 09:20:06
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:fc:01:00:00:00:65:00  May 01 09:20:06
2001:db8:200:101::254
     101        02:05:86:71:76:13  00:78:78:78:78:78:78:78:78:01  May 01 09:40:49
10.200.101.3

2001:db8:200:101::3

fe80::205:86ff:fe71:7613
     101        2c:6b:f5:9c:a0:f0  irb.101                        May 01 09:20:06
10.200.101.250

2001:db8:200:101::250

fe80::2e6b:f500:659c:a0f0
```

vMX8:

```
lab@vmx8# run show bgp summary group ibgp
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.7            64513         51         38       0       0        7:19 Establ
  inet.0: 0/5/5/0
  bgp.evpn.0: 20/20/20/0
  evpn1.evpn.0: 18/18/18/0
  __default_evpn__.evpn.0: 2/2/2/0

lab@vmx8# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:fc:01:00:00:00:64:00  May 01 09:20:06
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:fc:01:00:00:00:64:00  May 01 09:20:06
2001:db8:200:100::254
     100        02:05:86:71:76:1b  00:78:78:78:78:78:78:78:78:01  May 01 09:42:03
10.200.100.4

2001:db8:200:100::4

fe80::205:86ff:fe71:761b
     100        02:05:86:71:9b:13  00:78:78:78:78:78:78:78:78:02  May 01 09:40:52
10.200.100.5

2001:db8:200:100::5

fe80::205:86ff:fe71:9b13
     100        2c:6b:f5:2c:a7:f0  irb.100                        May 01 09:20:06
10.200.100.251

2001:db8:200:100::251

fe80::2e6b:f500:642c:a7f0
     101        00:00:5e:00:01:01  05:00:00:fc:01:00:00:00:65:00  May 01 09:20:06
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:fc:01:00:00:00:65:00  May 01 09:20:06
2001:db8:200:101::254
     101        02:05:86:71:76:13  00:78:78:78:78:78:78:78:78:01  May 01 09:40:49
10.200.101.3

2001:db8:200:101::3

fe80::205:86ff:fe71:7613
     101        2c:6b:f5:2c:a7:f0  irb.101                        May 01 09:20:06
10.200.101.251

2001:db8:200:101::251
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fe80::2e6b:f500:652c:a7f0
```

3) Ensure that all virtual gateway addresses reply to ICMP packets.

vMX7:

```
lab@vmx7# show interfaces irb
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.250/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::250/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.250/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::250/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

vMX8:

```
lab@vmx8# show interfaces irb
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.251/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::251/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.251/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::251/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

Verification:

```
lab@vQFX1# run ping 10.200.100.254 routing-instance server3
PING 10.200.100.254 (10.200.100.254): 56 data bytes
64 bytes from 10.200.100.254: icmp_seq=0 ttl=64 time=337.286 ms
..

{master:0}[edit]
lab@vQFX1# run ping 10.200.100.254 routing-instance user4
PING 10.200.100.254 (10.200.100.254): 56 data bytes
64 bytes from 10.200.100.254: icmp_seq=0 ttl=64 time=300.550 ms
..

{master:0}[edit]
lab@vQFX1# run ping 10.200.101.254 routing-instance user4
PING 10.200.101.254 (10.200.101.254): 56 data bytes
64 bytes from 10.200.101.254: icmp_seq=0 ttl=64 time=300.650 ms
..

{master:0}[edit]
lab@vQFX1# run ping 10.200.101.254 routing-instance server3
PING 10.200.101.254 (10.200.101.254): 56 data bytes
64 bytes from 10.200.101.254: icmp_seq=0 ttl=64 time=300.956 ms
```

Configurations for Task Completion:

vMX7:

```
set routing-options router-id 10.200.0.7
set protocols bgp group ibgp family evpn signaling
set interfaces ae0 aggregated-ether-options lacp system-id 00:78:78:78:78:01
set interfaces ae0 unit 0 esi 00:78:78:78:78:78:78:78:78:01
set interfaces ae0 unit 0 esi all-active
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ae1 aggregated-ether-options lacp system-id 00:78:78:78:78:02
set interfaces ae1 unit 0 esi 00:78:78:78:78:78:78:78:02
set interfaces ae1 unit 0 esi all-active
set interfaces irb unit 100 family inet address 10.200.100.250/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::250/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 family inet address 10.200.101.250/24 virtual-gateway-
address 10.200.101.254
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::250/64 virtual-
gateway-address 2001:0db8:0200:101::254
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface ae0.0
set routing-instances evpn1 interface ae1.0
set routing-instances evpn1 route-distinguisher 10.200.0.7:1
set routing-instances evpn1 vrf-target target:64513:1
set routing-instances evpn1 vrf-target auto
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list all
set routing-instances evpn1 protocols evpn default-gateway do-not-advertise
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 101 virtual-gateway-accept-data
```

vMX8:

```
set routing-options router-id 10.200.0.8
set protocols bgp group ibgp family evpn signaling
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:78:78:78:78:01
set interfaces ae0 unit 0 esi 00:78:78:78:78:78:78:78:01
set interfaces ae0 unit 0 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:78:78:78:78:02
set interfaces ae1 unit 0 esi 00:78:78:78:78:78:78:78:02
set interfaces ae1 unit 0 esi all-active
set interfaces irb unit 100 family inet address 10.200.100.251/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::251/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 family inet address 10.200.101.251/24 virtual-gateway-
address 10.200.101.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::251/64 virtual-
gateway-address 2001:0db8:0200:101::254
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface ae0.0
set routing-instances evpn1 interface ae1.0
set routing-instances evpn1 route-distinguisher 10.200.0.8:1
set routing-instances evpn1 vrf-target target:64513:1
set routing-instances evpn1 vrf-target auto
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list all
set routing-instances evpn1 protocols evpn default-gateway do-not-advertise
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 101 virtual-gateway-accept-data
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Chapter 7: Class of Service

This chapter is focused on Class of Services. You will configure and monitor MF and BA Classifiers, Two-color Policers, different Forwarding classes with respected Queues and Schedulers, DSCP Rewrite markers, RED Drop profiles and Class based forwarding.

COS Overview

In today's networking world, numerous applications and protocols are using the same packet switched network. These applications and protocols are very different in nature, but they have to share a common infrastructure.

Packet loss, delay and jitter can severely impact these applications and protocols. A set of QoS mechanisms can be used to give preferential treatment to sensitive applications and/or protocols. To be able to use the different QoS mechanisms on Juniper devices, we use the Class of Service (CoS) configuration construct.

The first CoS processing task that is performed on every router in the network is input traffic classification. By distributing packets in different virtual containers called Forwarding Classes (FC), we can define the order in which packets leave the router at a later stage. Packet classification is performed by means of a BA classifier or multi-field (MF) classifier.

A BA classifier examines the headers of incoming packets and then determines the Forwarding Class of those packets. The classification can be based on the value of any of the following fields: DSCP, IP Precedence and 802.1p. You can also define one more locally significant parameter called Packet Loss Priority (PLP). There are 4 possible values: high, medium-high, medium-low and low. PLP only comes into play during times of congestion and it can be used to determine what packets are dropped first. By default, a BA classifier is applied to every logical interface that is configured on the device. This can be changed by configuring and applying another BA classifier.

The MF classifier offers a more granular approach as you can use numerous firewall filter parameters as match criteria. The MF classifier can be used to assign packets to different Forwarding Classes, and it can be used to set PLP values. The MF classifier is commonly used on edge devices. Since they are processed after BA classifiers, the MF classifier can be used to overwrite any previous classifications.

Policing is a mechanism that can be used to limit the amount of traffic (mostly incoming) for a configured bandwidth value. Policing uses a token bucket algorithm which allows for traffic bursts as long as there are enough tokens in the bucket. Tokens are accumulated during physical interface silence and spent when the interface transmits packets at wire

© 2020 Juniper Networks, Inc. All rights reserved.

speed. Policers can be set to a hard-limit, which means that out-of-profile traffic is dropped. You can also configure soft-limit Policers. When a soft-limit Policer is used, out-of-profile traffic is assigned to a different (or worse) forwarding class and it is given a higher PLP.

Before transit traffic is sent through the fabric of the device, it can be subjected to CoS-based forwarding. This mechanism changes the way in which the outgoing interface is identified. Instead of a traditional destination IP lookup, the router uses class-to-next-hop mappings to forward traffic destined to a configured prefix. By using this approach, you can 'push' traffic across a more 'reliable' and 'faster' path.

Schedulers define how outgoing interface queues will be serviced. Every physical interface has 4 or 8 hardware queues that are mapped to forwarding classes. Every scheduler has a priority, interface bandwidth and buffer size configured. Each queue is serviced on a round-robin basis (from the highest priority to the lowest) according to its configured weight. That means that every Forwarding Class has a guaranteed rate (CIR) while unallocated bandwidth can be shared equally between other queues. You can also reference drop profiles in the JUNOS scheduler configuration, which can be used to enable the congestion avoidance mechanism called WRED.

As packets leave the interface of the device, the rewrite rule is the last mechanism applied during CoS processing. The rewrite rule can be used to set/ change different CoS field values according to the packet's forwarding class and PLP.

Classification

The BA and the multifield classifier are applied at different stages during CoS processing, and they are configured in different configuration hierarchies as well.

BA classifier: The BA classifier is configured, and applied, under the [**class-of-service classifiers**] hierarchy. In this hierarchy, you first choose what type of classifier you want to configure and then you name the classifier. Inside the classifier configuration, you simply associate forwarding classes with code-points and a loss-priority.

The following creates a classifier named IPv4, and associates a forwarding class called BE with a loss-priority of low and a code-point of '000000':

set class-of-service classifiers dscp IPV4 forwarding-class BE loss-priority low code-points 000000

You can, optionally, configure code-point-aliases under the [**class-of-service code-point-aliases**] hierarchy. This way, you can reference names instead of code-point values.

© 2020 Juniper Networks, Inc. All rights reserved.

Multifield classifier: The multifield classifier is configured using a firewall filter. Firewall filters can be configured for different address families, and the match conditions offer a lot more granularity then the BA classifier. Some match conditions include source IP, destination IP, source port, destination port and protocol. After specifying the match criteria (using **from**), the '**then forwarding-class**' can be used to actually classify the packet.

On EX, you can also easily classify VoIP VLAN traffic by configuring the [**ethernet-switching-options voip interface xx forwarding-class xxx**] command.

Loss Priority

Packet Loss Priority (**PLP**) can be used to decide what traffic is to be dropped first when congestion occurs. The PLP can be set using multifield classifiers, BA classifiers and policers. Usually, a higher loss priority is configured to be dropped first during congestion. How exactly this is done can be specified using a drop profile.

Scheduling

Schedulers configured under the [**class-of-service schedulers**] hierarchy, define how traffic inside a single queue is treated. Inside this configuration hierarchy, we can define the **transmission rate**, **queue priority**, **buffer size** and **drop profile** that is associated with the scheduler.

Scheduler maps configured under the [**class-of-service scheduler-maps**] hierarchy, can be used to apply the schedulers to a forwarding class. Scheduler-maps can be applied a logical interface directly, or they can be applied using a traffic-control-profile.

Example configuration of a scheduler called '**silver-sched**':

```
set class-of-service schedulers silver-sched transmit-rate percent 20
set class-of-service schedulers silver-sched buffer-size percent 16
set class-of-service schedulers silver-sched priority medium-high
set class-of-service schedulers silver-sched drop-profile-map loss-priority high
protocol any drop-profile high
```

A scheduler-map, called '**example**', that applies several schedulers to forwarding classes:

```
set class-of-service scheduler-maps example forwarding-class bronze scheduler bronze-
sched
set class-of-service scheduler-maps example forwarding-class silver scheduler silver-
sched
set class-of-service scheduler-maps example forwarding-class gold scheduler gold-sched
```

Drop profiles and RED

By default, Junos applies a 100% drop rate when there is a 100% fill. To start using Weighting RED (**WRED**), we have to define a drop profile. The drop profile is configured

© 2020 Juniper Networks, Inc. All rights reserved.

under the [**class-of-service drop-profiles**] hierarchy. Inside this configuration hierarchy, you can configure the drop probability for a fill-level.

By default, the drop profile is **segmented**. This means the drop profile will apply the drop probability as soon as a fill level is reached. It will do so up until the fill level rises to another configured level. When we specify the drop profile to be **interpolated**, we get a smooth graph. In this case, Junos automatically interpolates the different values that were configured.

For a drop profile to become active, it needs to be referenced in a scheduler.

The following is an example of an interpolated drop profile called '**HIGH**':

```
set class-of-service drop-profiles HIGH interpolate fill-level 50
set class-of-service drop-profiles HIGH interpolate fill-level 75
set class-of-service drop-profiles HIGH interpolate fill-level 85
set class-of-service drop-profiles HIGH interpolate drop-probability 0
set class-of-service drop-profiles HIGH interpolate drop-probability 50
set class-of-service drop-profiles HIGH interpolate drop-probability 100
```

Rewrite rules

Rewrite rules are used to mark packets on egress. Downstream nodes will be able to perform classification using only the BA classifier. The purpose of the rewrite rules is to convey a packet CoS profile to the next-hop router based on forwarding class and PLP. Rewrite rules are configured in the [**class-of-service rewrite-rules**] hierarchy and they are applied under the [**class-of-service interfaces**] hierarchy.

Policing

The Junos OS supports the following policers:

- Single-rate two-color

This policer consists of a CIR (Committed Information Rate) and a CBS (Committed Burst Size). Traffic exceeding the CIR and CBS can be marked or discarded.

- Single-rate tricolor

This policer uses the CIR, the CBS. It also introduces an additional burst size, the EBS (Excess Burst Size). The three different colors related to this policer are green, yellow and red. Green traffic is all traffic that falls within the CIR + CBS limit. Yellow exceeds the CBS but falls within the limits set by the EBS. Red is traffic that exceeds the EBS. Different packet loss priorities can be applied to each of these colors.

© 2020 Juniper Networks, Inc. All rights reserved.

- Two-rate tricolor

A two-rate tricolor policer takes the CIR and CBS from the single-rate two color policer, and adds both a Peak Information Rate (PIR) and a Peak Burst Size (PBS). Green is what conforms to the CIR + CBS. Yellow exceeds this, but it does not exceed the PIR + CBS. Red is what exceeds PIR + PBS.

A single-rate tricolor policer offers an additional burst-size over a single-rate two-color policer. A two-rate tricolor offers an additional rate over the additional burst size that a single-rate tricolor policer offers.

In the following example, we configure a single-rate two-color policer:

```
[edit firewall]
policer policer-name {
       if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then   {
policer-action (discard, set forwarding class, set PLP)
        }
}
```

We can apply the policer directly on an interface (under the family hierarchy) or within a firewall filter using the **then** statement. By default, when we reference a policer in a firewall term, the policer is 'term specific'. This means that the policer is an instance for that specific term. When can also define a 'filter-specific' policer. In this case, a single instance of a policer will police several terms.

CoS-Based Forwarding

CoS-Based Forwarding configuration requires the following configuration steps:

1) Configuring a [**class-of-service forwarding-policy next-hop map**]
2) Configuring a **routing-policy** to match interesting routes (**from protocol, route-filter** etc.) and point it to the previously configured forwarding policy (**then cos-next-hop-map**)
3) Exporting the configured routing policy to the **forwarding-table**

© 2020 Juniper Networks, Inc. All rights reserved.

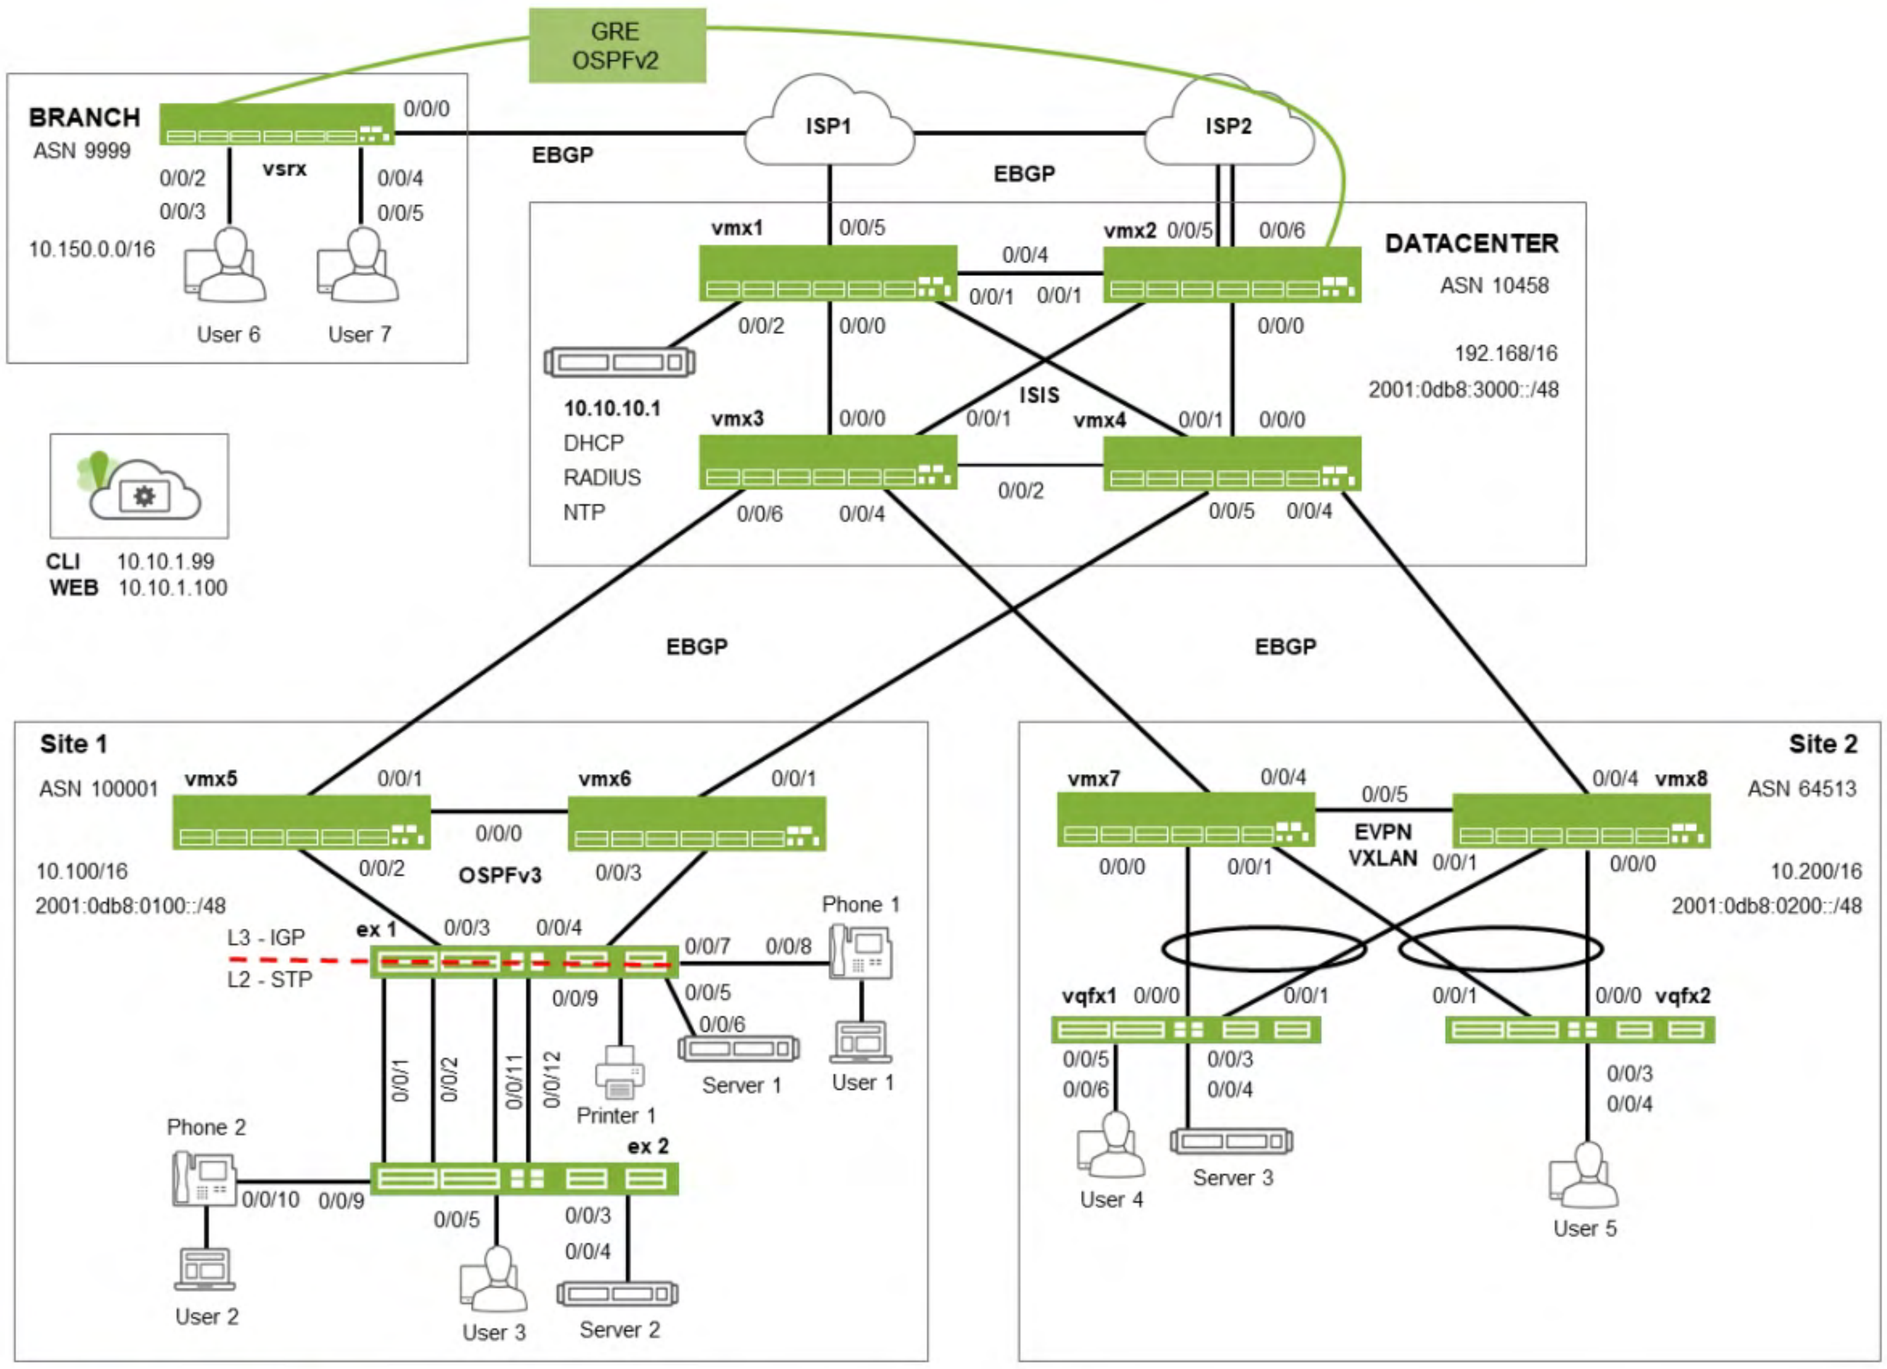

Figure 6 topology

Task 1: vMX Class of Service

In this task, you will configure all vMX devices in the data center to support Class of Service.

1) Configure Forwarding Classes and map them to the output Queues as indicated in the table below.

Forwarding Class	Queue	Scheduler
control	7	control-sched
voice	5	voice-sched
server-gold	3	server-gold-sched
server-silver	2	server-silver-sched
server-bronze	1	server-bronze-sched

© 2020 Juniper Networks, Inc. All rights reserved.

best-effort	0	best-effort-sched

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service forwarding-classes
queue 0 best-effort;
queue 1 server-bronze;
queue 2 server-silver;
queue 3 server-gold;
queue 5 voice;
queue 7 control;
```

2) Configure schedulers with the parameters in Table 7.

Scheduler	Parameter	Value
control-sched	priority	medium-high
	transmit rate	5%
	buffer-size	10%
voice-sched	priority	strict-high
	transmit rate	10%
	buffer-size	20 (milliseconds)
server-gold-sched	priority	high
	transmit rate	25%
	buffer-size	10% exact
server-silver-sched	priority	medium-high
	transmit rate	35%
	buffer-size	50% exact
	drop profile high	relaxed
server-bronze-sched	priority	medium-low
	transmit rate	remainder
	buffer-size	remainder
	drop profile low	relaxed
	drop profile high	aggressive
best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder
	drop profile any	aggressive

Table 9

vMX1, vMX2, vMX3, and vMX4:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show class-of-service schedulers
control-sched {
    transmit-rate percent 5;
    buffer-size percent 10;
    priority medium-high;
}
voice-sched {
    transmit-rate percent 10;
    buffer-size temporal 20k;
    priority strict-high;
}
server-gold-sched {
    transmit-rate percent 25;
    priority high;
}
server-silver-sched {
    transmit-rate percent 35;
    priority medium-high;
    drop-profile-map loss-priority high protocol any drop-profile relaxed;
}
server-bronze-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority medium-low;
    drop-profile-map loss-priority low protocol any drop-profile relaxed;
    drop-profile-map loss-priority high protocol any drop-profile aggressive;
}
best-effort-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile aggressive;
}
```

3) Create scheduler maps that link the schedulers with their respective forwarding class.

Forwarding Class	Scheduler
control	control-sched
voice	voice-sched
server-gold	server-gold-sched

© 2020 Juniper Networks, Inc. All rights reserved.

server-silver	server-silver-sched
server-bronze	server-bronze-sched
best-effort	best-effort-sched

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service scheduler-maps
cos-sched-map {
    forwarding-class control scheduler control-sched;
    forwarding-class voice scheduler voice-sched;
    forwarding-class server-gold scheduler server-gold-sched;
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
```

4) Configure two interpolated Drop profiles listed in Table 8.

Drop profile	Fill level (%)	Drop probability (%)
relaxed	75	0
	90	75
	100	100
aggressive	50	0
	75	50
	90	75
	100	100

Table 10

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service drop-profiles
relaxed {
    interpolate {
        fill-level [ 75 90 100 ];
        drop-probability [ 0 75 100 ];
    }
}
aggressive {
    interpolate {
        fill-level [ 50 75 90 100 ];
        drop-probability [ 0 50 75 100 ];
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

5) Apply the schedulers to all physical gigabit Ethernet interfaces.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
}
```

6) Configure IPv4 and IPv6 Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service classifiers
dscp ba-classifier {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
dscp-ipv6 ba-classifier-v6 {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
```

7) Apply the classifiers to all gigabit Ethernet logical interfaces.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
```

8) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	DSCP value	Loss-priority
control	110000	low

© 2020 Juniper Networks, Inc. All rights reserved.

	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service rewrite-rules
dscp rewrite-v4 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
dscp-ipv6 rewrite-v6 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
```

9) Apply the rewrite rules on all logical interfaces.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
```

Configurations for Task Completion:

vMX1, vMX2, vMX3, and vMX4:

```
set class-of-service forwarding-classes queue 0 best-effort
set class-of-service forwarding-classes queue 1 server-bronze
set class-of-service forwarding-classes queue 2 server-silver
set class-of-service forwarding-classes queue 3 server-gold
set class-of-service forwarding-classes queue 5 voice
set class-of-service forwarding-classes queue 7 control
set class-of-service schedulers control-sched transmit-rate percent 5
set class-of-service schedulers control-sched buffer-size percent 10
set class-of-service schedulers control-sched priority medium-high
set class-of-service schedulers voice-sched transmit-rate percent 10
set class-of-service schedulers voice-sched buffer-size temporal 20k
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers server-gold-sched transmit-rate percent 25
set class-of-service schedulers server-gold-sched priority high
set class-of-service schedulers server-silver-sched transmit-rate percent 35
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service schedulers server-silver-sched priority medium-high
set class-of-service schedulers server-silver-sched drop-profile-map loss-priority high
protocol any drop-profile relaxed
set class-of-service schedulers server-bronze-sched transmit-rate remainder
set class-of-service schedulers server-bronze-sched buffer-size remainder
set class-of-service schedulers server-bronze-sched priority medium-low
set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority low
protocol any drop-profile relaxed
set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority high
protocol any drop-profile aggressive
set class-of-service schedulers best-effort-sched transmit-rate remainder
set class-of-service schedulers best-effort-sched buffer-size remainder
set class-of-service schedulers best-effort-sched priority low
set class-of-service schedulers best-effort-sched drop-profile-map loss-priority any
protocol any drop-profile aggressive
set class-of-service scheduler-maps cos-sched-map forwarding-class control scheduler
control-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class voice scheduler
voice-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-gold scheduler
server-gold-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class best-effort scheduler
best-effort-sched
set class-of-service drop-profiles relaxed interpolate fill-level [ 75 90 100 ]
set class-of-service drop-profiles relaxed interpolate drop-probability [ 0 75 100 ]
set class-of-service drop-profiles aggressive interpolate fill-level [ 50 75 90 100 ]
set class-of-service drop-profiles aggressive interpolate drop-probability [ 0 50 75 100
]
set class-of-service interfaces ge-* scheduler-map cos-sched-map
set class-of-service classifiers dscp ba-classifier import default
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc1
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc2
set class-of-service classifiers dscp ba-classifier forwarding-class voice loss-priority
low code-points ef
set class-of-service classifiers dscp ba-classifier forwarding-class server-gold loss-
priority low code-points af41
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority low code-points af21
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority high code-points af23
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority low code-points af11
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority high code-points af13
set class-of-service classifiers dscp ba-classifier forwarding-class best-effort loss-
priority low code-points be
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 import default
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc1
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc2
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class voice loss-
priority low code-points ef
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-gold
loss-priority low code-points af41
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority low code-points af21
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority high code-points af23
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority low code-points af11
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority high code-points af13
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class best-effort
loss-priority low code-points be
set class-of-service interfaces ge-* unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-* unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class voice loss-priority
low code-point ef
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-gold loss-
priority low code-point af41
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority low code-point af21
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority high code-point af23
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority low code-point af11
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority high code-point af13
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class best-effort loss-
priority high code-point be
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class voice loss-
priority low code-point ef
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-gold
loss-priority low code-point af41
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority low code-point af21
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority high code-point af23
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority low code-point af11
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority high code-point af13
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class best-effort
loss-priority high code-point be
set class-of-service interfaces ge-* unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-* unit * rewrite-rules dscp-ipv6 rewrite-v6
```

Task 2: EX Class of Service

In this task, you will configure EX1 and EX2 in Site1 to support Class of Service.

1) Configure the following Forwarding Classes and map them to the output Queues as indicated in the table below.

Forwarding Class	Queue	Schedulers L2 and L3 core Interfaces (EX1)	Schedulers L2 access ports (EX2)
control	7	control-sched-core	control-sched-user
voice	5	voice-sched-core	voice-sched-user
server-gold	3	server-gold-sched	server-gold-sched
server-silver	2	server-silver-sched	server-silver-sched
server-bronze	1	server-bronze-sched	server-bronze-sched
best-effort	0	best-effort-sched	best-effort-sched

EX1 and EX2:

```
lab@ex1# show class-of-service forwarding-classes
queue 0 best-effort;
queue 1 server-bronze;
queue 2 server-silver;
queue 3 server-gold;
queue 5 voice;
queue 7 control;
```

2) Configure schedulers with the parameters as described in th table below.

© 2020 Juniper Networks, Inc. All rights reserved.

Scheduler	Parameter	Value
control-sched-core	priority	strict-high
	shaping rate	5%
	buffer-size	10%
control-sched-user	priority	strict-high
	shaping rate	2%
	buffer-size	10%
voice-sched-core	priority	strict-high
	shaping rate	10%
	buffer-size	5% exact
voice-sched-user	priority	strict-high
	shaping rate	5%
	buffer-size	5% exact
server-gold-sched	priority	strict-high
	shaping rate	25%
	buffer-size	10% exact
server-silver-sched	priority	low
	transmit rate	35%
	buffer-size	50%
server-bronze-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder
best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder

EX1 and EX2:

```
lab@ex1# show class-of-service schedulers
control-sched-core {
    shaping-rate percent 5;
    buffer-size percent 10;
    priority strict-high;
}
control-sched-user {
    shaping-rate percent 2;
    buffer-size percent 10;
    priority strict-high;
}
voice-sched-core {
    shaping-rate percent 10;
    buffer-size {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            percent 5;
            exact;
        }
        priority strict-high;
    }
    voice-sched-user {
        shaping-rate percent 5;
        buffer-size {
            percent 5;
            exact;
        }
        priority strict-high;
    }
    server-gold-sched {
        shaping-rate percent 25;
        buffer-size {
            percent 10;
            exact;
        }
        priority strict-high;
    }
    server-silver-sched {
        transmit-rate percent 35;
        buffer-size percent 50;
        priority low;
    }
    server-bronze-sched {
        transmit-rate {
            remainder;
        }
        buffer-size {
            remainder;
        }
        priority low;
    }
    best-effort-sched {
        transmit-rate {
            remainder;
        }
        buffer-size {
            remainder;
        }
        priority low;
    }
```

3) Create scheduler maps that link the schedulers with their respective forwarding class.

© 2020 Juniper Networks, Inc. All rights reserved.

Forwarding Class	Schedulers L2 access ports	Forwarding Class
control	control-sched-core	control-sched-user
voice	voice-sched-core	voice-sched-user
server-gold	server-gold-sched	server-gold-sched
server-silver	server-silver-sched	server-silver-sched
server-bronze	server-bronze-sched	server-bronze-sched
best-effort	best-effort-sched	best-effort-sched

EX1 and EX2:

```
lab@ex1# show class-of-service scheduler-maps
cos-sched-map-core {
    forwarding-class control scheduler control-sched-core;
    forwarding-class voice scheduler voice-sched-core;
    forwarding-class server-gold scheduler server-gold-sched;
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
cos-sched-map-user {
    forwarding-class control scheduler control-sched-user;
    forwarding-class voice scheduler voice-sched-user;
    forwarding-class server-gold scheduler server-gold-sched;
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
```

4) Apply the schedulers to the physical interfaces on EX1 and EX2. Make sure that OoB interfaces are not affected.

EX1:

```
lab@ex1# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/2 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/4 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    scheduler-map cos-sched-map-core;
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
}
ge-0/0/7 {
    scheduler-map cos-sched-map-user;
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
}
ae10 {
    scheduler-map cos-sched-map-core;
}
```

EX2:

```
lab@ex2# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/2 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/3 {
    scheduler-map cos-sched-map-user;
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
}
ae10 {
    scheduler-map cos-sched-map-core;
}
```

5) Configure IPv4, IPv6 and dot1p Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

EX1 should have IPv4, IPv6 and dot1q classifiers. EX2 will only need the dot1q BA since it is strictly Layer 2.

Forwarding class	Dot1p	Dot1p loss priority	DSCP value	Loss-priority
control	111	low	110000	low
			111000	High

© 2020 Juniper Networks, Inc. All rights reserved.

voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 37

EX1:

```
lab@ex1# show class-of-service classifiers
dscp ba-classifier {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
dscp-ipv6 ba-classifier-v6 {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
ieee-802.1 cos {
    forwarding-class best-effort {
        loss-priority low code-points 000;
    }
    forwarding-class server-bronze {
        loss-priority low code-points 001;
    }
    forwarding-class server-silver {
        loss-priority low code-points 010;
    }
    forwarding-class server-gold {
        loss-priority low code-points 011;
    }
    forwarding-class voice {
        loss-priority low code-points 101;
    }
    forwarding-class control {
        loss-priority low code-points 111;
    }
}
```

EX2:

```
lab@ex2# show class-of-service classifiers
ieee-802.1 cos {
    forwarding-class best-effort {
        loss-priority low code-points 000;
    }
    forwarding-class server-bronze {
        loss-priority low code-points 001;
    }
    forwarding-class server-silver {
        loss-priority low code-points 010;
    }
    forwarding-class server-gold {
        loss-priority low code-points 011;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    forwarding-class voice {
        loss-priority low code-points 101;
    }
    forwarding-class control {
        loss-priority low code-points 111;
    }
}
```

6) Apply the classifiers to all interfaces as described in the table below. Make sure that the OoB interfaces are not affected.

Device	Classifiers	Interfaces
EX1	BA IPv4, IPv6, and dot1p classifiers	All interfaces
EX2	Dot1p	All interfaces

EX1:

```
lab@ex1# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/2 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-0/0/4 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/7 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ae10 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
```

EX2:

```
lab@ex2# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/2 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/3 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ae10 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
```

7) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	Dot1p	Dot1p loss priority	DSCP value	Loss-priority
control	111	low	110000	low
			111000	high
voice	101		101110	low

© 2020 Juniper Networks, Inc. All rights reserved.

server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

EX1:

```
lab@ex1# show class-of-service rewrite-rules
dscp rewrite-v4 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
dscp-ipv6 rewrite-v6 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
ieee-802.1 cos {
    forwarding-class best-effort {
        loss-priority low code-point 000;
    }
    forwarding-class server-bronze {
        loss-priority low code-point 001;
    }
    forwarding-class server-silver {
        loss-priority low code-point 010;
    }
    forwarding-class server-gold {
        loss-priority low code-point 011;
    }
    forwarding-class voice {
        loss-priority low code-point 101;
    }
    forwarding-class control {
        loss-priority low code-point 111;
    }
}
```

EX2:

```
lab@ex2# show class-of-service rewrite-rules
ieee-802.1 cos {
    forwarding-class server-bronze {
        loss-priority low code-point 001;
    }
    forwarding-class server-silver {
        loss-priority low code-point 010;
    }
    forwarding-class server-gold {
        loss-priority low code-point 011;
    }
    forwarding-class voice {
        loss-priority low code-point 101;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

8) Apply the rewrite-rules to all interfaces as described in the table below. Make sure that the OoB interfaces are not affected.

Device	Classifiers	Interfaces
EX1	IPv4, IPv6 and dot1p rewrite-rules	All interfaces
EX2	Dot1p rewrite rules	All interfaces

Table 38

EX1:

```
lab@ex1# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/2 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-0/0/4 {
    scheduler-map cos-sched-map-core;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/7 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ae10 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
```

EX2:

```
lab@ex2# show class-of-service interfaces
ge-0/0/1 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/2 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/3 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/5 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/9 {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
ae10 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
```

9) Configure all voice enabled interfaces on EX1 and EX2 to advertise the Forwarding-Class ‘voice’ to the VoIP phones.

EX1:

```
lab@ex1# show switch-options
voip {
    interface ge-0/0/7.0 {
        vlan phone;
        forwarding-class voice;
    }
}
```

EX2:

```
lab@ex2# show switch-options
voip {
    interface ge-0/0/9.0 {
        vlan phone;
        forwarding-class voice;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
interface ge-0/0/9.0 {
    interface-mac-limit {
        10;
        packet-action drop;
    }
}
```

Configurations for Task Completion:

EX1:

```
set class-of-service forwarding-classes queue 0 best-effort
set class-of-service forwarding-classes queue 1 server-bronze
set class-of-service forwarding-classes queue 2 server-silver
set class-of-service forwarding-classes queue 3 server-gold
set class-of-service forwarding-classes queue 5 voice
set class-of-service forwarding-classes queue 7 control
set class-of-service schedulers control-sched-core shaping-rate percent 5
set class-of-service schedulers control-sched-core buffer-size percent 10
set class-of-service schedulers control-sched-core priority strict-high
set class-of-service schedulers control-sched-user shaping-rate percent 2
set class-of-service schedulers control-sched-user buffer-size percent 10
set class-of-service schedulers control-sched-user priority strict-high
set class-of-service schedulers voice-sched-core shaping-rate percent 10
set class-of-service schedulers voice-sched-core buffer-size percent 5
set class-of-service schedulers voice-sched-core buffer-size exact
set class-of-service schedulers voice-sched-core priority strict-high
set class-of-service schedulers voice-sched-user shaping-rate percent 5
set class-of-service schedulers voice-sched-user buffer-size percent 5
set class-of-service schedulers voice-sched-user buffer-size exact
set class-of-service schedulers voice-sched-user priority strict-high
set class-of-service schedulers server-gold-sched shaping-rate percent 25
set class-of-service schedulers server-gold-sched buffer-size percent 10
set class-of-service schedulers server-gold-sched buffer-size exact
set class-of-service schedulers server-gold-sched priority strict-high
set class-of-service schedulers server-silver-sched transmit-rate percent 35
set class-of-service schedulers server-silver-sched buffer-size percent 50
set class-of-service schedulers server-silver-sched priority low
set class-of-service schedulers server-bronze-sched transmit-rate remainder
set class-of-service schedulers server-bronze-sched buffer-size remainder
set class-of-service schedulers server-bronze-sched priority low
set class-of-service schedulers best-effort-sched transmit-rate remainder
set class-of-service schedulers best-effort-sched buffer-size remainder
set class-of-service schedulers best-effort-sched priority low
set class-of-service scheduler-maps cos-sched-map-core forwarding-class control
scheduler control-sched-core
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service scheduler-maps cos-sched-map-core forwarding-class voice scheduler
voice-sched-core
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class best-effort
scheduler best-effort-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class control
scheduler control-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class voice scheduler
voice-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class best-effort
scheduler best-effort-sched
set class-of-service interfaces ge-0/0/1 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/2 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/3 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/4 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/5 scheduler-map cos-sched-map-user
set class-of-service interfaces ge-0/0/7 scheduler-map cos-sched-map-user
set class-of-service interfaces ge-0/0/9 scheduler-map cos-sched-map-user
set class-of-service interfaces ae10 scheduler-map cos-sched-map-core
set class-of-service classifiers dscp ba-classifier import default
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc1
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc2
set class-of-service classifiers dscp ba-classifier forwarding-class voice loss-priority
low code-points ef
set class-of-service classifiers dscp ba-classifier forwarding-class server-gold loss-
priority low code-points af41
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority low code-points af21
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority high code-points af23
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority low code-points af11
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority high code-points af13
set class-of-service classifiers dscp ba-classifier forwarding-class best-effort loss-
priority low code-points be
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 import default
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc1
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc2
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class voice loss-
priority low code-points ef
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-gold
loss-priority low code-points af41
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority low code-points af21
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority high code-points af23
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority low code-points af11
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority high code-points af13
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class best-effort
loss-priority low code-points be
set class-of-service classifiers ieee-802.1 cos forwarding-class best-effort loss-
priority low code-points 000
set class-of-service classifiers ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-points 001
set class-of-service classifiers ieee-802.1 cos forwarding-class server-silver loss-
priority low code-points 010
set class-of-service classifiers ieee-802.1 cos forwarding-class server-gold loss-
priority low code-points 011
set class-of-service classifiers ieee-802.1 cos forwarding-class voice loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 cos forwarding-class control loss-priority
low code-points 111
set class-of-service interfaces ge-0/0/1 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/2 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-0/0/3 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-0/0/4 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-0/0/4 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-0/0/5 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/7 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/9 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ae10 unit * classifiers ieee-802.1 cos
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class voice loss-priority
low code-point ef
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-gold loss-
priority low code-point af41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority low code-point af21
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority high code-point af23
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority low code-point af11
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority high code-point af13
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class best-effort loss-
priority high code-point be
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class voice loss-
priority low code-point ef
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-gold
loss-priority low code-point af41
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority low code-point af21
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority high code-point af23
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority low code-point af11
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority high code-point af13
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class best-effort
loss-priority high code-point be
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class best-effort loss-
priority low code-point 000
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-point 001
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-silver loss-
priority low code-point 010
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-gold loss-
priority low code-point 011
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class voice loss-priority
low code-point 101
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class control loss-priority
low code-point 111
set class-of-service interfaces ge-0/0/1 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/2 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/3 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-0/0/5 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/7 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/9 unit * rewrite-rules ieee-802.1 cos
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service interfaces ae10 unit * rewrite-rules ieee-802.1 cos
set switch-options voip interface ge-0/0/7.0 forwarding-class voice
```

EX2:

```
set class-of-service forwarding-classes queue 0 best-effort
set class-of-service forwarding-classes queue 1 server-bronze
set class-of-service forwarding-classes queue 2 server-silver
set class-of-service forwarding-classes queue 3 server-gold
set class-of-service forwarding-classes queue 5 voice
set class-of-service forwarding-classes queue 7 control
set class-of-service schedulers control-sched-core shaping-rate percent 5
set class-of-service schedulers control-sched-core buffer-size percent 10
set class-of-service schedulers control-sched-core priority strict-high
set class-of-service schedulers control-sched-user shaping-rate percent 2
set class-of-service schedulers control-sched-user buffer-size percent 10
set class-of-service schedulers control-sched-user priority strict-high
set class-of-service schedulers voice-sched-core shaping-rate percent 10
set class-of-service schedulers voice-sched-core buffer-size percent 5
set class-of-service schedulers voice-sched-core buffer-size exact
set class-of-service schedulers voice-sched-core priority strict-high
set class-of-service schedulers voice-sched-user shaping-rate percent 5
set class-of-service schedulers voice-sched-user buffer-size percent 5
set class-of-service schedulers voice-sched-user buffer-size exact
set class-of-service schedulers voice-sched-user priority strict-high
set class-of-service schedulers server-gold-sched shaping-rate percent 25
set class-of-service schedulers server-gold-sched buffer-size percent 10
set class-of-service schedulers server-gold-sched buffer-size exact
set class-of-service schedulers server-gold-sched priority strict-high
set class-of-service schedulers server-silver-sched transmit-rate percent 35
set class-of-service schedulers server-silver-sched buffer-size percent 50
set class-of-service schedulers server-silver-sched priority low
set class-of-service schedulers server-bronze-sched transmit-rate remainder
set class-of-service schedulers server-bronze-sched buffer-size remainder
set class-of-service schedulers server-bronze-sched priority low
set class-of-service schedulers best-effort-sched transmit-rate remainder
set class-of-service schedulers best-effort-sched buffer-size remainder
set class-of-service schedulers best-effort-sched priority low
set class-of-service scheduler-maps cos-sched-map-core forwarding-class control
scheduler control-sched-core
set class-of-service scheduler-maps cos-sched-map-core forwarding-class voice scheduler
voice-sched-core
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-bronze
scheduler server-bronze-sched
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service scheduler-maps cos-sched-map-core forwarding-class best-effort
scheduler best-effort-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class control
scheduler control-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class voice scheduler
voice-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class best-effort
scheduler best-effort-sched
set class-of-service interfaces ge-0/0/1 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/2 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/3 scheduler-map cos-sched-map-user
set class-of-service interfaces ge-0/0/5 scheduler-map cos-sched-map-user
set class-of-service interfaces ge-0/0/9 scheduler-map cos-sched-map-user
set class-of-service interfaces ae10 scheduler-map cos-sched-map-core
set class-of-service classifiers ieee-802.1 cos forwarding-class best-effort loss-
priority low code-points 000
set class-of-service classifiers ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-points 001
set class-of-service classifiers ieee-802.1 cos forwarding-class server-silver loss-
priority low code-points 010
set class-of-service classifiers ieee-802.1 cos forwarding-class server-gold loss-
priority low code-points 011
set class-of-service classifiers ieee-802.1 cos forwarding-class voice loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 cos forwarding-class control loss-priority
low code-points 111
set class-of-service interfaces ge-0/0/1 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/2 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/5 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/9 unit * classifiers ieee-802.1 cos
set class-of-service interfaces ae10 unit * classifiers ieee-802.1 cos
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-point 001
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-silver loss-
priority low code-point 010
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-gold loss-
priority low code-point 011
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class voice loss-priority
low code-point 101
set class-of-service interfaces ge-0/0/1 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/2 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * rewrite-rules ieee-802.1 cos
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service interfaces ge-0/0/5 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/9 unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ae10 unit * rewrite-rules ieee-802.1 cos
set switch-options voip interface ge-0/0/9.0 forwarding-class voice
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Super Lab 1

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

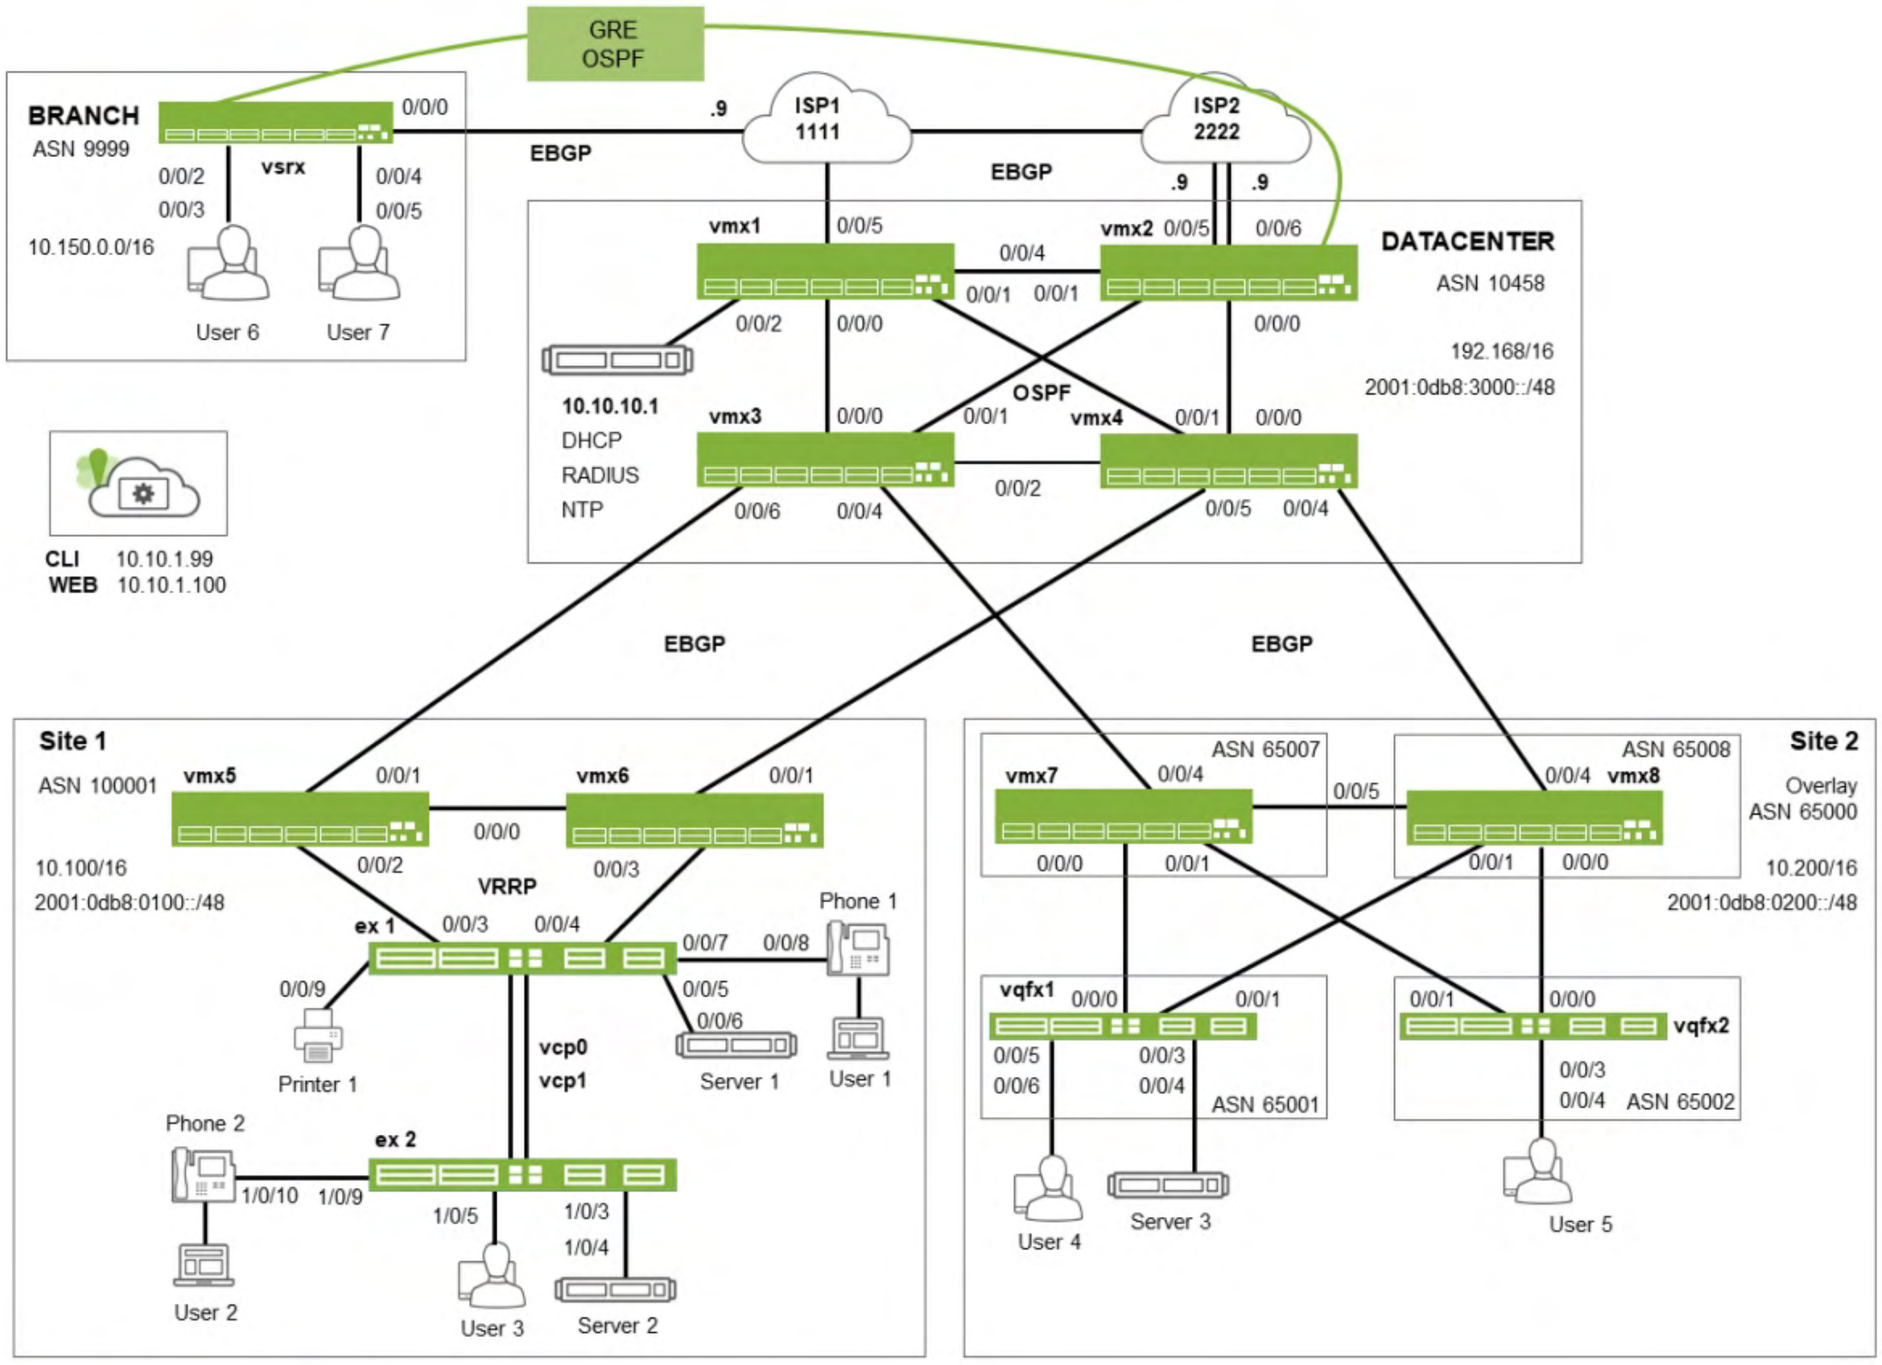

Task 1: General System Features

System Services

Note: For this section you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure all vMX and EX devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute.

Note: You will not need to configure EX2 in this section since it will be placed into a Virtual Chassis with EX1.

Configure the following on all vMX and EX devices:

```
[edit]
lab@vmx1# show system services
ssh {
    root-login allow;
    connection-limit 4;
    rate-limit 10;
}
telnet {
    connection-limit 4;
    rate-limit 10;
}
```

2) Configure server 10.10.10.1 as the DNS server.

Configure the following on all vMX and EX devices:

```
[edit]
lab@vmx1# show system name-server
10.10.10.1;
```

3) Set the time zone to *America/Los_Angeles* on all vMX devices and EX1..

Configure the following on all vMX devices, EX1 and vSRX

```
[edit]
lab@vmx1# show system time-zone
time-zone America/Los_Angeles;
```

4) Make sure that all vMX devices and EX1 synchronize the time with NTP server 10.10.10.1. You should use the local devices loopback address as the source.

NOTE: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1*

vMX1:

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show system ntp
server 10.10.10.1;
source-address 192.168.0.1;
```

vMX2:

```
[edit]
lab@vmx2# show system ntp
server 10.10.10.1;
source-address 192.168.0.2;
```

vMX3:

```
[edit]
lab@vmx3# show system ntp
server 10.10.10.1;
source-address 192.168.0.3;
```

vMX4:

```
[edit]
lab@vmx4# show system ntp
server 10.10.10.1;
source-address 192.168.0.4;
```

vMX5:

```
[edit]
lab@vmx5# show system ntp
server 10.10.10.1;
source-address 10.100.0.5;
```

vMX6:

```
[edit]
lab@vmx6# show system ntp
server 10.10.10.1;
source-address 10.100.0.6;
```

EX1:

```
[edit]
lab@ex1# show system ntp
server 10.10.10.1;
source-address 10.100.0.1;
```

vMX7:

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx7# show system ntp
server 10.10.10.1;
source-address 10.200.0.7;
```

vMX8:

```
[edit]
lab@vmx8# show system ntp
server 10.10.10.1;
source-address 10.200.0.8;
```

User Authentication and Authorization

In this section you will configure new users allowed to access the devices and define their privileges and permissions.

The following tasks should be configured on all vMX devices and EX1

5) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts 3 and timeout of 3 seconds. Use juniper123 as the RADIUS shared secret. Use the local loopback interface address as the source for authentication requests.

vMX1:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.1;
    }
}
```

vMX2:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        timeout 3;
        retry 3;
        source-address 192.168.0.2;
    }
}
```

vMX3:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.3;
    }
}
```

vMX4:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.4;
    }
}
```

vMX5:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.5;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX6:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.6;
    }
}
```

vMX7:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.200.0.7;
    }
}
```

vMX8:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.200.0.8;
    }
}
```

EX1:

```
[edit]
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.1;
    }
}
```

6) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices and EX1 as a super-user.

Configure the following on all vMXs and EX1:

```
lab@vmx1# show system login
user admin {
    uid 2003;
    class super-user;
}
...
```

SNMP Configuration

In this section you will configure SNMPv2. The following tasks should be configured on all vMX, vQFX, and vSRX devices in the topology.

7) The SNMP read-only community string is *jncie*. Make sure that only the Junos Space server can poll the devices in the data center, Site1, Site2, and the branch device.
8) Configure all devices to send SNMP traps to Junos Space, in case of link, chassis, or both events. Make sure that traps are only send using the OoB management interface.

vMX1:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.1;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

vMX2:

```
lab@vmx2# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.2;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX3:

```
lab@vmx3# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.3;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX4:

```
lab@vmx4# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.4;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX5:

```
lab@vmx5# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.5;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX6:

```
lab@vmx6# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.6;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX7:

```
lab@vmx7# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    client-list-name list;
}
trap-options {
    source-address 10.10.1.7;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX8:

```
lab@vmx8# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.8;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vSRX:

```
lab@vsrx# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.10;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vQFX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vqfx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.13;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vQFX2:

```
lab@vqfx2# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.14;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

Junos Space

In this section you will configure the Junos Space Network Director

9) You have a Junos Space Network Director instance running for device management. You must discover all vMX devices using Junos Space Network Director with SNMPv2 and the community string jncie. Ensure that the devices are present under the Inventory window.

Note: Junos Space can be accessed by browsing to the management address using HTTPS.

10) Back up all vMX configurations onto Junos Space Network Director.

© 2020 Juniper Networks, Inc. All rights reserved.

Login to Junos Space with username: *super* and password: *lab123*

In the left column navigate to “Devices / Device Discovery” and select “Device Discovery Profiles”

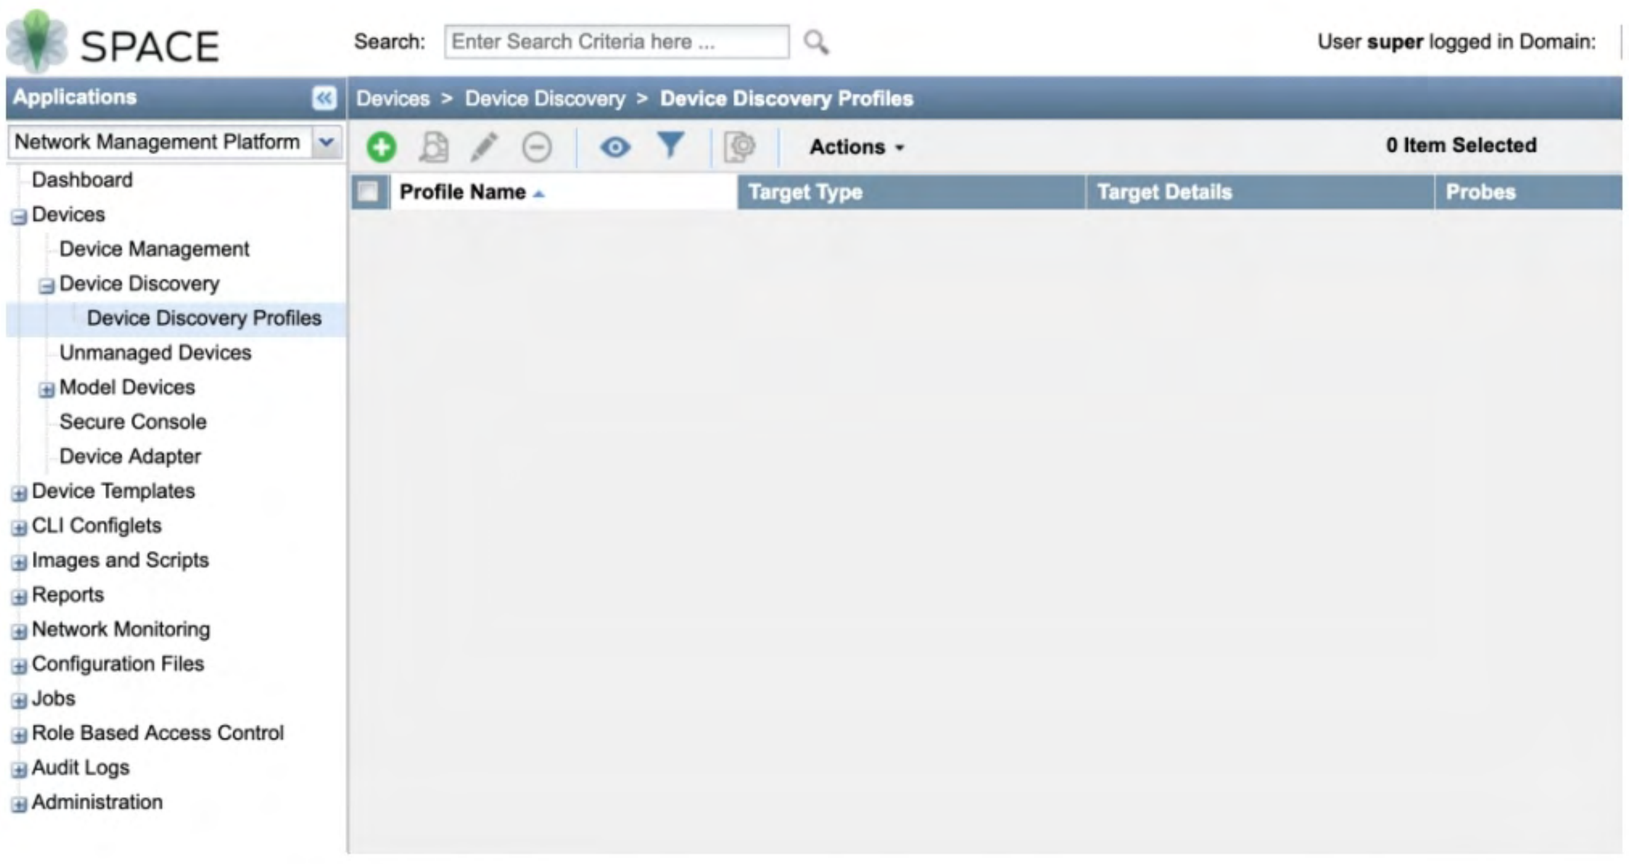

© 2020 Juniper Networks, Inc. All rights reserved.

Specify a Device Profile Name. In this case we used the name “mydevices”. The task requirement is that all vMXs in the topology should be discovered. The IPv4 management/ OOB range of the vMXs is 10.10.1.1 to 10.1.1.8. Select “Next”

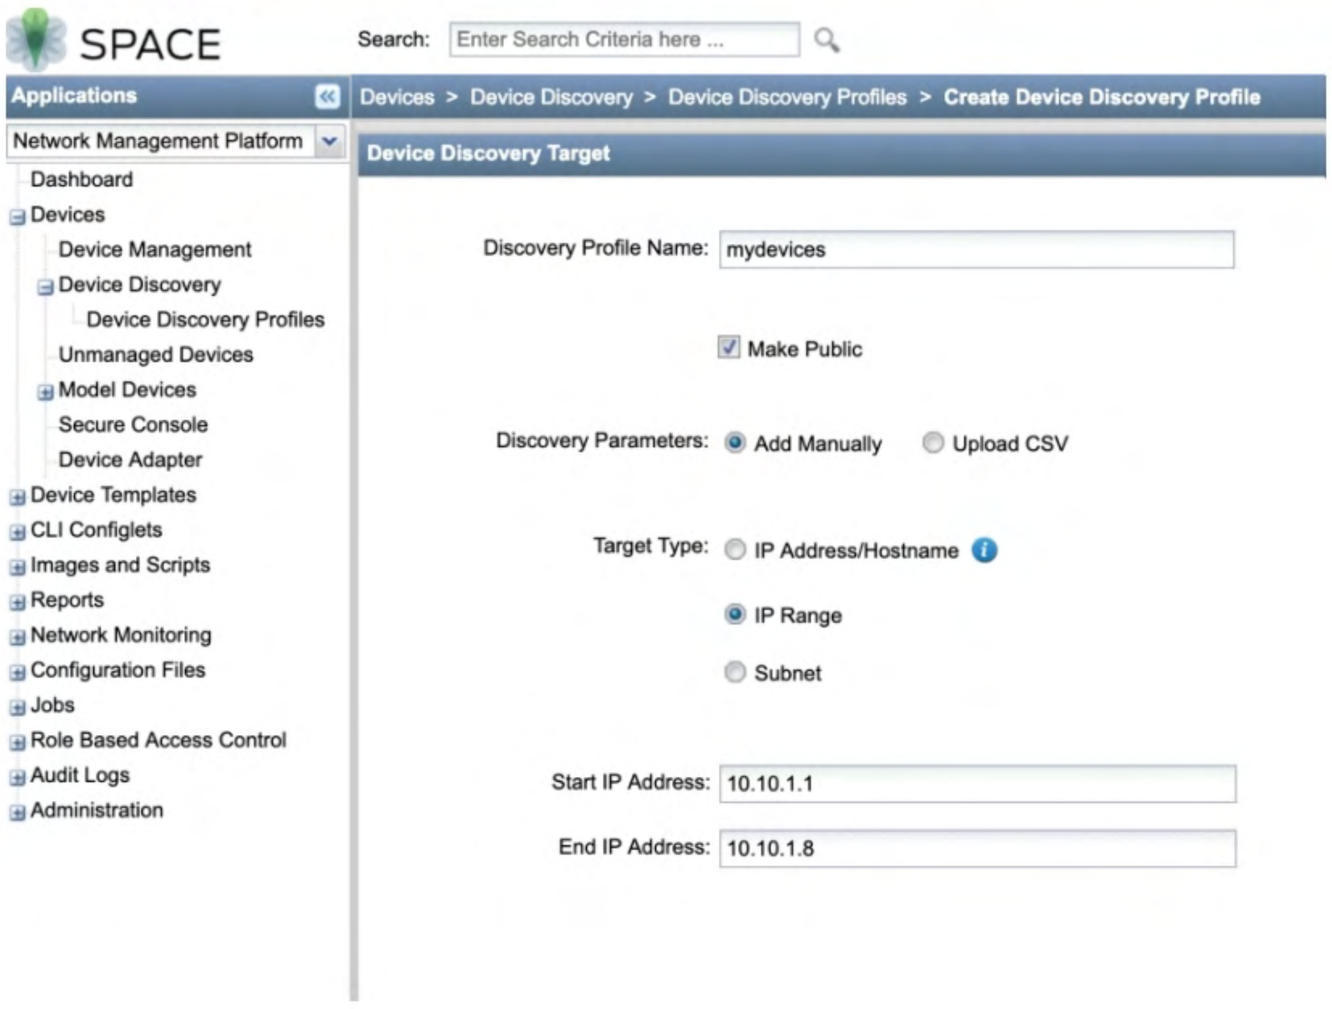

© 2020 Juniper Networks, Inc. All rights reserved.

Select and configure the probing mechanisms so that the vMX devices can be discovered. Make sure the specify the correct community name ‘jncie’ as configured on the vMX devices. Select “Next”

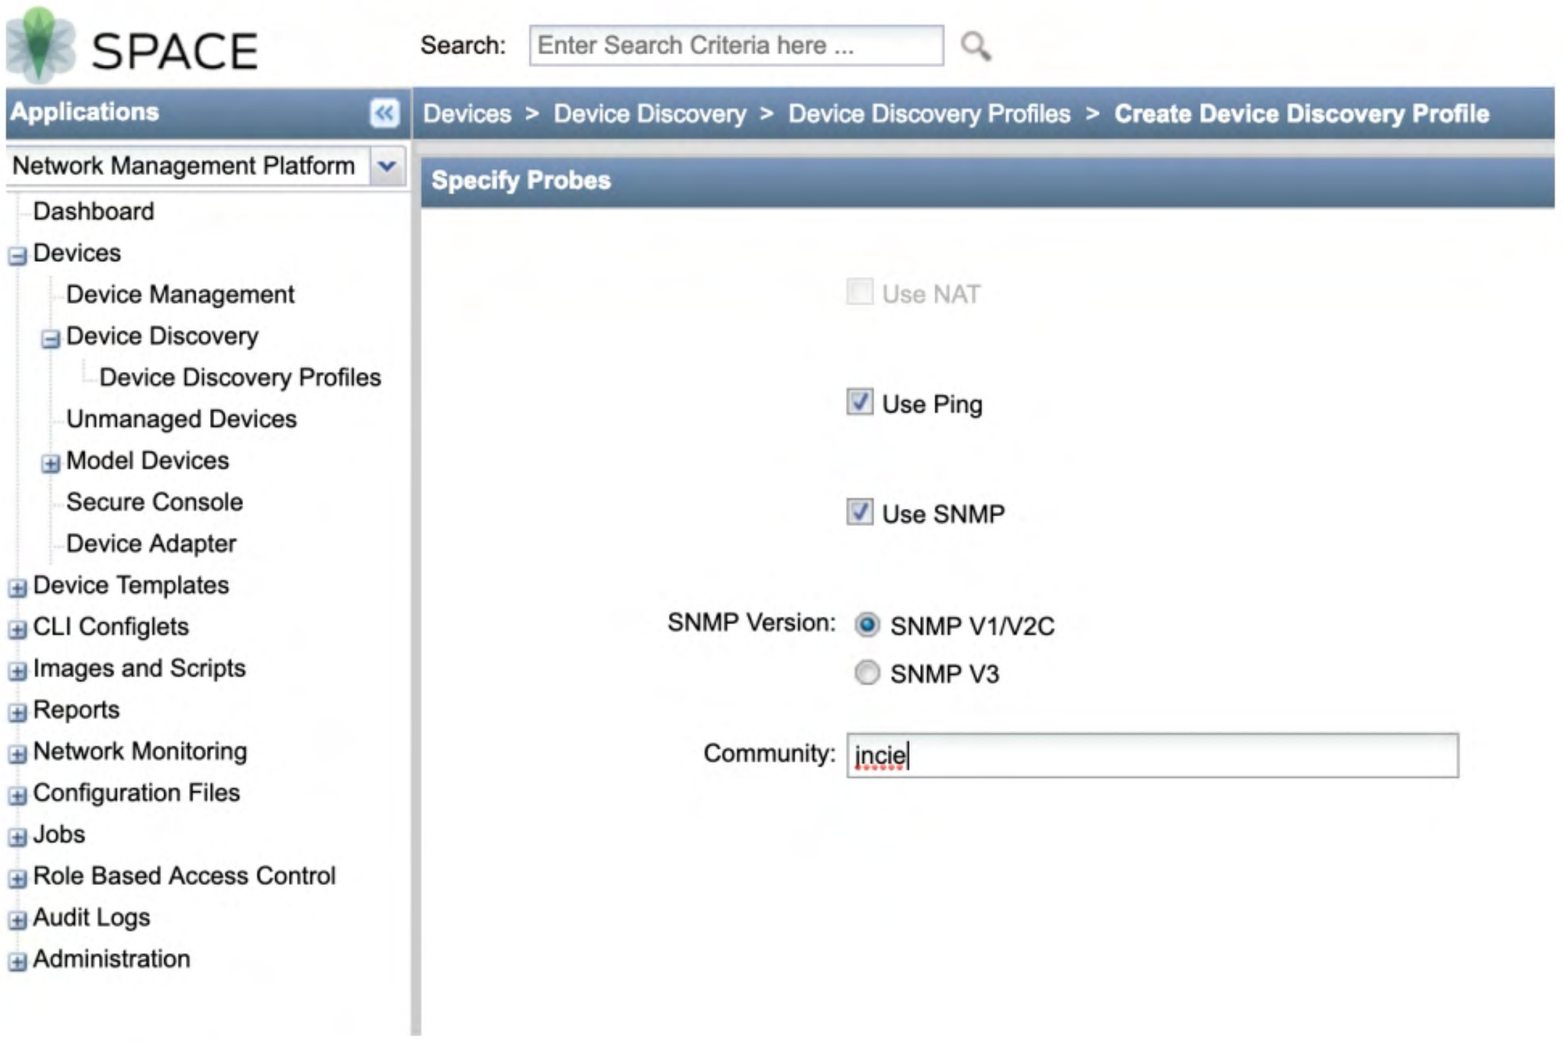

© 2020 Juniper Networks, Inc. All rights reserved.

Specify a user with super-user authorization for the vMX devices.

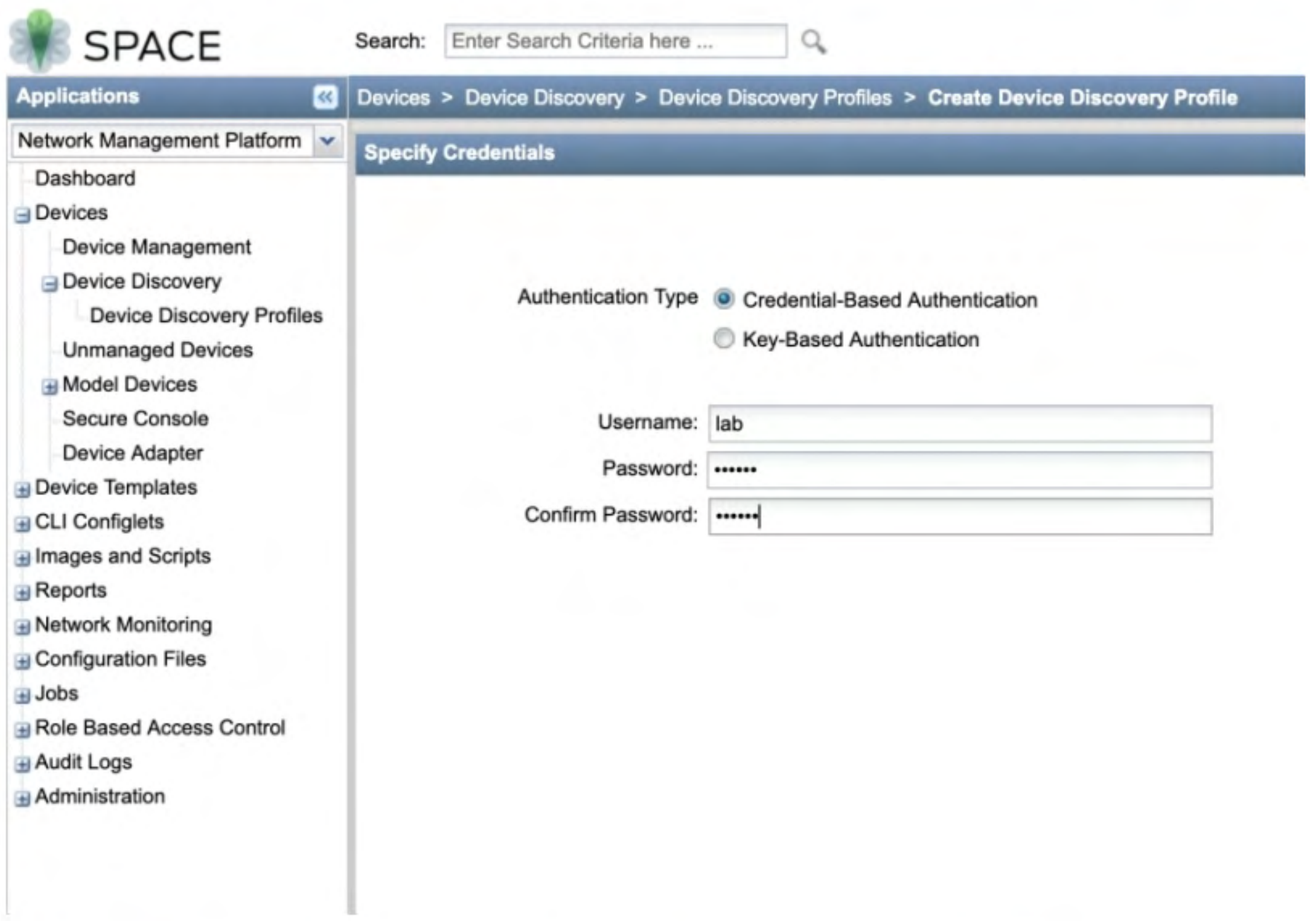

It is not needed to specify a ‘fingerprint’ as it is optional and not required by the task. Select ‘Next’.

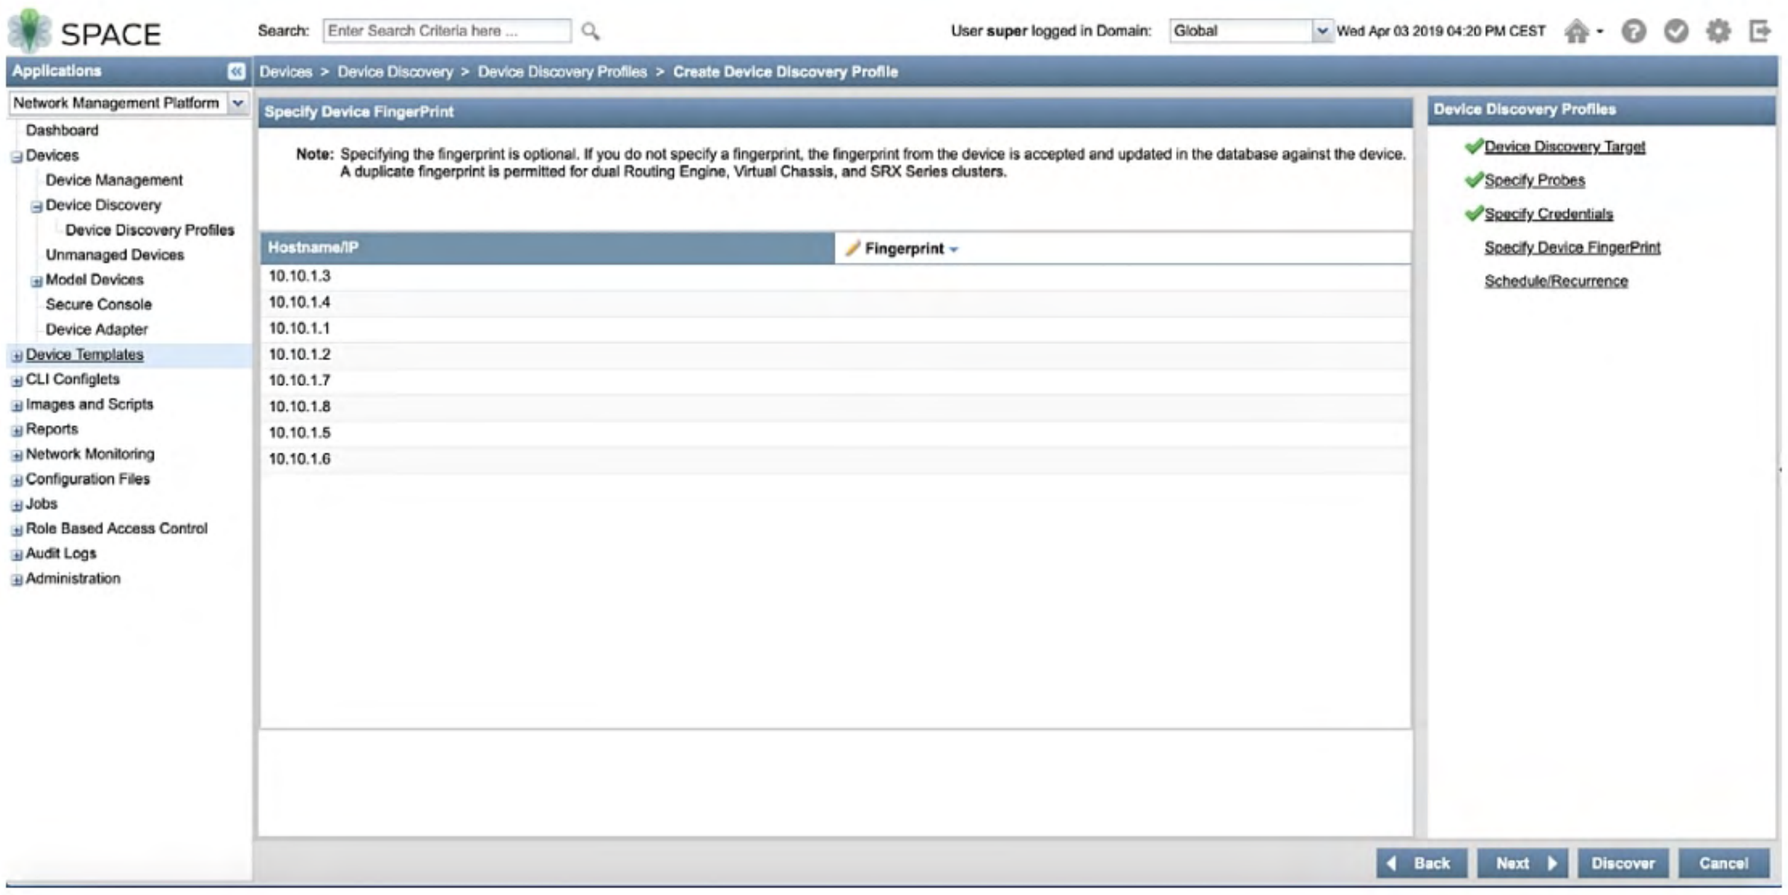

© 2020 Juniper Networks, Inc. All rights reserved.

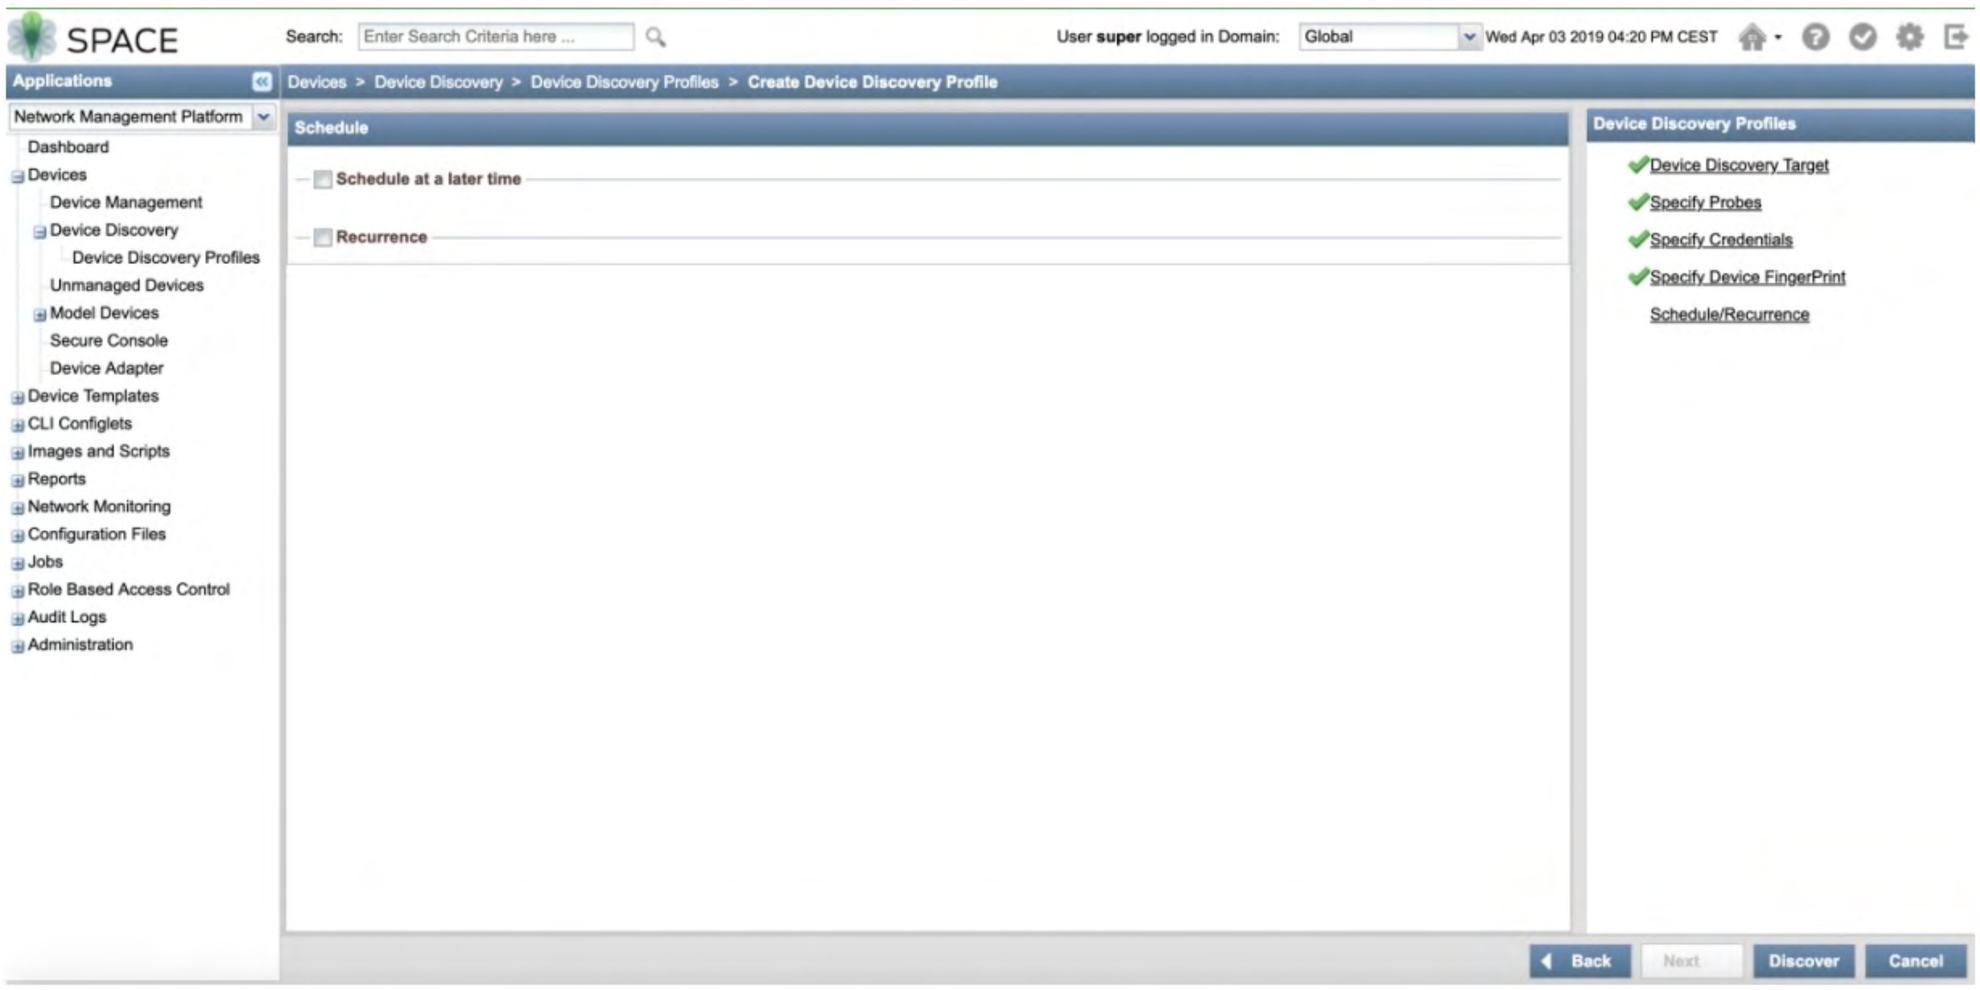

Now Junos Space start the actual discovery process. The Discovery Status shows if the devices where discovered of if the discovery process failed. Once the devices are fully discovered they will become “managed” as shown in the figure below.

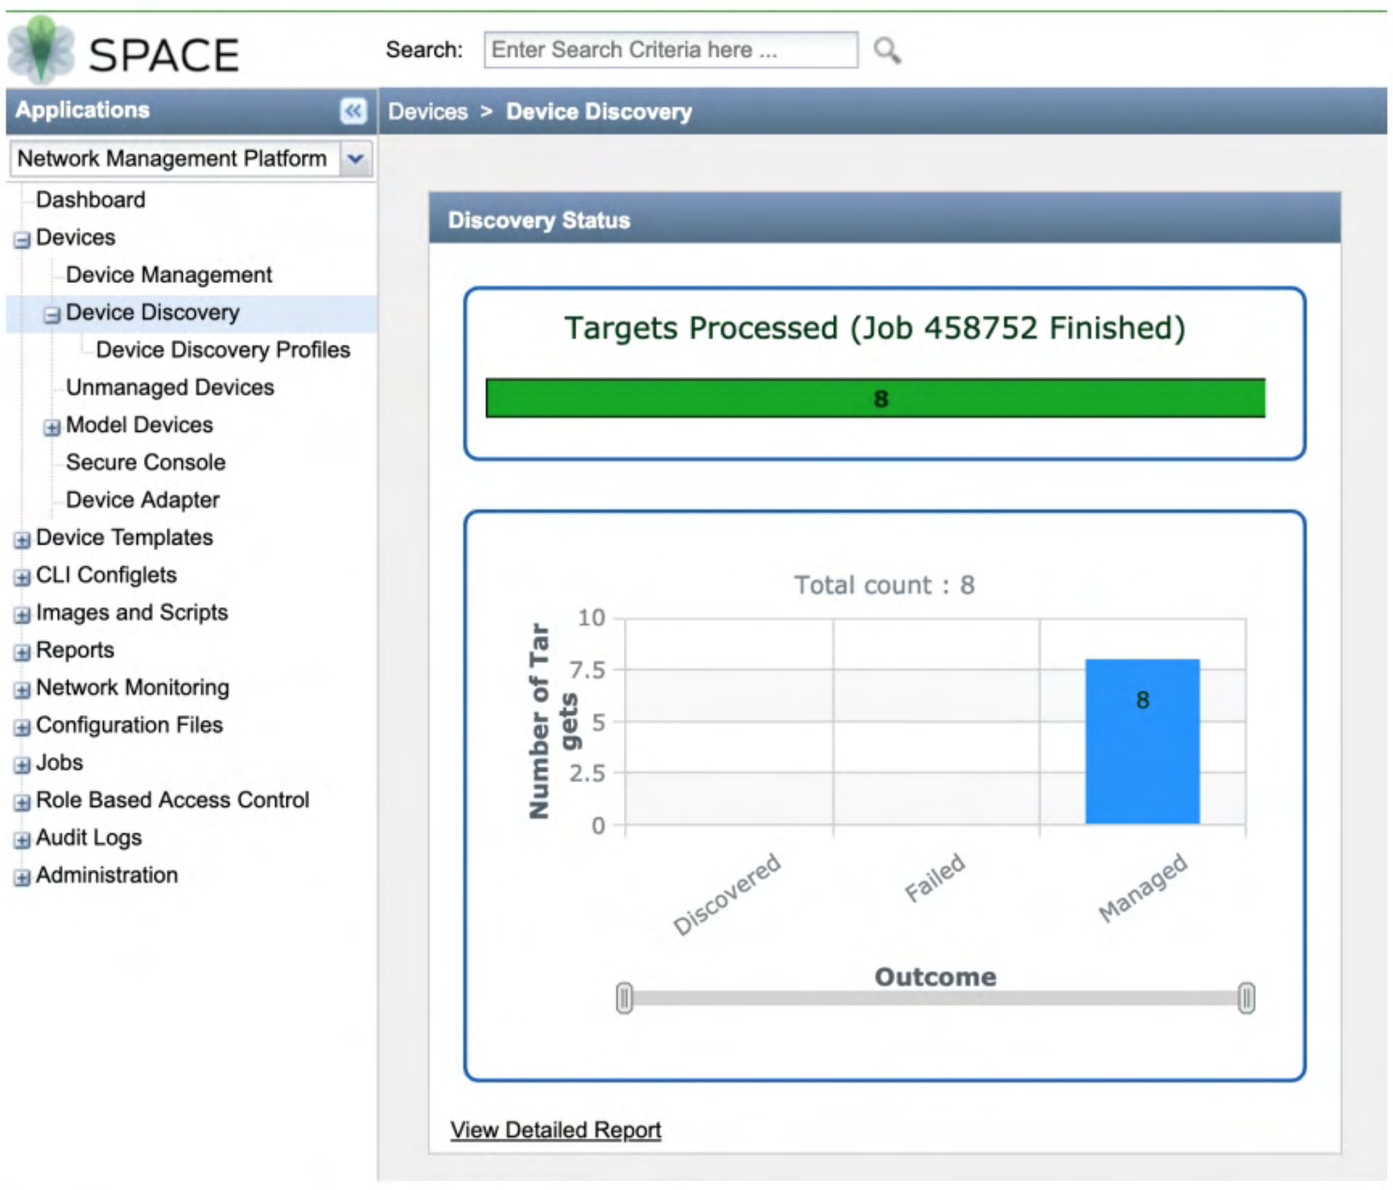

© 2020 Juniper Networks, Inc. All rights reserved.

In the left column navigate to “Devices/ Device Management” to verify if all vMX where successfully discovered. Like the figure below you should see all vMXs discovered and in Sync.

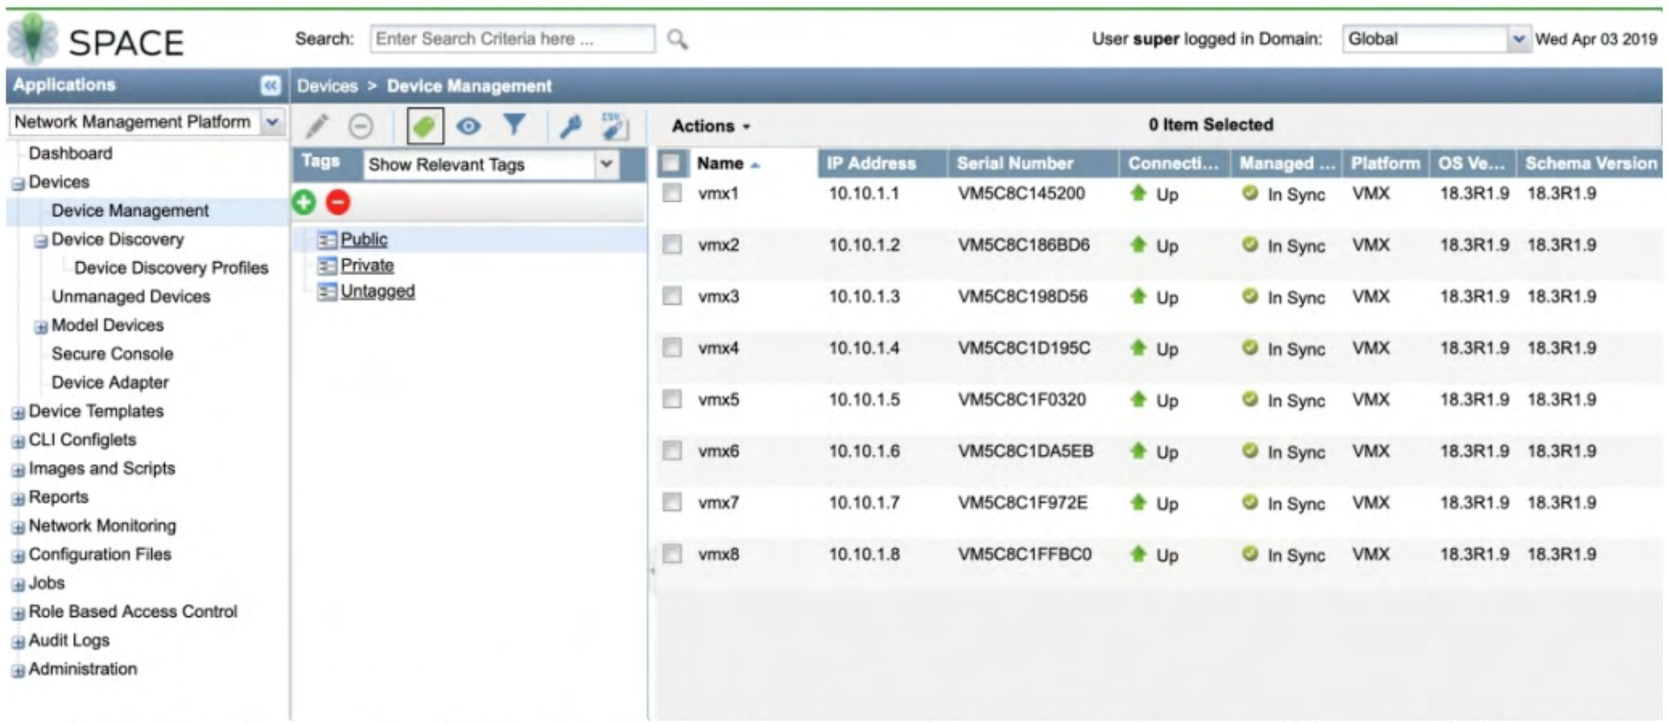

11) Back up all vMX configurations onto Junos Space Network Director.

Now that all vMX devices are discovered it is possible to backup the configuration files. In the left column, navigate to “Configuration Files” and select “Click here to backup devices configuration files”

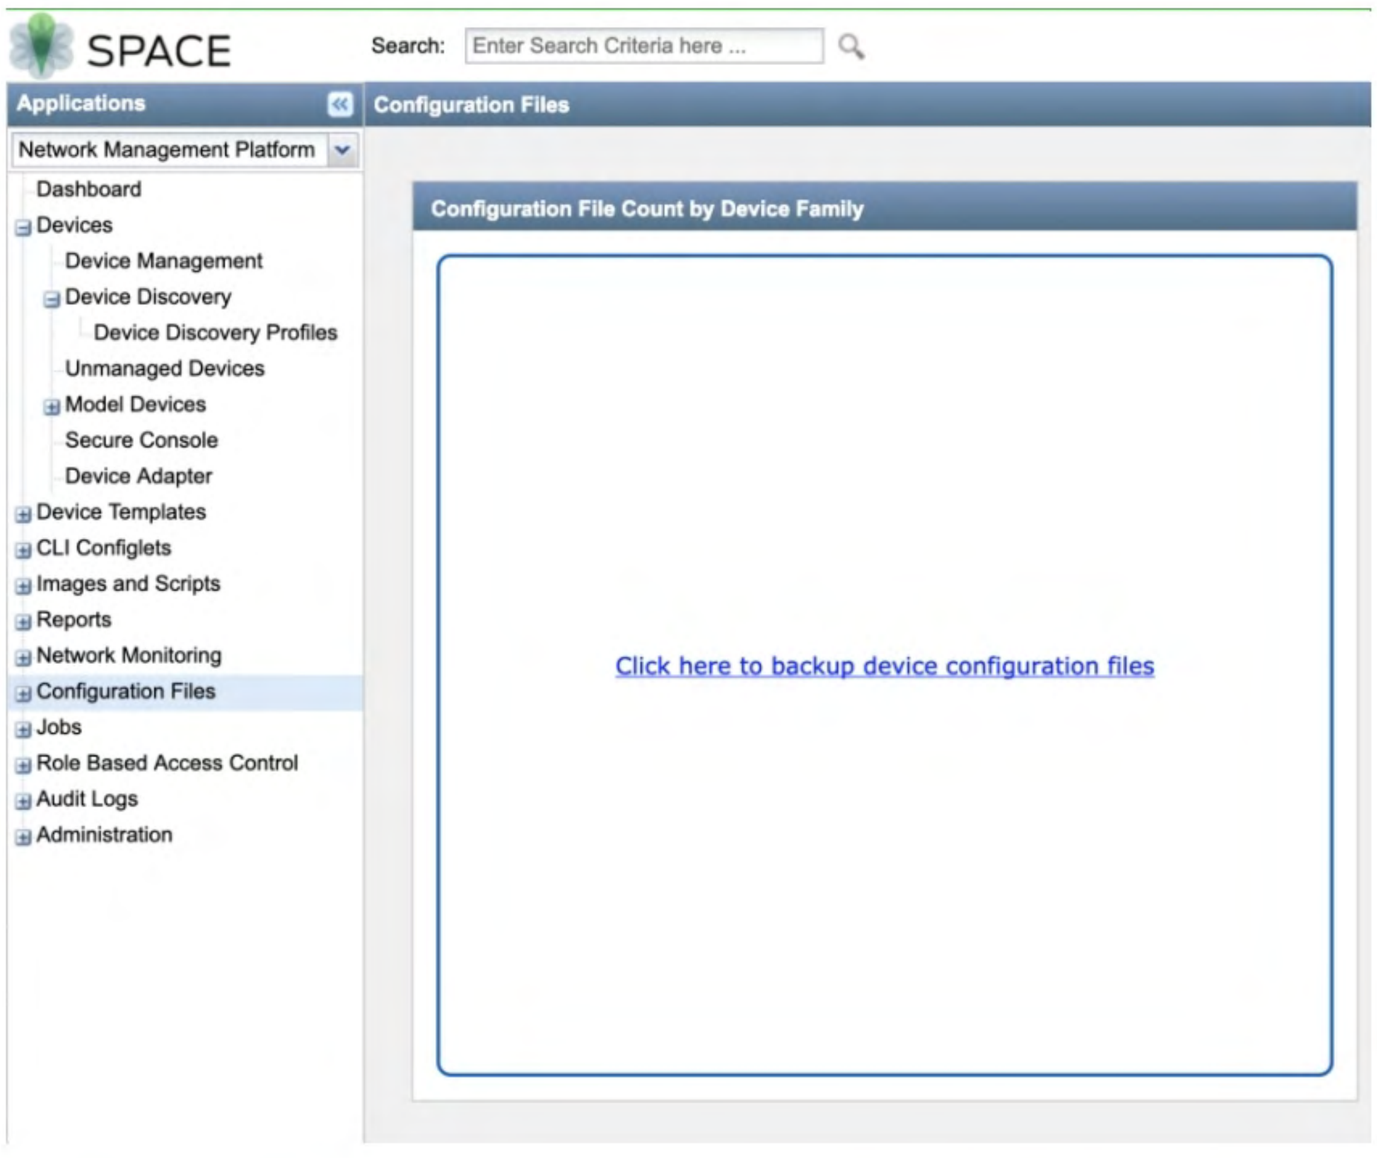

© 2020 Juniper Networks, Inc. All rights reserved.

Next, Junos Space provides a list of the devices that are eligible for configuration backup. You should see all vMX devices in the list. Scroll down and select “Backup”

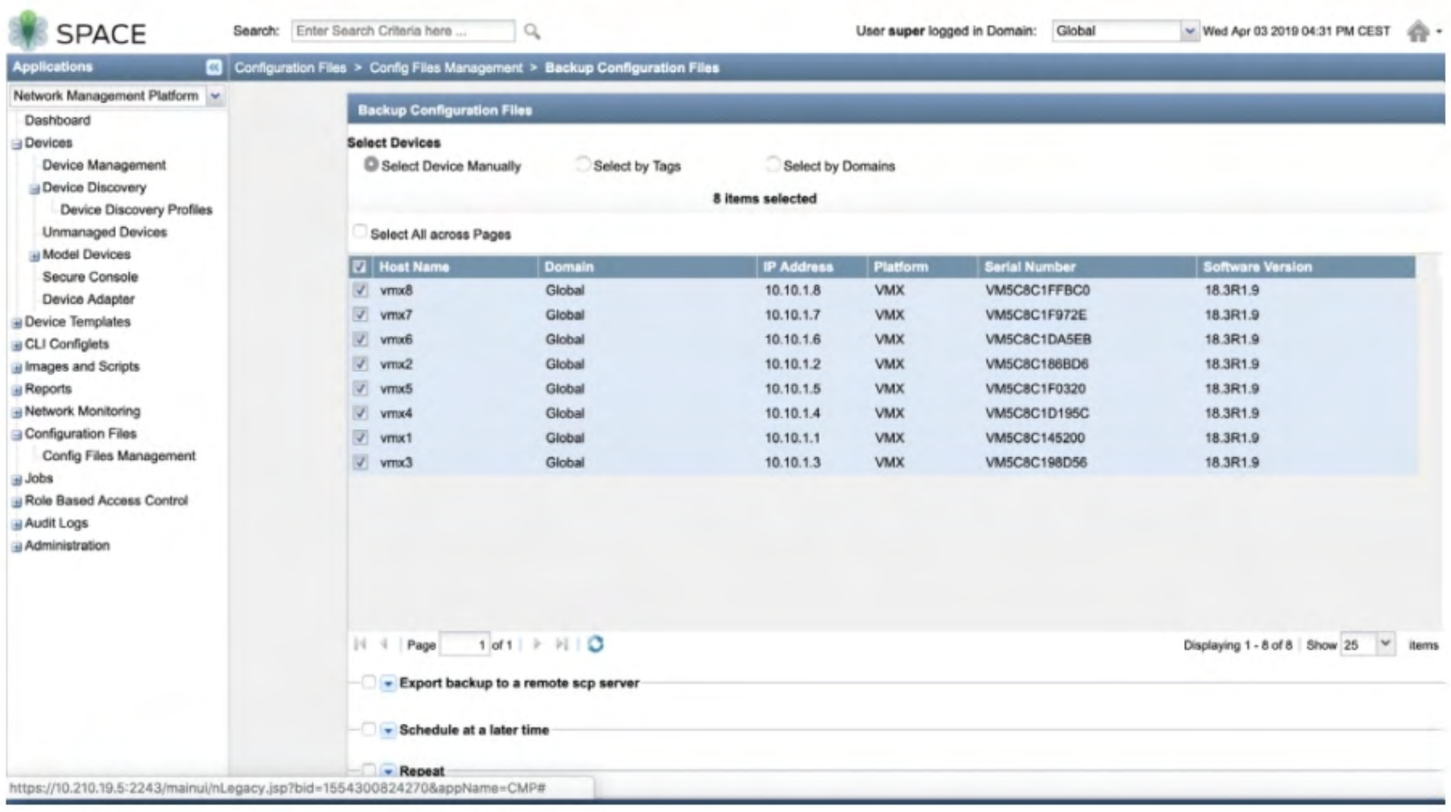

Junos Space will create a Job that runs in the background to download the devices configurations of all vMX devices. Each job is referenced by a job ID. By clicking on the Job ID you can see the details of each job if needed.

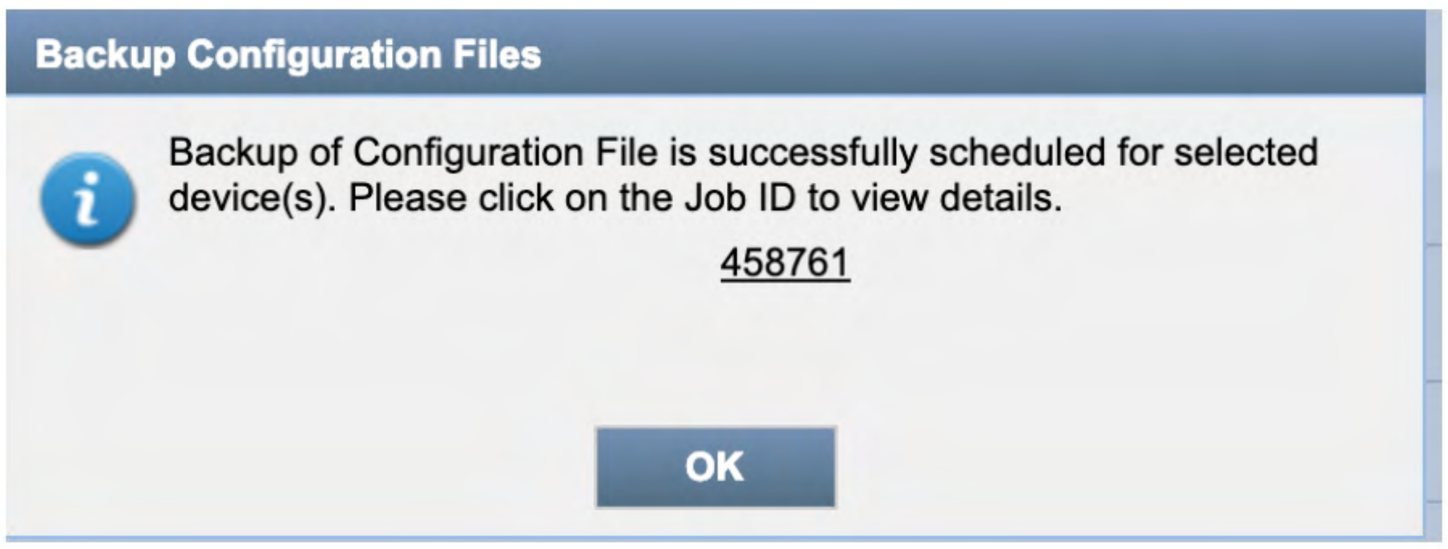

© 2020 Juniper Networks, Inc. All rights reserved.

In the left column navigate to “Configuration Files” and select “Config File Management”. You should see the configurations of all vMX devices.

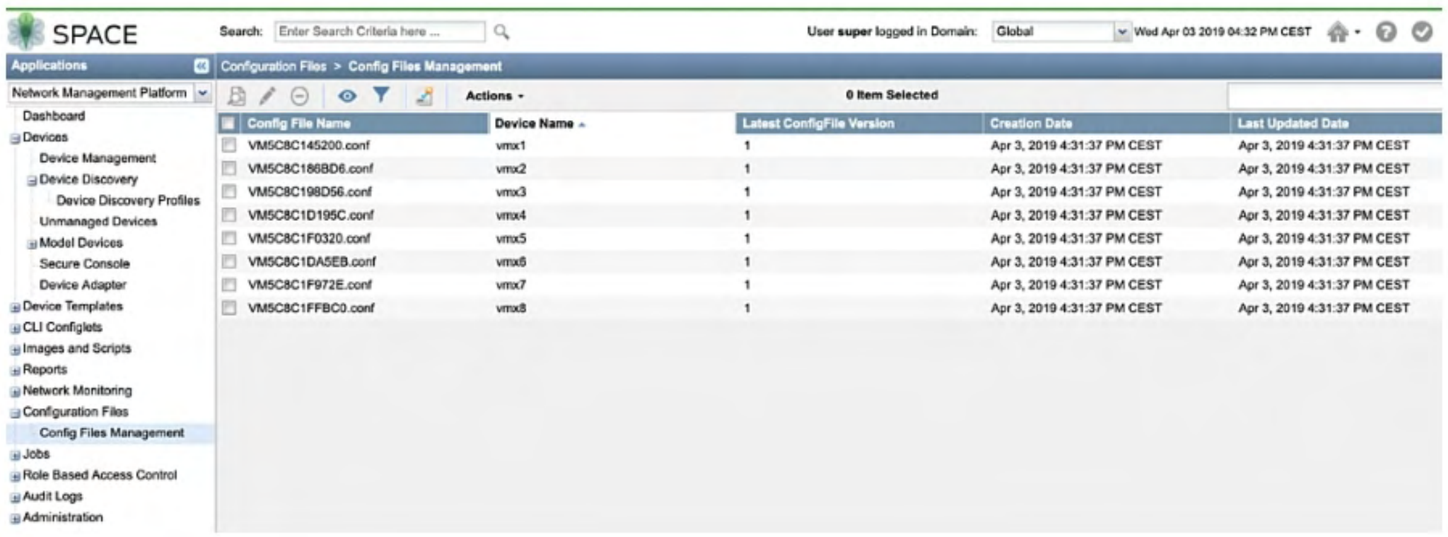

Netconf

12) Ensure that RFC 4741 compliant Python scripts can connect to vMX1 over secured shell session using Netconf.
13) Allow 6 concurrent connections and a maximum of 2 sessions per minute.
14) Prevent unsupported Junos configuration hierarchies from being exported in get-config RPC outputs

```
[edit]
lab@vmx1# show system services netconf
ssh {
    connection-limit 6;
    rate-limit 2;
}
rfc-compliant;
yang-compliant;
```

Firewall Filters

In this section you will configure RE (Routing Engine) protection using a firewall filter for all vMX devices in the data center.

15) Configure an IPv4 firewall filter which allows protocol messages from the BFD, VRRP, OSPF, BGP and GRE. Make sure that for the protocols VRRP and OSPF messages from directly connected subnets are accepted.
16) Future added BGP peers should be automatically allowed in the firewall filter.
17) Configure the firewall filter to accept RADIUS, DNS, SNMP management protocols only from the 10.10.10/24 network.
18) Accept NTP messages only from the 10.10.10/24 and 10.10.1/24 networks.
19) Accept SSH and Telnet protocols only from the 10.10.1/24 management network.

© 2020 Juniper Networks, Inc. All rights reserved.

20) Configure the devices to accept ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer, which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.
21) Protect the devices from TCP syn/rst/fin attacks. Use a policer with the following parameters: bandwidth 500kbps and burst size 125kB. The excess traffic must be dropped.
22) Allow traffic to the control plane that is required for other tasks.
23) Configure the firewall filter to discard any other traffic and count the packets.
24) Apply the firewall filter such as to ensure that it is used for the RE protection.

Configure the following on all vMX devices in the data center:

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show firewall | display omit
family inet {
    filter protect-re-ipv4 {
        apply-flags omit;
        term tcp-flood {
            from {
                protocol tcp;
                tcp-flags "(syn & !ack) | fin | rst";
            }
            then {
                policer tcp-flood-limit;
                accept;
            }
        }
        term bfd {
            from {
                source-prefix-list {
                    direct-subnets;
                }
                protocol udp;
                destination-port 3784;
            }
            then accept;
        }
        term vrrp {
            from {
                source-prefix-list {
                    direct-subnets;
                }
                protocol vrrp;
            }
            then accept;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
term ospf {
    from {
        source-prefix-list {
            direct-subnets;
        }
        protocol ospf;
    }
    then accept;
}
term bgp {
    from {
        prefix-list {
            bgp-peers;
            direct-subnets;
        }
        protocol tcp;
        port [ bgp 830 ];
    }
    then accept;
}
term ntp {
    from {
        source-address {
            10.10.10.0/24;
        }
        source-prefix-list {
            direct-subnets;
        }
        protocol udp;
        port ntp;
    }
    then accept;
}
term snmp {
    from {
        source-address {
            10.10.10.0/24;
            10.10.1.0/24;
        }
        protocol udp;
        destination-port snmp;
    }
    then accept;
}
term radius {
    from {
        source-address {
            10.10.10.0/24;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        protocol udp;
        source-port radius;
    }
    then accept;
}
term dns {
    from {
        source-address {
            10.10.10.0/24;
        }
        protocol udp;
        source-port domain;
    }
    then accept;
}
term ssh {
    from {
        source-address {
            10.10.1.0/24;
        }
        protocol tcp;
        port ssh;
    }
    then accept;
}
term telnet {
    from {
        source-address {
            10.10.1.0/24;
        }
        protocol tcp;
        port telnet;
    }
    then accept;
}
term dhcp {
    from {
        protocol udp;
        port dhcp;
    }
    then accept;
}
term icmp {
    from {
        protocol icmp;
    }
    then {
        policer re-policer;
        accept;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
    term traceroute {
        from {
            protocol udp;
            destination-port 33434-33534;
        }
        then {
            policer re-policer;
            accept;
        }
    }
    term last {
            then {
                count dropped-ipv4;
                discard;
            }
        }
    }
}
policer re-policer {
    apply-flags omit;
    if-exceeding {
        bandwidth-limit 100k;
        burst-size-limit 25k;
    }
    then discard;
}
policer tcp-flood-limit {
    apply-flags omit;
    if-exceeding {
        bandwidth-limit 500k;
        burst-size-limit 125k;
    }
    then discard;
}

[edit]
lab@vmx1# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re-ipv4;
        }
    }
}

lab@vmx1# show policy-options
prefix-list direct-subnets {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    apply-path "interfaces <*> unit <*> family inet address <*>";
}
prefix-list bgp-peers {
    apply-path "protocols bgp group <*> neighbor <*>";
}
```

Configurations for Task Completion:

vMX1:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.1
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.1
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.1
set snmp trap-group traps targets 10.10.1.99
set system services netconf ssh connection-limit 6
set system services netconf ssh rate-limit 2
set system services netconf rfc-compliant
set system services netconf yang-compliant
set firewall family inet filter protect-re-ipv4 apply-flags omit
set firewall family inet filter protect-re-ipv4 term tcp-flood from protocol tcp
set firewall family inet filter protect-re-ipv4 term tcp-flood from tcp-flags "(syn &
!ack) | fin | rst"
set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-
limit
set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
set firewall family inet filter protect-re-ipv4 term bfd from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term bfd from protocol udp
set firewall family inet filter protect-re-ipv4 term bfd from destination-port 3784
set firewall family inet filter protect-re-ipv4 term bfd then accept
set firewall family inet filter protect-re-ipv4 term vrrp from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term vrrp from protocol vrrp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term vrrp then accept
set firewall family inet filter protect-re-ipv4 term ospf from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ospf from protocol ospf
set firewall family inet filter protect-re-ipv4 term ospf then accept
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list bgp-peers
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list direct-subnets
set firewall family inet filter protect-re-ipv4 term bgp from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp from port bgp
set firewall family inet filter protect-re-ipv4 term bgp from port 830
set firewall family inet filter protect-re-ipv4 term bgp then accept
set firewall family inet filter protect-re-ipv4 term ntp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term ntp from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term ntp from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp from port ntp
set firewall family inet filter protect-re-ipv4 term ntp then accept
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term snmp from protocol udp
set firewall family inet filter protect-re-ipv4 term snmp from destination-port snmp
set firewall family inet filter protect-re-ipv4 term snmp then accept
set firewall family inet filter protect-re-ipv4 term radius from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term radius from protocol udp
set firewall family inet filter protect-re-ipv4 term radius from source-port radius
set firewall family inet filter protect-re-ipv4 term radius then accept
set firewall family inet filter protect-re-ipv4 term dns from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term dns from protocol udp
set firewall family inet filter protect-re-ipv4 term dns from source-port domain
set firewall family inet filter protect-re-ipv4 term dns then accept
set firewall family inet filter protect-re-ipv4 term ssh from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh from port ssh
set firewall family inet filter protect-re-ipv4 term ssh then accept
set firewall family inet filter protect-re-ipv4 term telnet from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet from port telnet
set firewall family inet filter protect-re-ipv4 term telnet then accept
set firewall family inet filter protect-re-ipv4 term dhcp from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp from port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp then accept
set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer
set firewall family inet filter protect-re-ipv4 term icmp then accept
set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp
set firewall family inet filter protect-re-ipv4 term traceroute from destination-port
33434-33534
set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
set firewall family inet filter protect-re-ipv4 term traceroute then accept
set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4
set firewall family inet filter protect-re-ipv4 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set firewall policer tcp-flood-limit apply-flags omit
set firewall policer tcp-flood-limit if-exceeding bandwidth-limit 500k
set firewall policer tcp-flood-limit if-exceeding burst-size-limit 125k
set firewall policer tcp-flood-limit then discard
set policy-options prefix-list direct-subnets apply-path "interfaces <*> unit <*> family
inet address <*>"
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*>"
set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

vMX2:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.2
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.2
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.2
set snmp trap-group traps targets 10.10.1.99
set firewall family inet filter protect-re-ipv4 apply-flags omit
set firewall family inet filter protect-re-ipv4 term tcp-flood from protocol tcp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term tcp-flood from tcp-flags "(syn &
!ack) | fin | rst"
set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-
limit
set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
set firewall family inet filter protect-re-ipv4 term bfd from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term bfd from protocol udp
set firewall family inet filter protect-re-ipv4 term bfd from destination-port 3784
set firewall family inet filter protect-re-ipv4 term bfd then accept
set firewall family inet filter protect-re-ipv4 term vrrp from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term vrrp from protocol vrrp
set firewall family inet filter protect-re-ipv4 term vrrp then accept
set firewall family inet filter protect-re-ipv4 term ospf from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ospf from protocol ospf
set firewall family inet filter protect-re-ipv4 term ospf then accept
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list bgp-peers
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list direct-subnets
set firewall family inet filter protect-re-ipv4 term bgp from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp from port bgp
set firewall family inet filter protect-re-ipv4 term bgp from port 830
set firewall family inet filter protect-re-ipv4 term bgp then accept
set firewall family inet filter protect-re-ipv4 term ntp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term ntp from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term ntp from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp from port ntp
set firewall family inet filter protect-re-ipv4 term ntp then accept
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term snmp from protocol udp
set firewall family inet filter protect-re-ipv4 term snmp from destination-port snmp
set firewall family inet filter protect-re-ipv4 term snmp then accept
set firewall family inet filter protect-re-ipv4 term radius from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term radius from protocol udp
set firewall family inet filter protect-re-ipv4 term radius from source-port radius
set firewall family inet filter protect-re-ipv4 term radius then accept
set firewall family inet filter protect-re-ipv4 term dns from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term dns from protocol udp
set firewall family inet filter protect-re-ipv4 term dns from source-port domain
set firewall family inet filter protect-re-ipv4 term dns then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term ssh from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh from port ssh
set firewall family inet filter protect-re-ipv4 term ssh then accept
set firewall family inet filter protect-re-ipv4 term telnet from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet from port telnet
set firewall family inet filter protect-re-ipv4 term telnet then accept
set firewall family inet filter protect-re-ipv4 term dhcp from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp from port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp then accept
set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp
set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer
set firewall family inet filter protect-re-ipv4 term icmp then accept
set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp
set firewall family inet filter protect-re-ipv4 term traceroute from destination-port
33434-33534
set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
set firewall family inet filter protect-re-ipv4 term traceroute then accept
set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4
set firewall family inet filter protect-re-ipv4 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set firewall policer tcp-flood-limit apply-flags omit
set firewall policer tcp-flood-limit if-exceeding bandwidth-limit 500k
set firewall policer tcp-flood-limit if-exceeding burst-size-limit 125k
set firewall policer tcp-flood-limit then discard
set policy-options prefix-list direct-subnets apply-path "interfaces <*> unit <*> family
inet address <*>"
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*>"
set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

vMX3:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.3
set system authentication-order radius
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.3
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.3
set snmp trap-group traps targets 10.10.1.99
set firewall family inet filter protect-re-ipv4 apply-flags omit
set firewall family inet filter protect-re-ipv4 term tcp-flood from protocol tcp
set firewall family inet filter protect-re-ipv4 term tcp-flood from tcp-flags "(syn &
!ack) | fin | rst"
set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-
limit
set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
set firewall family inet filter protect-re-ipv4 term bfd from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term bfd from protocol udp
set firewall family inet filter protect-re-ipv4 term bfd from destination-port 3784
set firewall family inet filter protect-re-ipv4 term bfd then accept
set firewall family inet filter protect-re-ipv4 term vrrp from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term vrrp from protocol vrrp
set firewall family inet filter protect-re-ipv4 term vrrp then accept
set firewall family inet filter protect-re-ipv4 term ospf from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ospf from protocol ospf
set firewall family inet filter protect-re-ipv4 term ospf then accept
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list bgp-peers
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list direct-subnets
set firewall family inet filter protect-re-ipv4 term bgp from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp from port bgp
set firewall family inet filter protect-re-ipv4 term bgp from port 830
set firewall family inet filter protect-re-ipv4 term bgp then accept
set firewall family inet filter protect-re-ipv4 term ntp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term ntp from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term ntp from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp from port ntp
set firewall family inet filter protect-re-ipv4 term ntp then accept
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term snmp from protocol udp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term snmp from destination-port snmp
set firewall family inet filter protect-re-ipv4 term snmp then accept
set firewall family inet filter protect-re-ipv4 term radius from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term radius from protocol udp
set firewall family inet filter protect-re-ipv4 term radius from source-port radius
set firewall family inet filter protect-re-ipv4 term radius then accept
set firewall family inet filter protect-re-ipv4 term dns from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term dns from protocol udp
set firewall family inet filter protect-re-ipv4 term dns from source-port domain
set firewall family inet filter protect-re-ipv4 term dns then accept
set firewall family inet filter protect-re-ipv4 term ssh from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh from port ssh
set firewall family inet filter protect-re-ipv4 term ssh then accept
set firewall family inet filter protect-re-ipv4 term telnet from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet from port telnet
set firewall family inet filter protect-re-ipv4 term telnet then accept
set firewall family inet filter protect-re-ipv4 term dhcp from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp from port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp then accept
set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp
set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer
set firewall family inet filter protect-re-ipv4 term icmp then accept
set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp
set firewall family inet filter protect-re-ipv4 term traceroute from destination-port
33434-33534
set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
set firewall family inet filter protect-re-ipv4 term traceroute then accept
set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4
set firewall family inet filter protect-re-ipv4 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set firewall policer tcp-flood-limit apply-flags omit
set firewall policer tcp-flood-limit if-exceeding bandwidth-limit 500k
set firewall policer tcp-flood-limit if-exceeding burst-size-limit 125k
set firewall policer tcp-flood-limit then discard
set policy-options prefix-list direct-subnets apply-path "interfaces <*> unit <*> family
inet address <*>"
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*>"
set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX4:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.4
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.4
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.4
set snmp trap-group traps targets 10.10.1.99
set firewall family inet filter protect-re-ipv4 apply-flags omit
set firewall family inet filter protect-re-ipv4 term tcp-flood from protocol tcp
set firewall family inet filter protect-re-ipv4 term tcp-flood from tcp-flags "(syn &
!ack) | fin | rst"
set firewall family inet filter protect-re-ipv4 term tcp-flood then policer tcp-flood-
limit
set firewall family inet filter protect-re-ipv4 term tcp-flood then accept
set firewall family inet filter protect-re-ipv4 term bfd from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term bfd from protocol udp
set firewall family inet filter protect-re-ipv4 term bfd from destination-port 3784
set firewall family inet filter protect-re-ipv4 term bfd then accept
set firewall family inet filter protect-re-ipv4 term vrrp from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term vrrp from protocol vrrp
set firewall family inet filter protect-re-ipv4 term vrrp then accept
set firewall family inet filter protect-re-ipv4 term ospf from source-prefix-list
direct-subnets
set firewall family inet filter protect-re-ipv4 term ospf from protocol ospf
set firewall family inet filter protect-re-ipv4 term ospf then accept
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list bgp-peers
set firewall family inet filter protect-re-ipv4 term bgp from prefix-list direct-subnets
set firewall family inet filter protect-re-ipv4 term bgp from protocol tcp
set firewall family inet filter protect-re-ipv4 term bgp from port bgp
set firewall family inet filter protect-re-ipv4 term bgp from port 830
set firewall family inet filter protect-re-ipv4 term bgp then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet filter protect-re-ipv4 term ntp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term ntp from source-prefix-list direct-
subnets
set firewall family inet filter protect-re-ipv4 term ntp from protocol udp
set firewall family inet filter protect-re-ipv4 term ntp from port ntp
set firewall family inet filter protect-re-ipv4 term ntp then accept
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term snmp from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term snmp from protocol udp
set firewall family inet filter protect-re-ipv4 term snmp from destination-port snmp
set firewall family inet filter protect-re-ipv4 term snmp then accept
set firewall family inet filter protect-re-ipv4 term radius from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term radius from protocol udp
set firewall family inet filter protect-re-ipv4 term radius from source-port radius
set firewall family inet filter protect-re-ipv4 term radius then accept
set firewall family inet filter protect-re-ipv4 term dns from source-address
10.10.10.0/24
set firewall family inet filter protect-re-ipv4 term dns from protocol udp
set firewall family inet filter protect-re-ipv4 term dns from source-port domain
set firewall family inet filter protect-re-ipv4 term dns then accept
set firewall family inet filter protect-re-ipv4 term ssh from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term ssh from protocol tcp
set firewall family inet filter protect-re-ipv4 term ssh from port ssh
set firewall family inet filter protect-re-ipv4 term ssh then accept
set firewall family inet filter protect-re-ipv4 term telnet from source-address
10.10.1.0/24
set firewall family inet filter protect-re-ipv4 term telnet from protocol tcp
set firewall family inet filter protect-re-ipv4 term telnet from port telnet
set firewall family inet filter protect-re-ipv4 term telnet then accept
set firewall family inet filter protect-re-ipv4 term dhcp from protocol udp
set firewall family inet filter protect-re-ipv4 term dhcp from port dhcp
set firewall family inet filter protect-re-ipv4 term dhcp then accept
set firewall family inet filter protect-re-ipv4 term icmp from protocol icmp
set firewall family inet filter protect-re-ipv4 term icmp then policer re-policer
set firewall family inet filter protect-re-ipv4 term icmp then accept
set firewall family inet filter protect-re-ipv4 term traceroute from protocol udp
set firewall family inet filter protect-re-ipv4 term traceroute from destination-port
33434-33534
set firewall family inet filter protect-re-ipv4 term traceroute then policer re-policer
set firewall family inet filter protect-re-ipv4 term traceroute then accept
set firewall family inet filter protect-re-ipv4 term last then count dropped-ipv4
set firewall family inet filter protect-re-ipv4 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set firewall policer tcp-flood-limit apply-flags omit
set firewall policer tcp-flood-limit if-exceeding bandwidth-limit 500k
set firewall policer tcp-flood-limit if-exceeding burst-size-limit 125k
set firewall policer tcp-flood-limit then discard
set policy-options prefix-list direct-subnets apply-path "interfaces <*> unit <*> family
inet address <*>"
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*>"
set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

vMX5:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.5
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.5
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.5
set snmp trap-group traps targets 10.10.1.99
```

vMX6:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.6
set system authentication-order radius
set system authentication-order password
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.6
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.6
set snmp trap-group traps targets 10.10.1.99
```

vMX7:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.7
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.7
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.7
set snmp trap-group traps targets 10.10.1.99
```

vMX8:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.8
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.8
set system login user admin uid 2003
set system login user admin class super-user
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.8
set snmp trap-group traps targets 10.10.1.99
```

EX1:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system name-server 10.10.10.1
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.1
set system authentication-order radius
set system authentication-order password
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.1
set system login user admin uid 2003
set system login user admin class super-user
```

vQFX1:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.13
set snmp trap-group traps targets 10.10.1.99
```

vQFX2:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.14
set snmp trap-group traps targets 10.10.1.99
```

vSRX:

```
set system services ssh root-login allow
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.10
set snmp trap-group traps targets 10.10.1.99
```

Task 2: Layer 2 Configuration and Features

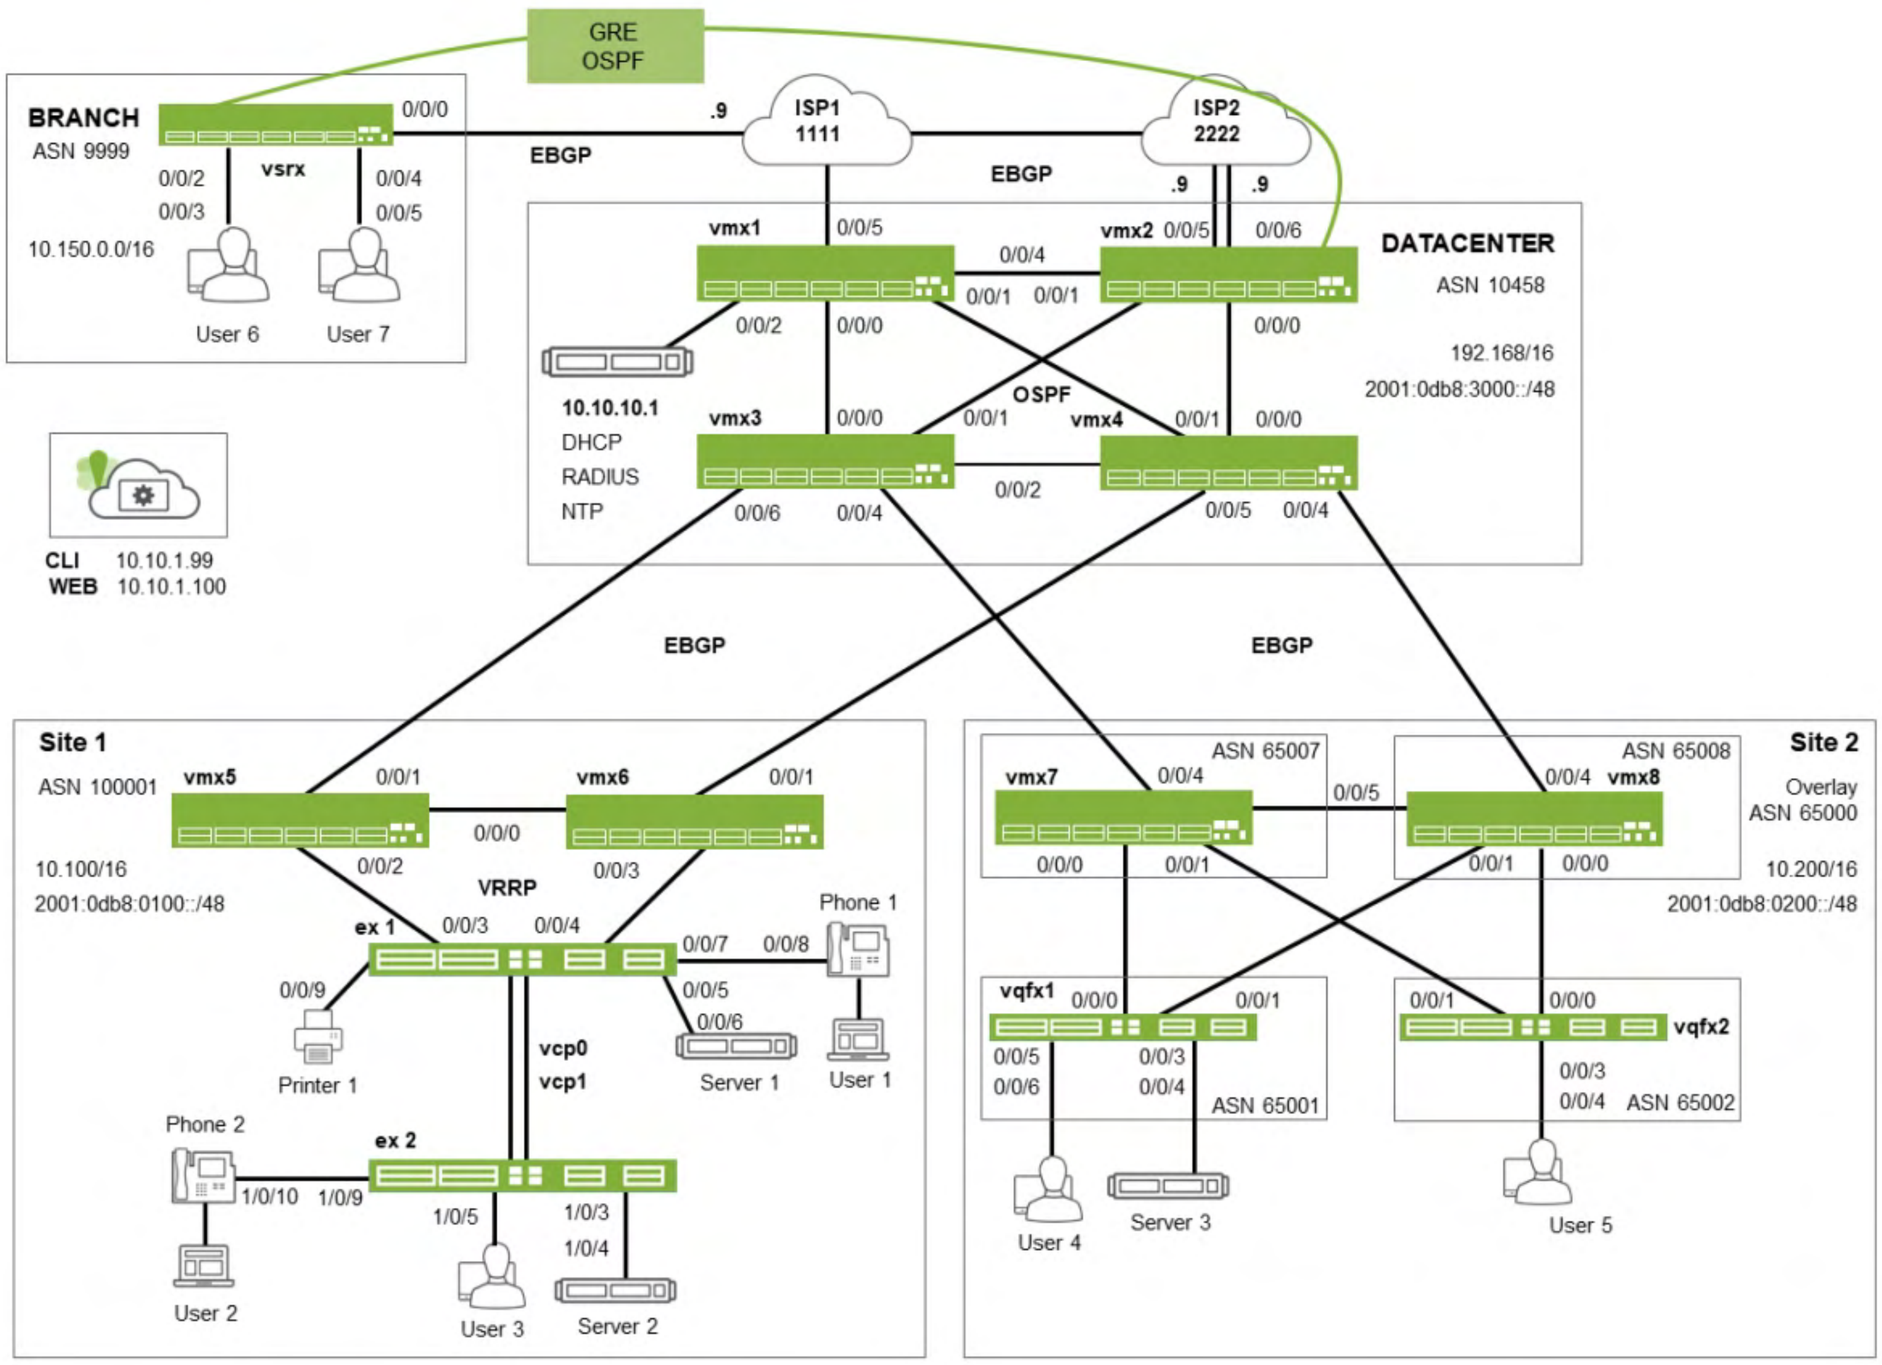

Virtual Chassis

In this section you will configure a Virtual chassis consisting of EX1 and EX2.

1) Configure EX1 and EX2 as a virtual chassis. Ensure EX1 is FPC0 and EX2 is FPC1. Ensure the hostname for the virtual-chassis is EX1
2) Configure Juniper recommended priority settings for a two-member Virtual Chassis.
3) Ensure that the master device retains mastership by not taking the linecard role in case the backup device fails.
4) Virtual chassis master election events should be logged in detail to a file called “me”.

© 2020 Juniper Networks, Inc. All rights reserved.

EX1:

```
lab@ex1> request virtual-chassis vc-port set interface vcp-255/1/0 member 0

lab@ex1> request virtual-chassis vc-port set interface vcp-255/1/1 member 0

lab@ex1# show virtual-chassis
no-split-detection;
traceoptions {
    file me;
    flag me detail;
}
member 0 {
    mastership-priority 255;
}
member 1 {
    mastership-priority 255;
}
```

EX2:

```
lab@ex2> request virtual-chassis vc-port set interface vcp-255/1/0

lab@ex2> request virtual-chassis vc-port set interface vcp-255/1/1
```

Verification:

```
lab@ex1> show virtual-chassis vc-port
fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
1/0         Configured          5    Up           40000        1   vcp-255/1/0
1/1         Configured          5    Up           40000        1   vcp-255/1/1
1/3         Configured               Absent
1/2         Configured               Absent

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
1/0         Configured          5    Up           40000        0   vcp-255/1/0
1/1         Configured          5    Up           40000        0   vcp-255/1/1
1/3         Configured               Absent
1/2         Configured               Absent
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@ex1> show virtual-chassis status

Virtual Chassis ID: 8f14.33f4.8988
Virtual Chassis Mode: Enabled
                                                  Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model            prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    PG3718460071 ex4300-24t       255   Master*      N  VC   1  vcp-
255/1/0
                                                                           1  vcp-
255/1/1
1 (FPC 1)  Prsnt    PG3718470055 ex4300-24t       255   Backup       N  VC   0  vcp-
255/1/0
                                                                           0  vcp-
255/1/1
```

Add the interface configuration and routing instances for the interfaces that are connected to EX2 on the Virtual Chassis master (EX1).

```
lab@ex1# show interfaces | find ge-1
ge-1/0/3 {
    description "to server2";
}
ge-1/0/4 {
    description "server 2";
}
ge-1/0/5 {
    description "to user3";
}
ge-1/0/9 {
    description "to phone2 user2";
}
ge-1/0/10 {
    description "phone2 user2";
    unit 0 {
        family inet {
            address 10.100.100.2/24;
        }
        family inet6 {
            address 2001:0db8:0100:100::2/64;
        }
    }
}
...

lab@ex1# show routing-instances
server1 {
    instance-type virtual-router;
    interface ge-0/0/6.0;
    routing-options {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        rib server1.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:101::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.101.254;
        }
    }
}
server2 {
    instance-type virtual-router;
    interface ge-1/0/4.0;
    routing-options {
        rib server2.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:101::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.101.254;
        }
    }
}
user1 {
    instance-type virtual-router;
    interface ge-0/0/8.0;
    routing-options {
        rib user1.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:100::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.100.254;
        }
    }
}
user2 {
    instance-type virtual-router;
    interface ge-1/0/10.0;
    routing-options {
        rib user2.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:100::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.100.254;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
}
```

VLAN Configuration

In this section you will configure (VLAN) interfaces on the switches.

VLAN name	VLAN ID	Interface	Devices
user	100	ge-0/0/7 untagged	EX1/ VC
		ge-0/0/5 untagged	
		ge-1/0/9 untagged	
		ge-0/0/5 untagged	vQFX1
		ge-0/0/3 untagged	vQFX2
server	101	ge-0/0/5 untagged	EX1/ VC
		ge-1/0/3 untagged	
		ge-0/0/3 untagged	vQFX1
printer	200	ge-0/0/9 untagged	
phone	201	ge-0/0/7 tagged	EX1 / VC
		ge-1/0/9 tagged	
guest	666	Dot1x guest vlan	
vmx5vmx6ex1	156	ge-0/0/3	
		ge-0/0/4	

Table 39 VLANs Site1 and Site 2

5) Configure the VLANs and corresponding interfaces in Site1 and Site 2 as shown in the table above

EX1:

```
lab@ex1# show vlans
default {
    vlan-id 1;
}
guest {
    vlan-id 666;
}
phone {
    vlan-id 201;
    l3-interface irb.201;
}
printer {
    vlan-id 200;
    l3-interface irb.200;
}
server {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    vlan-id 101;
    l3-interface irb.101;
}
user {
    vlan-id 100;
    l3-interface irb.100;
    switch-options {
        interface ge-0/0/7.0;
    }
}
vmx5vmx6ex1 {
    vlan-id 156;
    l3-interface irb.156;
}

lab@ex1# show interfaces
ge-0/0/3 {
    description "to vMX5";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members vmx5vmx6ex1;
            }
        }
    }
}
ge-0/0/4 {
    description "to vMX4";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members vmx5vmx6ex1;
            }
        }
    }
}
ge-0/0/5 {
    description "to server1";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members server;
            }
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/7 {
    description "to phone1 user1";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members user;
            }
        }
    }
}
ge-0/0/9 {
    description "to printer1";
    unit 0 {
        family ethernet-switching {
            vlan {
                members printer;
            }
        }
    }
}
ge-1/0/3 {
    description "to server2";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members server;
            }
        }
    }
}
ge-1/0/5 {
    description "to user3";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members user;
            }
        }
    }
}
ge-1/0/9 {
    description "to phone2 user2";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                members user;
            }
        }
    }
}
```

vQFX1:

```
lab@vQFX1# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
}
user {
    vlan-id 100;
}

lab@vQFX1# show interfaces xe-0/0/3
description "to server3";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members server;
        }
    }
}

lab@vQFX1# show interfaces xe-0/0/5
description "to user4";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```

vQFX2:

```
lab@vQFX2# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
user {
    vlan-id 100;
}

lab@vQFX2# show interfaces xe-0/0/3
description "to user5";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```

6) Configure interface ge-0/0/7 and interface ge-1/0/9 on virtual chassis EX1 for voice VLAN support. Make sure that tagged frames go to VLAN201 whereas untagged frames go to VLAN100.

```
lab@ex1# show switch-options

voip {
    interface ge-0/0/7.0 {
        vlan phone;
    }
    interface ge-1/0/9.0 {
        vlan phone;
    }
}

lab@ex1# show vlans user
vlan-id 100;
l3-interface irb.100;
switch-options {
    interface ge-0/0/7.0;
    interface ge-1/0/9.0;
}
```

Layer 2 Features

7) Ensure Server1 receives an IPv4 address from the DHCP server.

Missing Command split into two commands.

```
lab@ex1# show interfaces ge-0/0/6
description server1;
unit 0 {
    family inet {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        dhcp;
    }
    family inet6 {
        address 2001:0db8:0100:101::1/64;
    }
}

lab@ex1# show interfaces ge-1/0/4
description "server 2";
unit 0 {
    family inet {
        dhcp;
    }
}

lab@ex1# show forwarding-options
dhcp-relay {
    overrides;
    server-group {
        dhcp-server {
            10.10.10.1;
        }
    }
    group dhcp-server {
        active-server-group dhcp-server;
        interface irb.101;
    }
}

lab@ex1# show interfaces irb.101
family inet {
    address 10.100.101.254/24;
}
family inet6 {
    address 2001:0db8:0100:101::254/64;
}
```

Verification:

Note: You will not be able to verify this task until after you have enabled routing between Site 1 and the data center. The following outputs were taken after completing the BGP configurations in Task 5.

Note: You might need to deactivate and activate the ge-0/0/4 interface to trigger DHCP requests.

```
lab@ex1# run show dhcp client binding routing-instance server1
IP address        Hardware address   Expires     State      Interface
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
10.100.101.50     cc:e1:94:d2:6e:49  24930       BOUND      ge-0/0/6.0

{master:0}[edit]
lab@ex1# run show dhcp client binding routing-instance server2
IP address        Hardware address   Expires     State      Interface
10.100.101.51     b8:c2:53:97:97:a7  24931       BOUND      ge-1/0/4.0
```

8) Printer1 must be authenticated dynamically, but EX1 must not attempt 802.1x authentication. If authentication fails or the server is unreachable, the printer must be given guest access.

Note: You can use the RADIUS server at 10.10.10.1 using the secret 'juniper123'. You should source the requests from the IRB interface that is assigned to the printer VLAN.

```
lab@ex1# show access
radius-server {
    10.10.10.1 {
        secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"; ## SECRET-DATA
        source-address 10.100.200.254;
    }
}
profile jncie {
    authentication-order radius;
    radius {
        authentication-server 10.10.10.1;
    }
}

lab@ex1# show protocols dot1x
authenticator {
    authentication-profile-name jncie;
    interface {
        ge-0/0/9.0 {
            mac-radius {
                restrict;
            }
            guest-vlan guest;
            server-fail vlan-name guest;
        }
    }
}
```

Verification:

Note: You will not be able to verify this task until after you have enabled routing between Site 1 and the data center. The following outputs were taken after completing the BGP configurations in Task 5.

© 2020 Juniper Networks, Inc. All rights reserved.

Log into the printer VM to initiate traffic to EX1's IRB interface in the printer VLAN.

Printer:

```
-bash-3.2$ ping 10.100.200.254
PING 10.100.200.254 (10.100.200.254) 56(84) bytes of data.
64 bytes from 10.100.200.254: icmp_seq=1 ttl=64 time=1003 ms
64 bytes from 10.100.200.254: icmp_seq=2 ttl=64 time=3.02 ms
64 bytes from 10.100.200.254: icmp_seq=3 ttl=64 time=2.35 ms

--- 10.100.200.254 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 2.351/336.355/1003.687/471.875 ms, pipe 2
-bash-3.2$
```

This will initiate traffic which should cause EX1 to authenticate the MAC address with the RADIUS server.

EX1:

```
{master:0}[edit]
lab@ex1# run show dot1x interface
802.1X Information:
Interface     Role           State          MAC address          User
ge-0/0/9.0    Authenticator  Authenticated  00:50:56:00:00:04    005056000004
```

9) Configure interface ge-0/0/7 on EX1 and interface ge-1/0/9 on EX2 for POE and LLDP-MED support. Send location-id TLVs. No more then 20 WATT should be allocated the POE ports.

```
lab@ex1# show poe
interface ge-0/0/7 {
    maximum-power 20;
}
interface ge-1/0/9 {
    maximum-power 20;
}

lab@ex1# show protocols lldp-med
interface ge-0/0/7 {
    tlv-select location-id;
}
```

10) The maximum IP/MAC bindings should be 5 on EX2 port ge-1/0/9.

```
lab@ex1# show switch-options
interface ge-1/0/9.0 {
    interface-mac-ip-limit {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        5;
    }
}
```

11) Configure EX1 to shutdown a port if Unknown Unicast or Broadcast traffic consumes more than 5% of the interface bandwidth. Multicast traffic should not be included in your solution.

```
lab@ex1# show forwarding-options
storm-control-profiles default {
    all {
        bandwidth-percentage 5;
        no-multicast;
    }
    action-shutdown;
}
```

Configurations for Task Completion:

EX1:

```
request virtual-chassis vc-port set interface vcp-255/1/0 member 0
request virtual-chassis vc-port set interface vcp-255/1/1 member 0
set virtual-chassis no-split-detection
set virtual-chassis member 0 mastership-priority 255
set virtual-chassis member 1 mastership-priority 255
set virtual-chassis traceoptions file me
set virtual-chassis traceoptions flag me detail
set interfaces ge-1/0/3 description "to server2"
set interfaces ge-1/0/4 description "server 2"
set interfaces ge-1/0/5 description "to user3"
set interfaces ge-1/0/9 description "to phone2 user2"
set interfaces ge-1/0/10 description "phone2 user2"
set interfaces ge-1/0/10 unit 0 family inet address 10.100.100.2/24
set interfaces ge-1/0/10 unit 0 family inet6 address 2001:0db8:0100:100::2/64
set routing-instances server2 instance-type virtual-router
set routing-instances server2 interface ge-1/0/4.0
set routing-instances server2 routing-options rib server2.inet6.0 static route 0::/0
next-hop 2001:0db8:0100:101::254
set routing-instances server2 routing-options static route 0.0.0.0/0 next-hop
10.100.101.254
set routing-instances user2 instance-type virtual-router
set routing-instances user2 interface ge-1/0/10.0
set routing-instances user2 routing-options rib user2.inet6.0 static route 0::/0 next-
hop 2001:0db8:0100:100::254
set routing-instances user2 routing-options static route 0.0.0.0/0 next-hop
10.100.100.254
set vlans guest vlan-id 666
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set vlans phone vlan-id 201
set vlans phone l3-interface irb.201
set vlans printer vlan-id 200
set vlans printer l3-interface irb.200
set vlans server vlan-id 101
set vlans server l3-interface irb.101
set vlans user vlan-id 100
set vlans user l3-interface irb.100
set vlans vmx5vmx6ex1 vlan-id 156
set vlans vmx5vmx6ex1 l3-interface irb.156
set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members vmx5vmx6ex1
set interfaces ge-0/0/4 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members vmx5vmx6ex1
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members server
set interfaces ge-0/0/7 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/7 unit 0 family ethernet-switching vlan members user
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members printer
set interfaces ge-1/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/3 unit 0 family ethernet-switching vlan members server
set interfaces ge-1/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/5 unit 0 family ethernet-switching vlan members user
set interfaces ge-1/0/9 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/9 unit 0 family ethernet-switching vlan members user
set switch-options voip interface ge-0/0/7.0 vlan phone
set switch-options voip interface ge-1/0/9.0 vlan phone
set vlans user switch-options interface ge-0/0/7.0
set vlans user switch-options interface ge-1/0/9.0
set interfaces ge-0/0/6 unit 0 family inet dhcp
set interfaces ge-1/0/4 unit 0 family inet dhcp
set forwarding-options dhcp-relay overrides
set forwarding-options dhcp-relay server-group dhcp-server 10.10.10.1
set forwarding-options dhcp-relay group dhcp-server active-server-group dhcp-server
set forwarding-options dhcp-relay group dhcp-server interface irb.101
set access radius-server 10.10.10.1 secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"
set access radius-server 10.10.10.1 source-address 10.100.200.254
set access profile jncie authentication-order radius
set access profile jncie radius authentication-server 10.10.10.1
set protocols dot1x authenticator authentication-profile-name jncie
set protocols dot1x authenticator interface ge-0/0/9.0 mac-radius restrict
set protocols dot1x authenticator interface ge-0/0/9.0 guest-vlan guest
set protocols dot1x authenticator interface ge-0/0/9.0 server-fail vlan-name guest
set poe interface ge-0/0/7 maximum-power 20
set poe interface ge-1/0/9 maximum-power 20
set protocols lldp-med interface ge-0/0/7 tlv-select location-id
set switch-options interface ge-1/0/9.0 interface-mac-ip-limit 5
set forwarding-options storm-control-profiles default all bandwidth-percentage 5
set forwarding-options storm-control-profiles default all no-multicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set forwarding-options storm-control-profiles default action-shutdown
```

EX2:

```
request virtual-chassis vc-port set interface vcp-255/1/0
request virtual-chassis vc-port set interface vcp-255/1/1
```

vQFX1:

```
set vlans user vlan-id 100
set vlans server vlan-id 101
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members server
set interfaces xe-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching vlan members user
```

vQFX2:

```
set vlans user vlan-id 100
set vlans server vlan-id 101
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members user
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: Protocol Independent Features

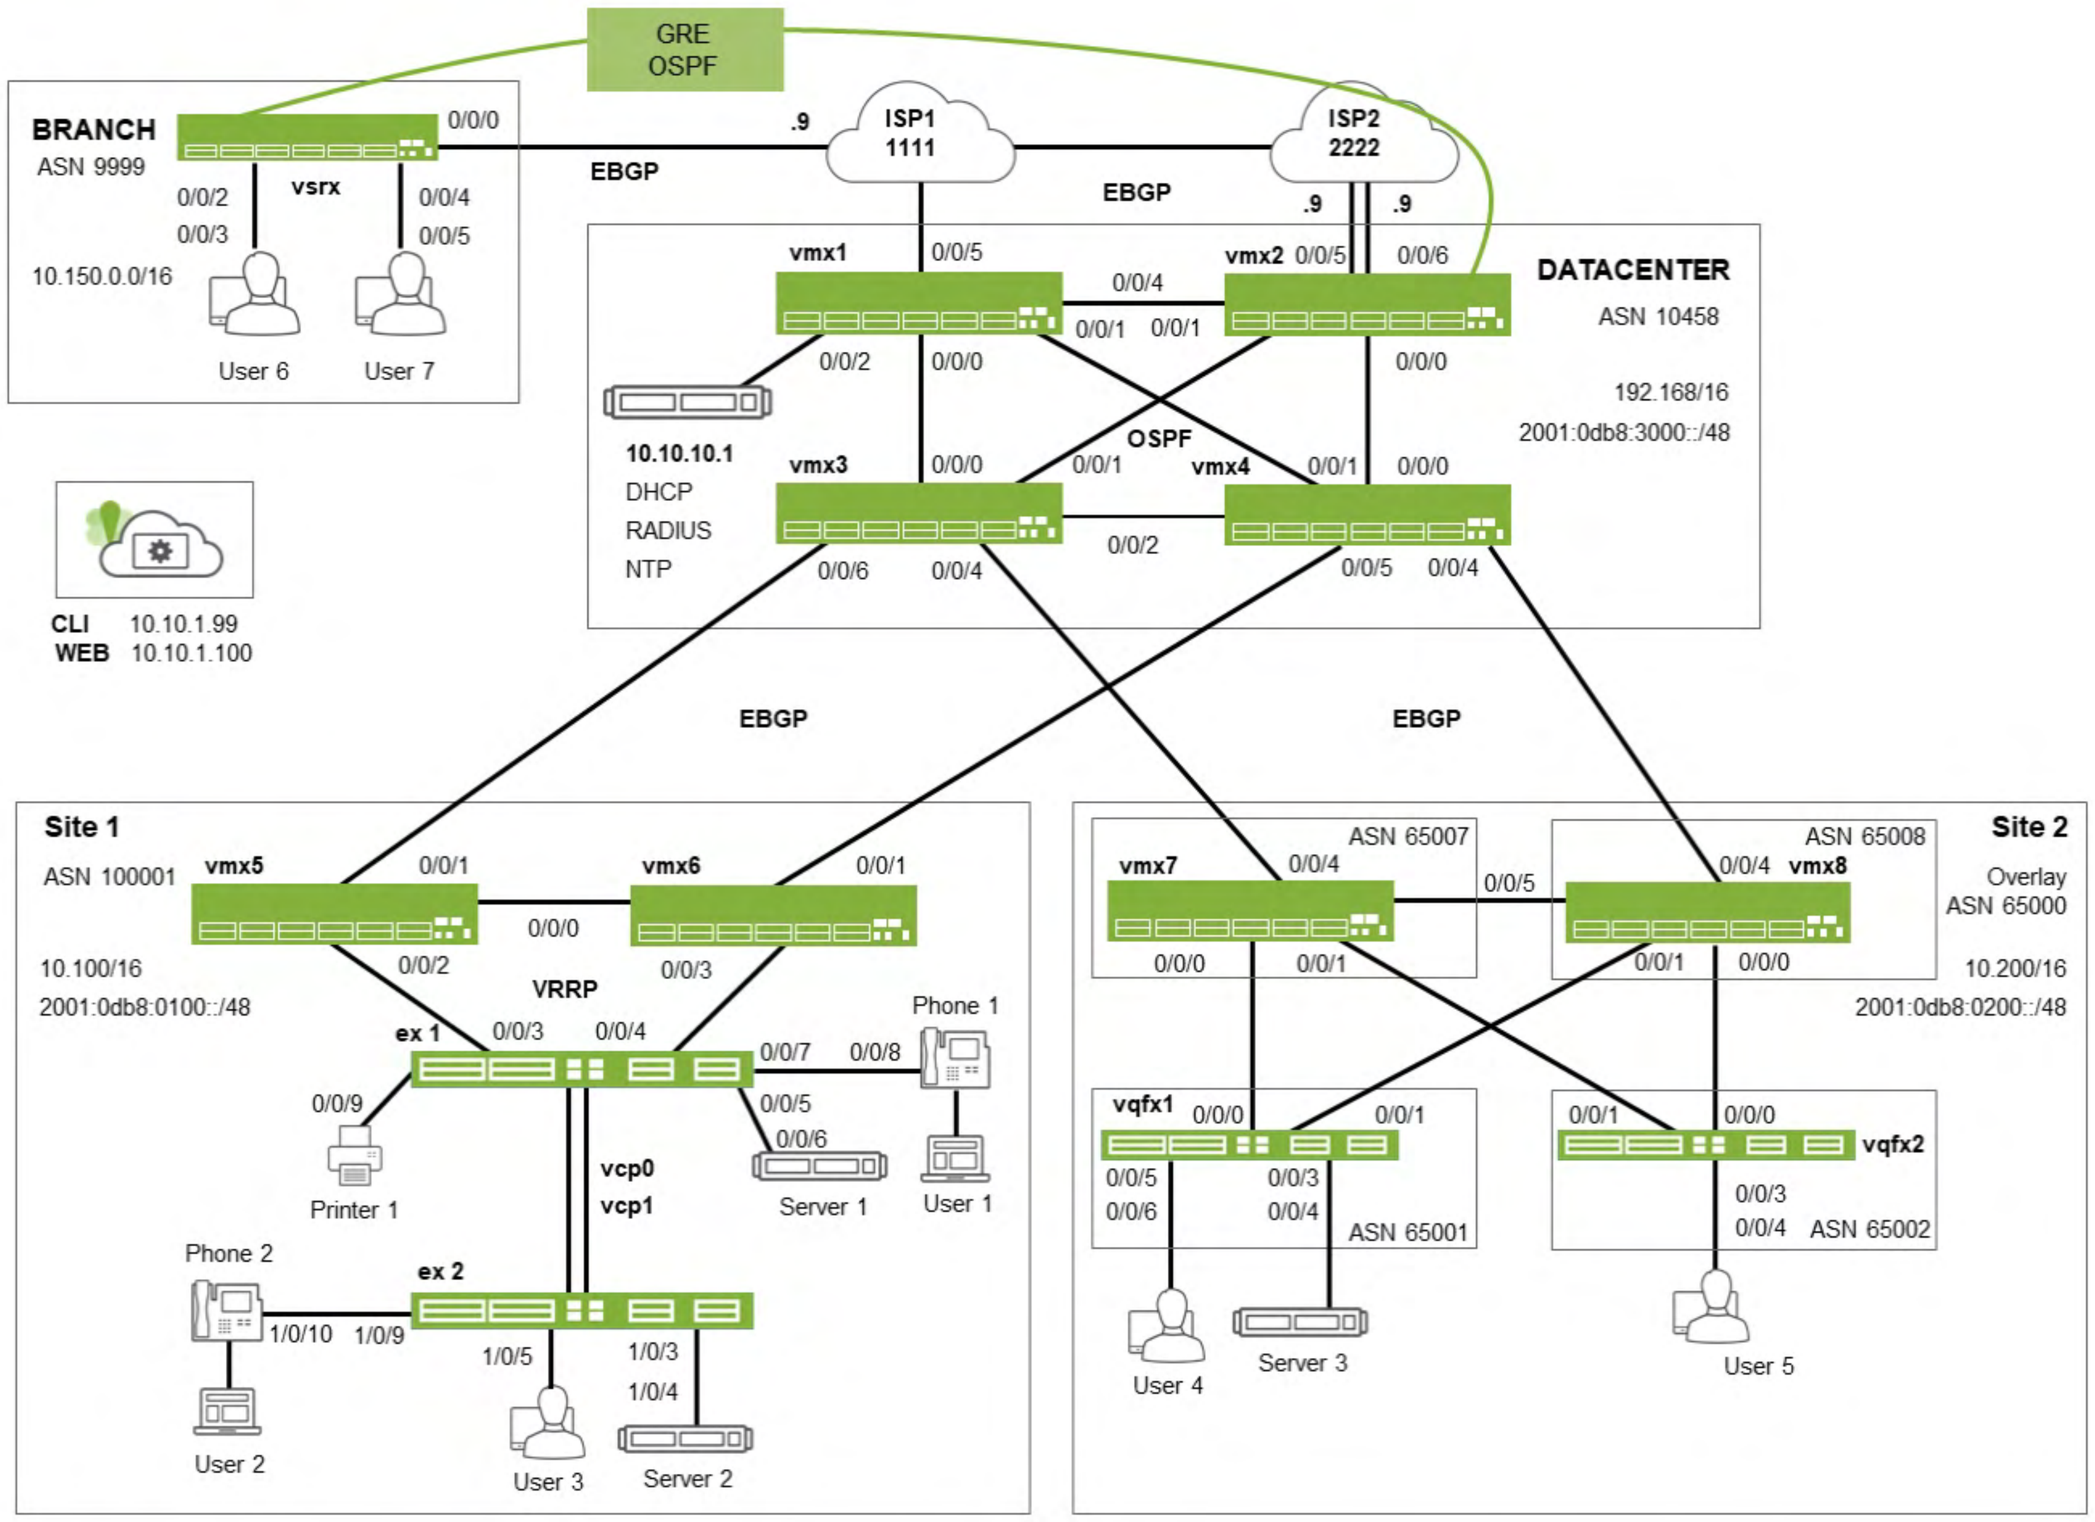

GRE Tunnel

1) Configure a GRE tunnel between the branch office and vMX2 in the data center. Use the IP addressing as shown in the Lab setup Information chapter. Use the loopback interface as GRE tunnel source address.

vMX2:

```
lab@vmx2# show chassis
fpc 0 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    lite-mode;
}
```

```
lab@vmx2# run show interfaces terse | match gr
gr-0/0/10               up    up
```

```
lab@vmx2# show interfaces gr-0/0/10
unit 0 {
    tunnel {
        source 192.168.0.2;
        destination 10.150.0.99;
    }
    family inet {
        address 173.17.99.2/24;
    }
}
```

vSRX:

```
lab@vsrx# show interfaces gr-0/0/0
unit 0 {
    tunnel {
        source 10.150.0.99;
        destination 192.168.0.2;
    }
    family inet {
        address 173.17.99.99/24;
    }
}
```

```
lab@vsrx# show routing-options
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
```

Filter-Based Forwarding

2) Ensure that vMX2 forwards ssh and unencrypted webtraffic received from Site1 and Site2 destinated to the 18.42.206.0/24 network using the ge-0/0/6 interface. If interface ge-0/0/6 is unavailable, then interface ge-0/0/5 must be used. Other traffic should follow the routing requirements as mentioned in the other tasks

vMX2:

```
lab@vmx2# show routing-options
interface-routes {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    rib-group inet interfaces;
}
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 ];
    }
}

lab@vmx2# show firewall filter fbf-v4 | display omit
apply-flags omit;
term 1 {
    from {
        source-address {
            10.200.0.0/16;
            10.100.0.0/16;
        }
        destination-address {
            18.42.206.0/24;
        }
        destination-port [ 80 22 ];
    }
    then {
        routing-instance fbf;
    }
}
term 2 {
    then accept;
}

lab@vmx2# show routing-instances
fbf {
    instance-type forwarding;
    routing-options {
        static {
            route 0.0.0.0/0 {
                next-hop 172.16.3.9;
                qualified-next-hop 172.16.2.9 {
                    preference 6;
                }
            }
        }
    }
}


lab@vmx2# show interfaces ge-0/0/0
unit 0 {
    family inet {
        filter {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            input fbf-v4;
        }
        address 192.168.24.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0024::2/64;
    }
}

[edit]
lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.23.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
}

[edit]
lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
}
```

Verification:

```
lab@vmx2# run show route table fbf.inet

fbf.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:14:41
                    > to 172.16.3.9 via ge-0/0/6.0
                    [Static/6] 00:14:41
                    > to 172.16.2.9 via ge-0/0/5.0
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Commit Script

3) vMX1 has a commit script stored that needs to be enabled. The commit script provides a warning message if vMX1 has IPv4 control plane filtering disabled. The commit script is currently not working and needs to be fixed.

Locate the commit script:

```
lab@vmx1# run start shell
% cd /var/db/scripts/commit/
% ls
Loopback
% exit
```

There is only one script (loopback) located in the well known commit script folder.

Activate the commit script on vMX1:

```
lab@vmx1# show system scripts commit
file loopback;

[edit]
lab@vmx1# commit
/var/db/scripts/commit/loopback:2:(0) Start tag expected, '<' not found
error: error reading stylesheet: loopback
error: 2 errors reported by commit scripts
error: commit script failure

[edit]
lab@vmx1#
```

It clearly shows that there is an issue with the “loopback” commit script as we cannot commit the configuration.

Investigate the loopback script:

```
lab@vmx1# run file show /var/db/scripts/commit/loopback
#!/usr/bin/pythons
from junos import Junos_Configuration
import jcs


def main():
    root = Junos_Configuration

    lo0_interface = root.find("./interfaces/interface[name='lo0']")
    lo0_interface_fwv4 = lo0_interface.find("./unit[name='0']/family/inet/filter/input")
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    lo0_interface_fwv6 =
lo0_interface.find("./unit[name='0']/family/inet6/filter/input")

    if lo0_interface and not(lo0_interface_fwv4):
        jcs.emit_warning("IPv4 control plane protection disabled")


if __name__ == '__main__':
    main()
```

Let start at the top of the script. Junos must understand which interpreter to use to execute the commit script. The loopback script shows that the interpreter “pythons” is used, and that this interpreter should be located in the /usr/bin/ directory. Let’s verify the interpreter.

```
lab@vmx1# run start shell
% ls /usr/bin/pythons
ls: /usr/bin/pythons: No such file or directory
```

The interpreter does not exist. Let’s find out which python interpreter is used for our environment.

```
% which python
/usr/bin/python
```

There is clearly a typo in the interpreter specified in the loopback commit script. The correct interpreter name is: “python” and not “pythons”, which must be corrected:

```
#!/usr/bin/python

from junos import Junos_Configuration
import jcs


def main():
    root = Junos_Configuration

    lo0_interface = root.find("./interfaces/interface[name='lo0']")
    lo0_interface_fwv4 = lo0_interface.find("./unit[name='0']/family/inet/filter/input")
    lo0_interface_fwv6 =
lo0_interface.find("./unit[name='0']/family/inet6/filter/input")

    if lo0_interface and not(lo0_interface_fwv4):
        jcs.emit_warning("IPv4 control plane protection disabled")


if __name__ == '__main__':
    main()
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
% exit
```

Let's verify again if the commit script works.

```
lab@vmx1# commit
warning: [edit system scripts language python] not configured
error: Unable to run Python script 'loopback'
error: 1 error reported by commit scripts
error: commit script failure
```

Based on the script coding format and the interpreter used it should be clear at this stage that this is a Python commit script. For Junos to accept Python as a commit script language you must explicitly configure this. See below:

```
lab@vmx1# show system scripts
commit {
    file loopback;
}
language python;

[edit]
lab@vmx1# commit
commit complete
```

Finally Junos accepts our commit script. Verify if the commit script actually provides a warning if we disable IPv4 control plane protection.

```
[edit]
lab@vmx1# deactivate interfaces lo0.0 family inet filter

[edit]
lab@vmx1# commit
warning: IPv4 control plane protection disabled
commit complete

[edit]
lab@vmx1#
```

It seems that the commit script is working as expected. Do not forget to reapply the control-plane protection to meet the firewall filter task requirements.

```
[edit]
lab@vmx1# rollback 1
load complete

[edit]
lab@vmx1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

On-Box Scripting

4) Ensure that interface ge-0/0/0 on vMX1 is disabled by the Junos OS when it goes down 2 times within 1 minute. Do not use an external script
5) Ensure that event-policies are not allowed to be executed simultaneously.
6) Ensure that Python commit scripts are allowed to be stored off-box. Other commit scripts should only be allowed to be stored on-box.

vMX1:

```
lab@vmx1# show event-options
max-policies 0;
policy jncie {
    events snmp_trap_link_down;
    within 60 {
        trigger after 1;
        events snmp_trap_link_up;
    }
    attributes-match {
        snmp_trap_link_down.interface-name equals snmp_trap_link_up.interface-name;
        "{$$.interface-name}" matches ge-0/0/0;
    }
    then {
        change-configuration {
            commands {
                "set interfaces ge-0/0/0 disable";
            }
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Static Routing and VRRP

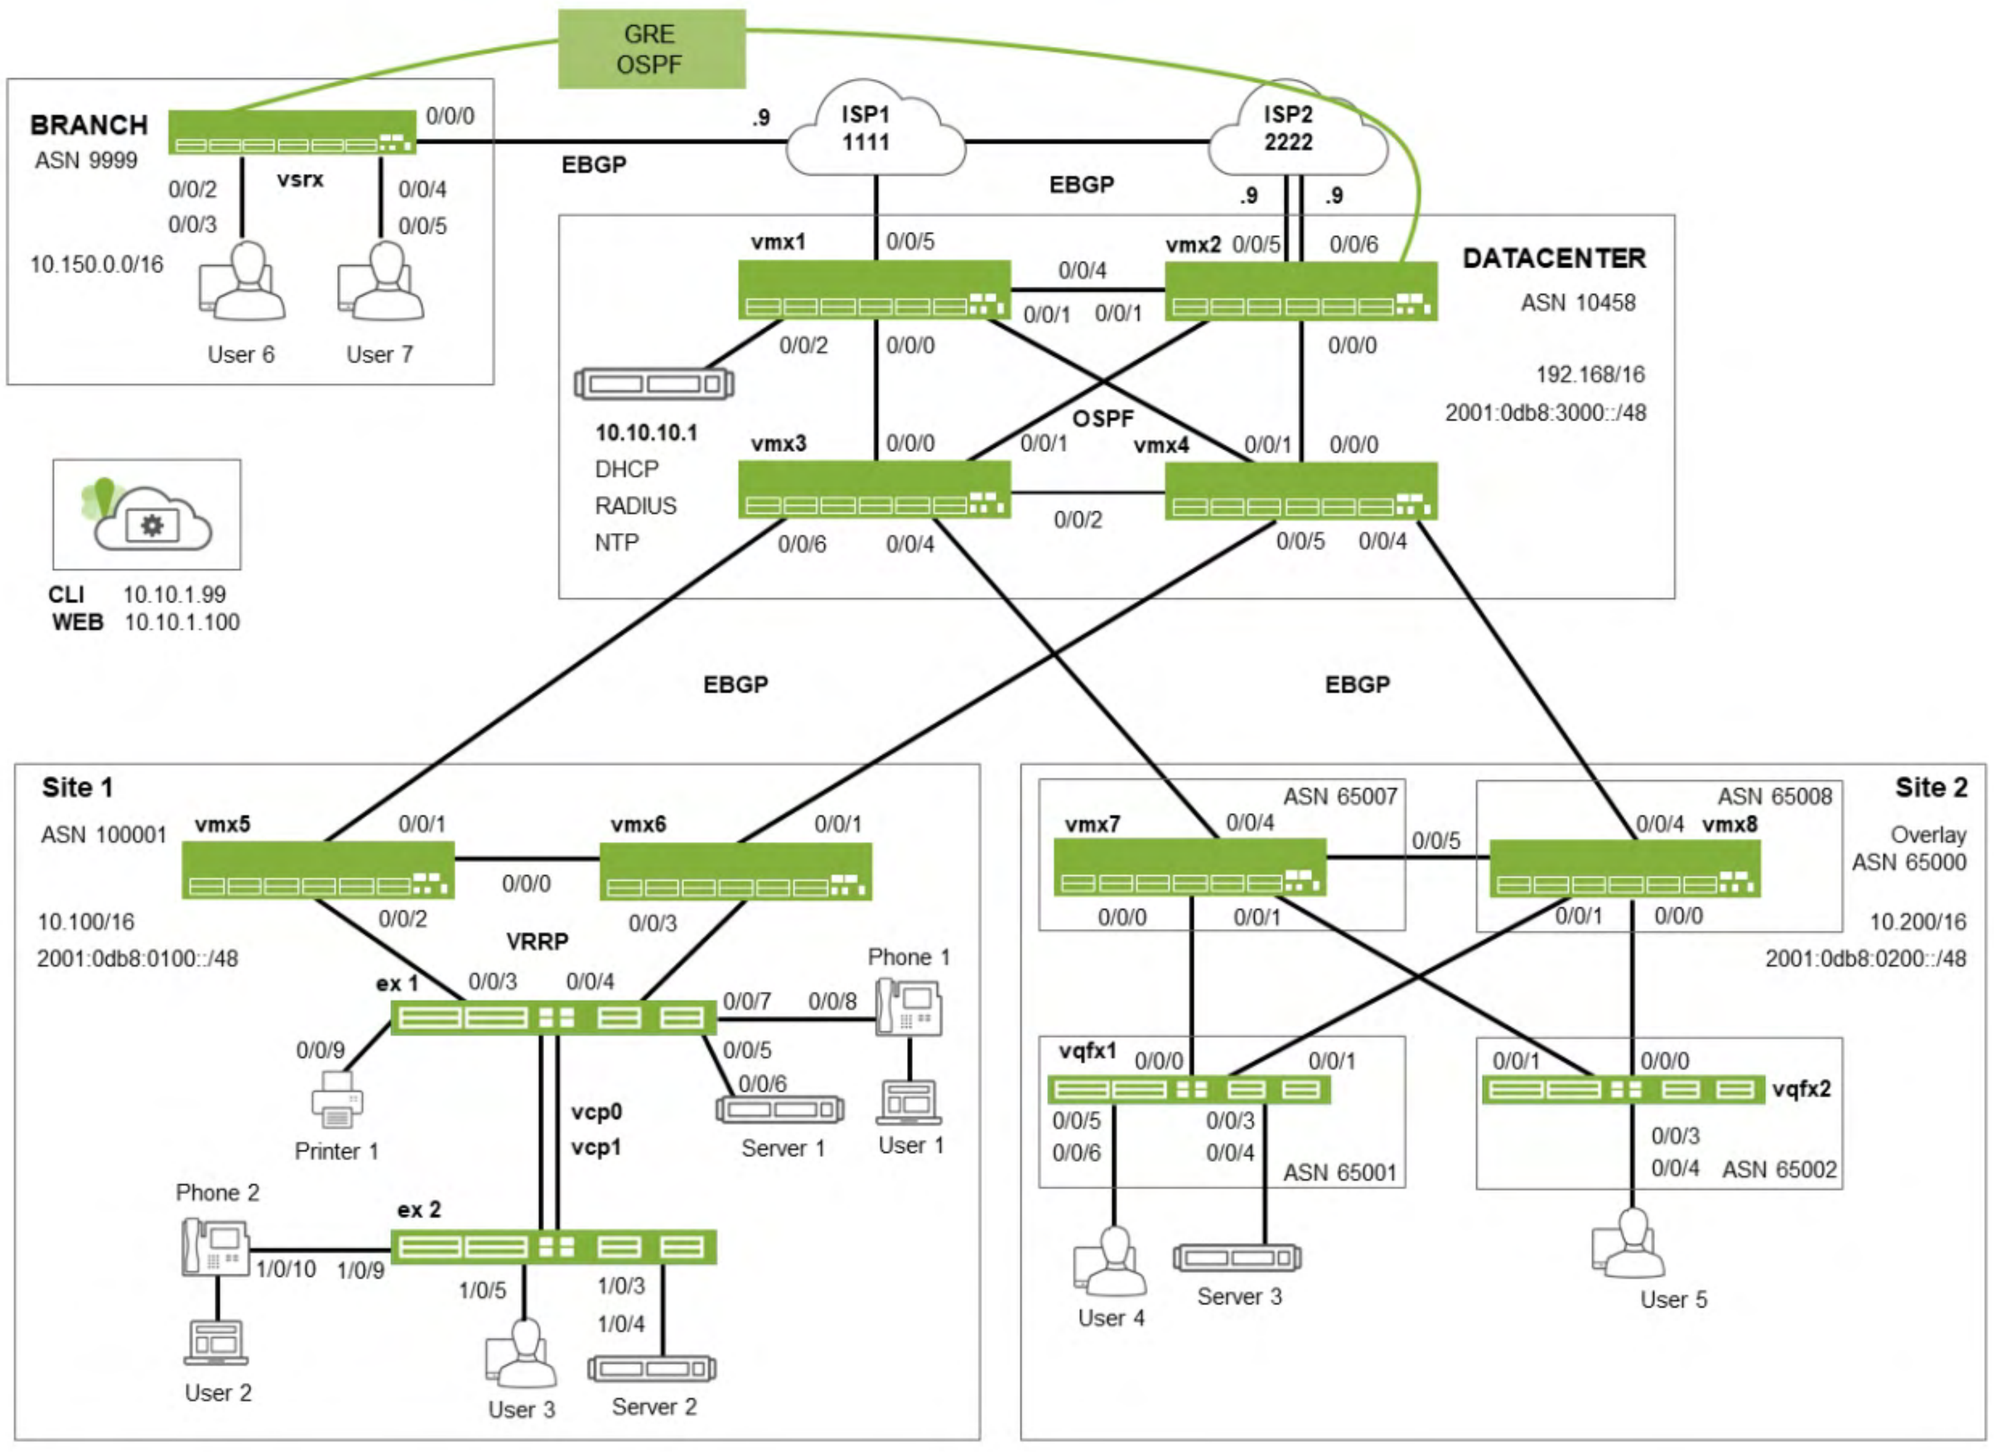

Site 1:

1) Configure static routing over the ge-0/0/0 interface to ensure vMX5 and vMX6 can reach each other loopback address

vMX5:

```
lab@vmx5# show routing-options
rib inet6.0 {
    static {
        route 2001:0db8:0100:0000::6/128 next-hop 2001:0db8:0100:0056::6;
    }
}
static {
    route 10.100.0.6/32 next-hop 10.100.56.6;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

vMX6:

```
lab@vmx6# show routing-options
rib inet6.0 {
    static {
        route 2001:0db8:0100:0000::5/128 next-hop 2001:0db8:0100:0056::5;
    }
}
static {
    route 10.100.0.5/32 next-hop 10.100.56.5;
}
```

Verification:

```
lab@vmx5# run show route 2001:0db8:0100:0000::6

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::6/128*[Static/5] 2d 21:36:08
                    > to 2001:db8:100:56::6 via ge-0/0/0.0

[edit]
lab@vmx5# run show route 10.100.0.6

inet.0: 21 destinations, 30 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.6/32      *[Static/5] 2d 21:36:19
                    > to 10.100.56.6 via ge-0/0/0.0
```

2) Create two static summary routes on vMX5 and vMX6 that encompass the EX1 IRB interfaces.

vMX5 and vMX6:

```
lab@vmx5# show routing-options
rib inet6.0 {
    static {
        route 2001:0db8:0100::/48 next-hop 2001:0db8:0100:0156::1;
    }
}
static {
    route 10.100.0.0/16 next-hop 10.100.156.1;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

vMX5:

```
[edit]
lab@vmx5# run show route 2001:0db8:0100::/48 exact

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::/48  *[Static/5] 2d 22:28:29
                    > to 2001:db8:100:156::1 via ge-0/0/2.0
```

```
[edit]
lab@vmx5# run show route 10.100/16 exact

inet.0: 21 destinations, 30 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/16      *[Static/5] 2d 22:28:51
                    > to 10.100.156.1 via ge-0/0/2.0
```

vMX6:

```
[edit]
lab@vmx6# run show route 2001:0db8:0100::/48 exact

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::/48  *[Static/5] 2d 22:29:03
                    > to 2001:db8:100:156::1 via ge-0/0/3.0

[edit]
lab@vmx6# run show route 10.100/16 exact

inet.0: 20 destinations, 29 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/16      *[Static/5] 2d 22:29:22
                    > to 10.100.156.1 via ge-0/0/3.0
```

3) Configure VRRP on vMX5 and vMX6 to provide a redundant default gateway for EX1. Make sure that the virtual address responds to icmp packets and enable MD5 authentication with password “juniper”

vMX5:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx5# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 10.100.156.5/24 {
            vrrp-group 1 {
                virtual-address 10.100.156.254;
                accept-data;
                authentication-type md5;
                authentication-key "$9$jOkmT69pRhrz3hrev7Nik."; ## SECRET-DATA
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::5/64 {
            vrrp-inet6-group 2 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                accept-data;
            }
        }
    }
}
```

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3
unit 0 {
    family inet {
        address 10.100.156.6/24 {
            vrrp-group 1 {
                virtual-address 10.100.156.254;
                accept-data;
                authentication-type md5;
                authentication-key "$9$jOkmT69pRhrz3hrev7Nik."; ## SECRET-DATA
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 2 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                accept-data;
            }
        }
    }
}
```

4) vMX5 should be the VRRP master for IPv4 traffic. vMX6 should be the master for IPv6 traffic. After recovery of a failure the VRRP master should revert back to its role.

© 2020 Juniper Networks, Inc. All rights reserved.

vMX5:

```
lab@vmx5# show interfaces ge-0/0/2.0 family inet
address 10.100.156.5/24 {
    vrrp-group 1 {
        virtual-address 10.100.156.254;
        priority 150;
        preempt;
        accept-data;
        authentication-type md5;
        authentication-key "$9$jOkmT69pRhrz3hrev7Nik."; ## SECRET-DATA
    }
}
```

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3.0 family inet6
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 2 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                priority 150;
                preempt;
                accept-data;
            }
        }
    }
}
```

5) vMX5 should become VRRP master for IPv6 in case vMX6 interface ge-0/0/1 fails then

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3.0 family inet6
address 2001:0db8:0100:0156::6/64 {
    vrrp-inet6-group 2 {
        virtual-inet6-address 2001:0db8:0100:0156::254;
        priority 150;
        preempt;
        accept-data;
        track {
            interface ge-0/0/1 {
                priority-cost 51;
            }
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

vMX5:

```
lab@vmx5# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/2.0    up              1   master         Active     lcl    10.100.156.5
                                                             vip     10.100.156.254
ge-0/0/2.0    up              2   backup         Active     lcl    2001:db8:100:156::5
                                                             vip
fe80::200:5eff:fe00:263
                                                             vip    2001:db8:100:156::254
```

vMX6:

```
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up              1   backup         Active     lcl    10.100.156.6
                                                             vip     10.100.156.254
ge-0/0/3.0    up              2   master         Active     lcl    2001:db8:100:156::6
                                                             vip
fe80::200:5eff:fe00:263
                                                             vip    2001:db8:100:156::254

edit]
lab@vmx6# deactivate interfaces ge-0/0/1

[edit]
lab@vmx6# commit
commit complete

[edit]
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up              1   backup         Active     lcl    10.100.156.6
                                                             vip     10.100.156.254
ge-0/0/3.0    up              2   backup         Active     lcl    2001:db8:100:156::6
                                                             vip
fe80::200:5eff:fe00:263
                                                             vip    2001:db8:100:156::254

[edit]
lab@vmx6# rollback 1
load complete

[edit]
lab@vmx6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up              1   backup         Active      lcl     10.100.156.6
                                                              vip      10.100.156.254
ge-0/0/3.0    up              2   master         Active      lcl     2001:db8:100:156::6
                                                              vip
fe80::200:5eff:fe00:263
                                                              vip     2001:db8:100:156::254
```

If interface ge-0/0/1 on vMX6 is disabled vMX5 becomes the VRRP master for group 2. vMX6 reclaims VRRP mastership when interface ge-0/0/1 is up again. This means VRRP preempt is working as expected.

6) Create two static default routes on EX1 towards the VRRP virtual address (.254) of vMX5 and vMX6.

EX1:

```
lab@ex1# show routing-options
rib inet6.0 {
    static {
        route ::/0 next-hop 2001:0db8:0100:0156::254;
    }
}
static {
    route 0.0.0.0/0 next-hop 10.100.156.254;
}
```

Verification:

```
lab@ex1# run show route protocol static table inet.0

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

0.0.0.0/0          *[Static/5] 22:00:49
                    >  to 10.100.156.254 via irb.156

lab@ex1# run show route protocol static table inet6.0

inet6.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0               *[Static/5] 22:00:41
                    >  to 2001:db8:100:156::254 via irb.156
```

Configurations for Task Completion:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX1:

```
set system scripts commit file loopback
set system scripts language python
set event-options max-policies 0
set event-options policy jncie events snmp_trap_link_down
set event-options policy jncie within 60 trigger after
set event-options policy jncie within 60 trigger 1
set event-options policy jncie within 60 events snmp_trap_link_up
set event-options policy jncie attributes-match snmp_trap_link_down.interface-name
equals snmp_trap_link_up.interface-name
set event-options policy jncie attributes-match "{$$.interface-name}" matches ge-0/0/0
set event-options policy jncie then change-configuration commands "set interfaces ge-
0/0/0 disable"
```

vMX2:

```
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set interfaces gr-0/0/10 unit 0 tunnel source 192.168.0.2
set interfaces gr-0/0/10 unit 0 tunnel destination 10.150.0.99
set interfaces gr-0/0/10 unit 0 family inet address 173.17.99.2/24
set routing-options interface-routes rib-group inet interfaces
set routing-options rib-groups interfaces import-rib [ inet.0 fbf.inet.0 ]
set firewall filter fbf-v4 apply-flags omit
set firewall filter fbf-v4 term 1 from source-address 10.200.0.0/16
set firewall filter fbf-v4 term 1 from source-address 10.100.0.0/16
set firewall filter fbf-v4 term 1 from destination-address 18.42.206.0/24
set firewall filter fbf-v4 term 1 from destination-port [ 80 22 ]
set firewall filter fbf-v4 term 1 then routing-instance fbf
set firewall filter fbf-v4 term 2 then accept
set routing-instances fbf instance-type forwarding
set routing-instances fbf routing-options static route 0.0.0.0/0 next-hop 172.16.3.9
set routing-instances fbf routing-options static route 0.0.0.0/0 qualified-next-hop
172.16.2.9 preference 6
set interfaces ge-0/0/0 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/1 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/4 unit 0 family inet filter input fbf-v4
```

vMX5:

```
set routing-options rib inet6.0 static route 2001:0db8:0100:0000::6/128 next-hop
2001:0db8:0100:0056::6
set routing-options static route 10.100.0.6/32 next-hop 10.100.56.6
set routing-options rib inet6.0 static route 2001:0db8:0100::/48 next-hop
2001:0db8:0100:0156::1
set routing-options static route 10.100.0.0/16 next-hop 10.100.156.1
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1 virtual-
address 10.100.156.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1 accept-
data
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1
authentication-type md5
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1
authentication-key "$9$jOkmT69pRhrz3hrev7Nik."
set interfaces ge-0/0/2 unit 0 family inet6 address 2001:0db8:0100:0156::5/64 vrrp-
inet6-group 2 virtual-inet6-address 2001:0db8:0100:0156::254
set interfaces ge-0/0/2 unit 0 family inet6 address 2001:0db8:0100:0156::5/64 vrrp-
inet6-group 2 accept-data
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1 priority
150
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 1 preempt
```

vMX6:

```
set routing-options rib inet6.0 static route 2001:0db8:0100:0000::5/128 next-hop
2001:0db8:0100:0056::5
set routing-options static route 10.100.0.5/32 next-hop 10.100.56.5
set routing-options rib inet6.0 static route 2001:0db8:0100::/48 next-hop
2001:0db8:0100:0156::1
set routing-options static route 10.100.0.0/16 next-hop 10.100.156.1
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 1 virtual-
address 10.100.156.254
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 1 accept-
data
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 1
authentication-type md5
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 1
authentication-key "$9$jOkmT69pRhrz3hrev7Nik."
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 2 virtual-inet6-address 2001:0db8:0100:0156::254
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 2 accept-data
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 2 priority 150
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 2 preempt
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 2 track interface ge-0/0/1 priority-cost 51
```

vSRX:

```
set interfaces gr-0/0/0 unit 0 tunnel source 10.150.0.99
set interfaces gr-0/0/0 unit 0 tunnel destination 192.168.0.2
set interfaces gr-0/0/0 unit 0 family inet address 173.17.99.99/24
set routing-options static route 192.168.0.2/32 next-hop 10.1.1.1
```

EX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 static route ::/0 next-hop 2001:0db8:0100:0156::254
set routing-options static route 0.0.0.0/0 next-hop 10.100.156.254
```

Task 4: IGP

This task focuses on the implementation of various IGPs like OSPFv2 and OSPv3.

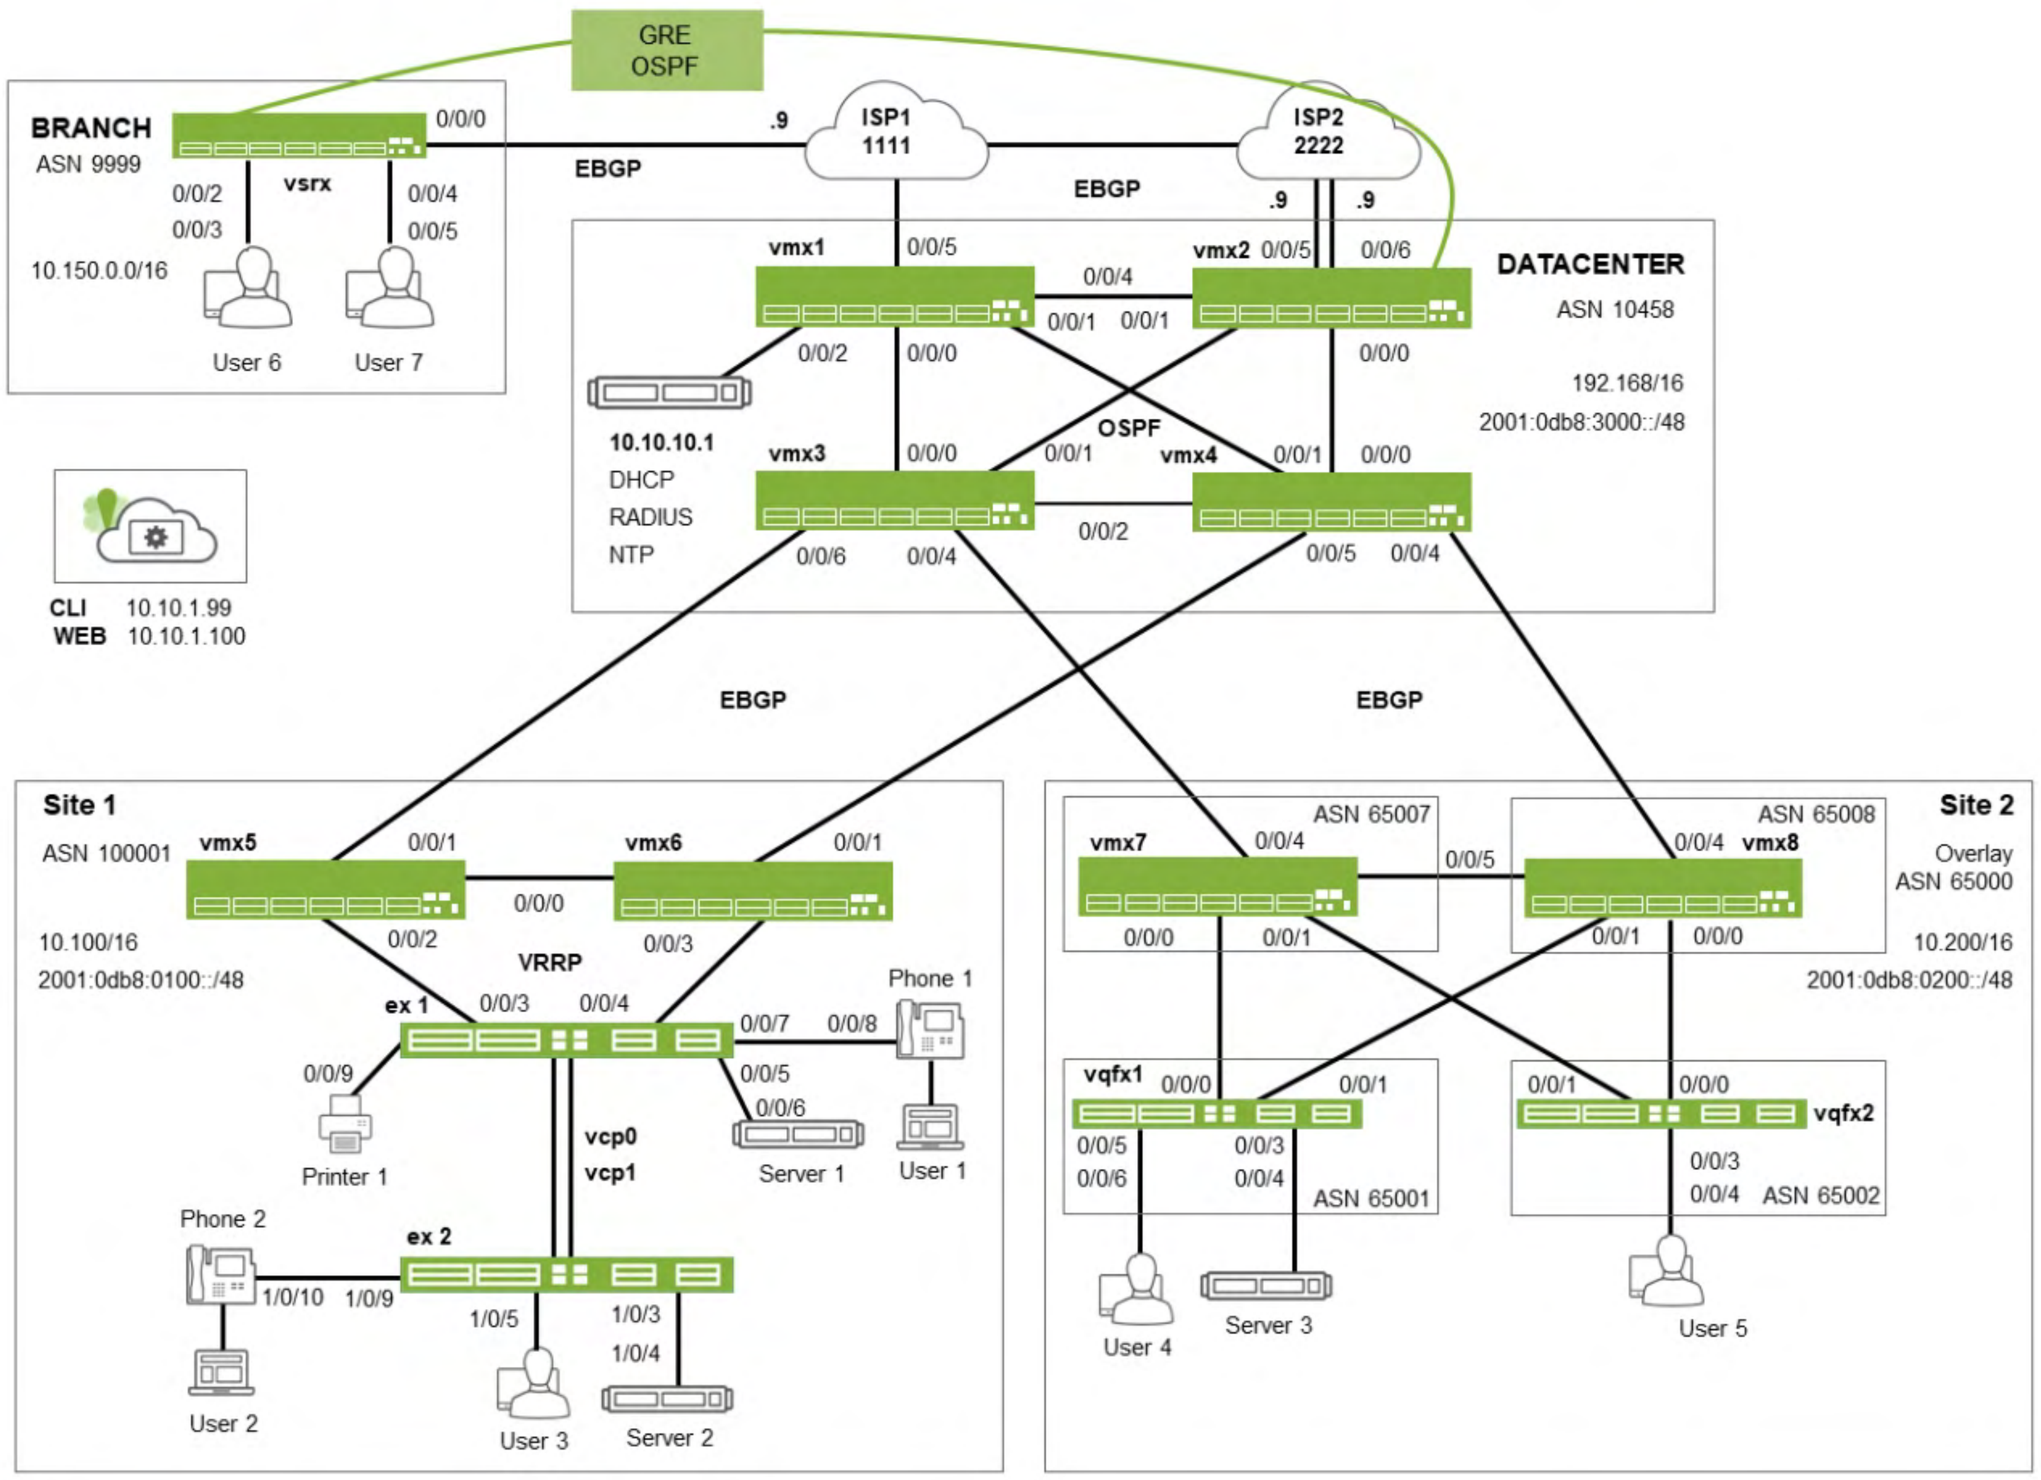

OSPFv3

Enable OSPFv3 in the data center for IPv6 routing.

Data Center:

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX2	lo0	0.0.0.0
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX3	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1
vMX4	lo0	0.0.0.0
	ge-0/0/2	0.0.0.1
	ge-0/0/0	0.0.0.1
	ge-0/0/1	0.0.0.1

Table 40 OSPF Information Data Center

1) Configure OSPFv3 in the data center using the information as shown in the tables above.
2) Ensure no LSA type 0x2002 are found in the OSPFv3 databases in Site 1.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four missed hello packets.

vMX1:

```
lab@vmx1# show protocols ospf3
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
    interface ge-0/0/4.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/2.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        interface-type p2p;
        hello-interval 1;
    }
}
```

vMX2:

```
lab@vmx2# show protocols ospf3
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
    interface ge-0/0/4.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
area 0.0.0.1 {
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
```

vMX3:

```
lab@vmx3# show protocols ospf3
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

vMX4:

```
lab@vmx4# show protocols ospf3
area 0.0.0.0 {
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.1 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
```

Verification:

vMX1:

```
lab@vmx1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
192.168.0.2     ge-0/0/4.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3b9f
192.168.0.3     ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3aaa
192.168.0.4     ge-0/0/1.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:12b5
```

4) Ensure that each vMX devices in the data center learn each other's IPv6 loopback address.

The loopback addresses of the vMX devices participate in the OSPF backbone area (0), while the links between the vMXs participate in area 1. This is an invalid OSPF design as the OSPF backbone should be contiguous. Make the OSPF backbone area contiguous by configuring virtual links between the vMX devices.

vMX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show protocols ospf3 area 0
virtual-link neighbor-id 192.168.0.3 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.4 transit-area 0.0.0.1;
interface lo0.0 {
    passive;
}
interface ge-0/0/4.0 {
    interface-type p2p;
    hello-interval 1;
}
interface ge-0/0/2.0 {
    passive;
}
```

vMX2:

```
lab@vmx2# show protocols ospf3 area 0
virtual-link neighbor-id 192.168.0.3 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.4 transit-area 0.0.0.1;
interface lo0.0 {
    passive;
}
interface ge-0/0/4.0 {
    interface-type p2p;
    hello-interval 1;
}
```

vMX3:

```
lab@vmx3# show protocols ospf3 area 0
virtual-link neighbor-id 192.168.0.1 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.2 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.4 transit-area 0.0.0.1;
interface lo0.0 {
    passive;
}
```

vMX4:

```
lab@vmx4# show protocols ospf3 area 0
virtual-link neighbor-id 192.168.0.1 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.2 transit-area 0.0.0.1;
virtual-link neighbor-id 192.168.0.3 transit-area 0.0.0.1;
interface lo0.0 {
    passive;
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX1:

```
lab@vmx1# run show ospf3 neighbor
ID              Interface              State     Pri   Dead
192.168.0.2     ge-0/0/4.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3b9f
192.168.0.3     vl-192.168.0.3         Full        0     35
  Neighbor-address 2001:db8:3000:13::3
192.168.0.4     vl-192.168.0.4         Full        0     38
  Neighbor-address 2001:db8:3000:14::4
192.168.0.3     ge-0/0/0.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3aaa
192.168.0.4     ge-0/0/1.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:12b5

lab@vmx1# run show route protocol ospf3 | match 128
2001:db8:3000::2/128
2001:db8:3000::3/128
2001:db8:3000::4/128
2001:db8:3000:23::2/128
2001:db8:3000:34::3/128
2001:db8:3000:34::4/128
...
```

vMX2:

```
lab@vmx2# run show ospf3 neighbor
ID              Interface              State     Pri   Dead
192.168.0.1     ge-0/0/4.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4466
192.168.0.3     vl-192.168.0.3         Full        0     32
  Neighbor-address 2001:db8:3000:23::3
192.168.0.4     vl-192.168.0.4         Full        0     31
  Neighbor-address 2001:db8:3000:24::4
192.168.0.4     ge-0/0/0.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:51c6
192.168.0.3     ge-0/0/1.0             Full      128      3
  Neighbor-address fe80::250:56ff:fea9:32cf

lab@vmx2# run show route protocol ospf3 | match 128
2001:db8:3000::1/128
2001:db8:3000::3/128
2001:db8:3000::4/128
2001:db8:3000:14::1/128
2001:db8:3000:34::3/128
2001:db8:3000:34::4/128
...
```

vMX3:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx3# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
192.168.0.1     vl-192.168.0.1        Full        0     32
  Neighbor-address 2001:db8:3000:13::1
192.168.0.2     vl-192.168.0.2        Full        0     35
  Neighbor-address 2001:db8:3000:23::2
192.168.0.4     vl-192.168.0.4        Full        0     36
  Neighbor-address 2001:db8:3000:34::4
192.168.0.1     ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3ed3
192.168.0.2     ge-0/0/1.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:35aa
192.168.0.4     ge-0/0/2.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:36

lab@vmx3# run show route protocol ospf3 | match 128
2001:db8:3000::1/128
2001:db8:3000::2/128
2001:db8:3000::4/128
2001:db8:3000:14::1/128
2001:db8:3000:23::2/128
2001:db8:3000:34::4/128
```

vMX4:

```
lab@vmx4# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
192.168.0.1     vl-192.168.0.1        Full        0     30
  Neighbor-address 2001:db8:3000:14::1
192.168.0.2     vl-192.168.0.2        Full        0     33
  Neighbor-address 2001:db8:3000:24::2
192.168.0.3     vl-192.168.0.3        Full        0     35
  Neighbor-address 2001:db8:3000:34::3
192.168.0.2     ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3944
192.168.0.1     ge-0/0/1.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:427f
192.168.0.3     ge-0/0/2.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:6e2a

lab@vmx4# run show route protocol ospf3 | match 128
2001:db8:3000::1/128
2001:db8:3000::2/128
2001:db8:3000::3/128
2001:db8:3000:14::1/128
2001:db8:3000:23::2/128
2001:db8:3000:34::3/128
```


© 2020 Juniper Networks, Inc. All rights reserved.

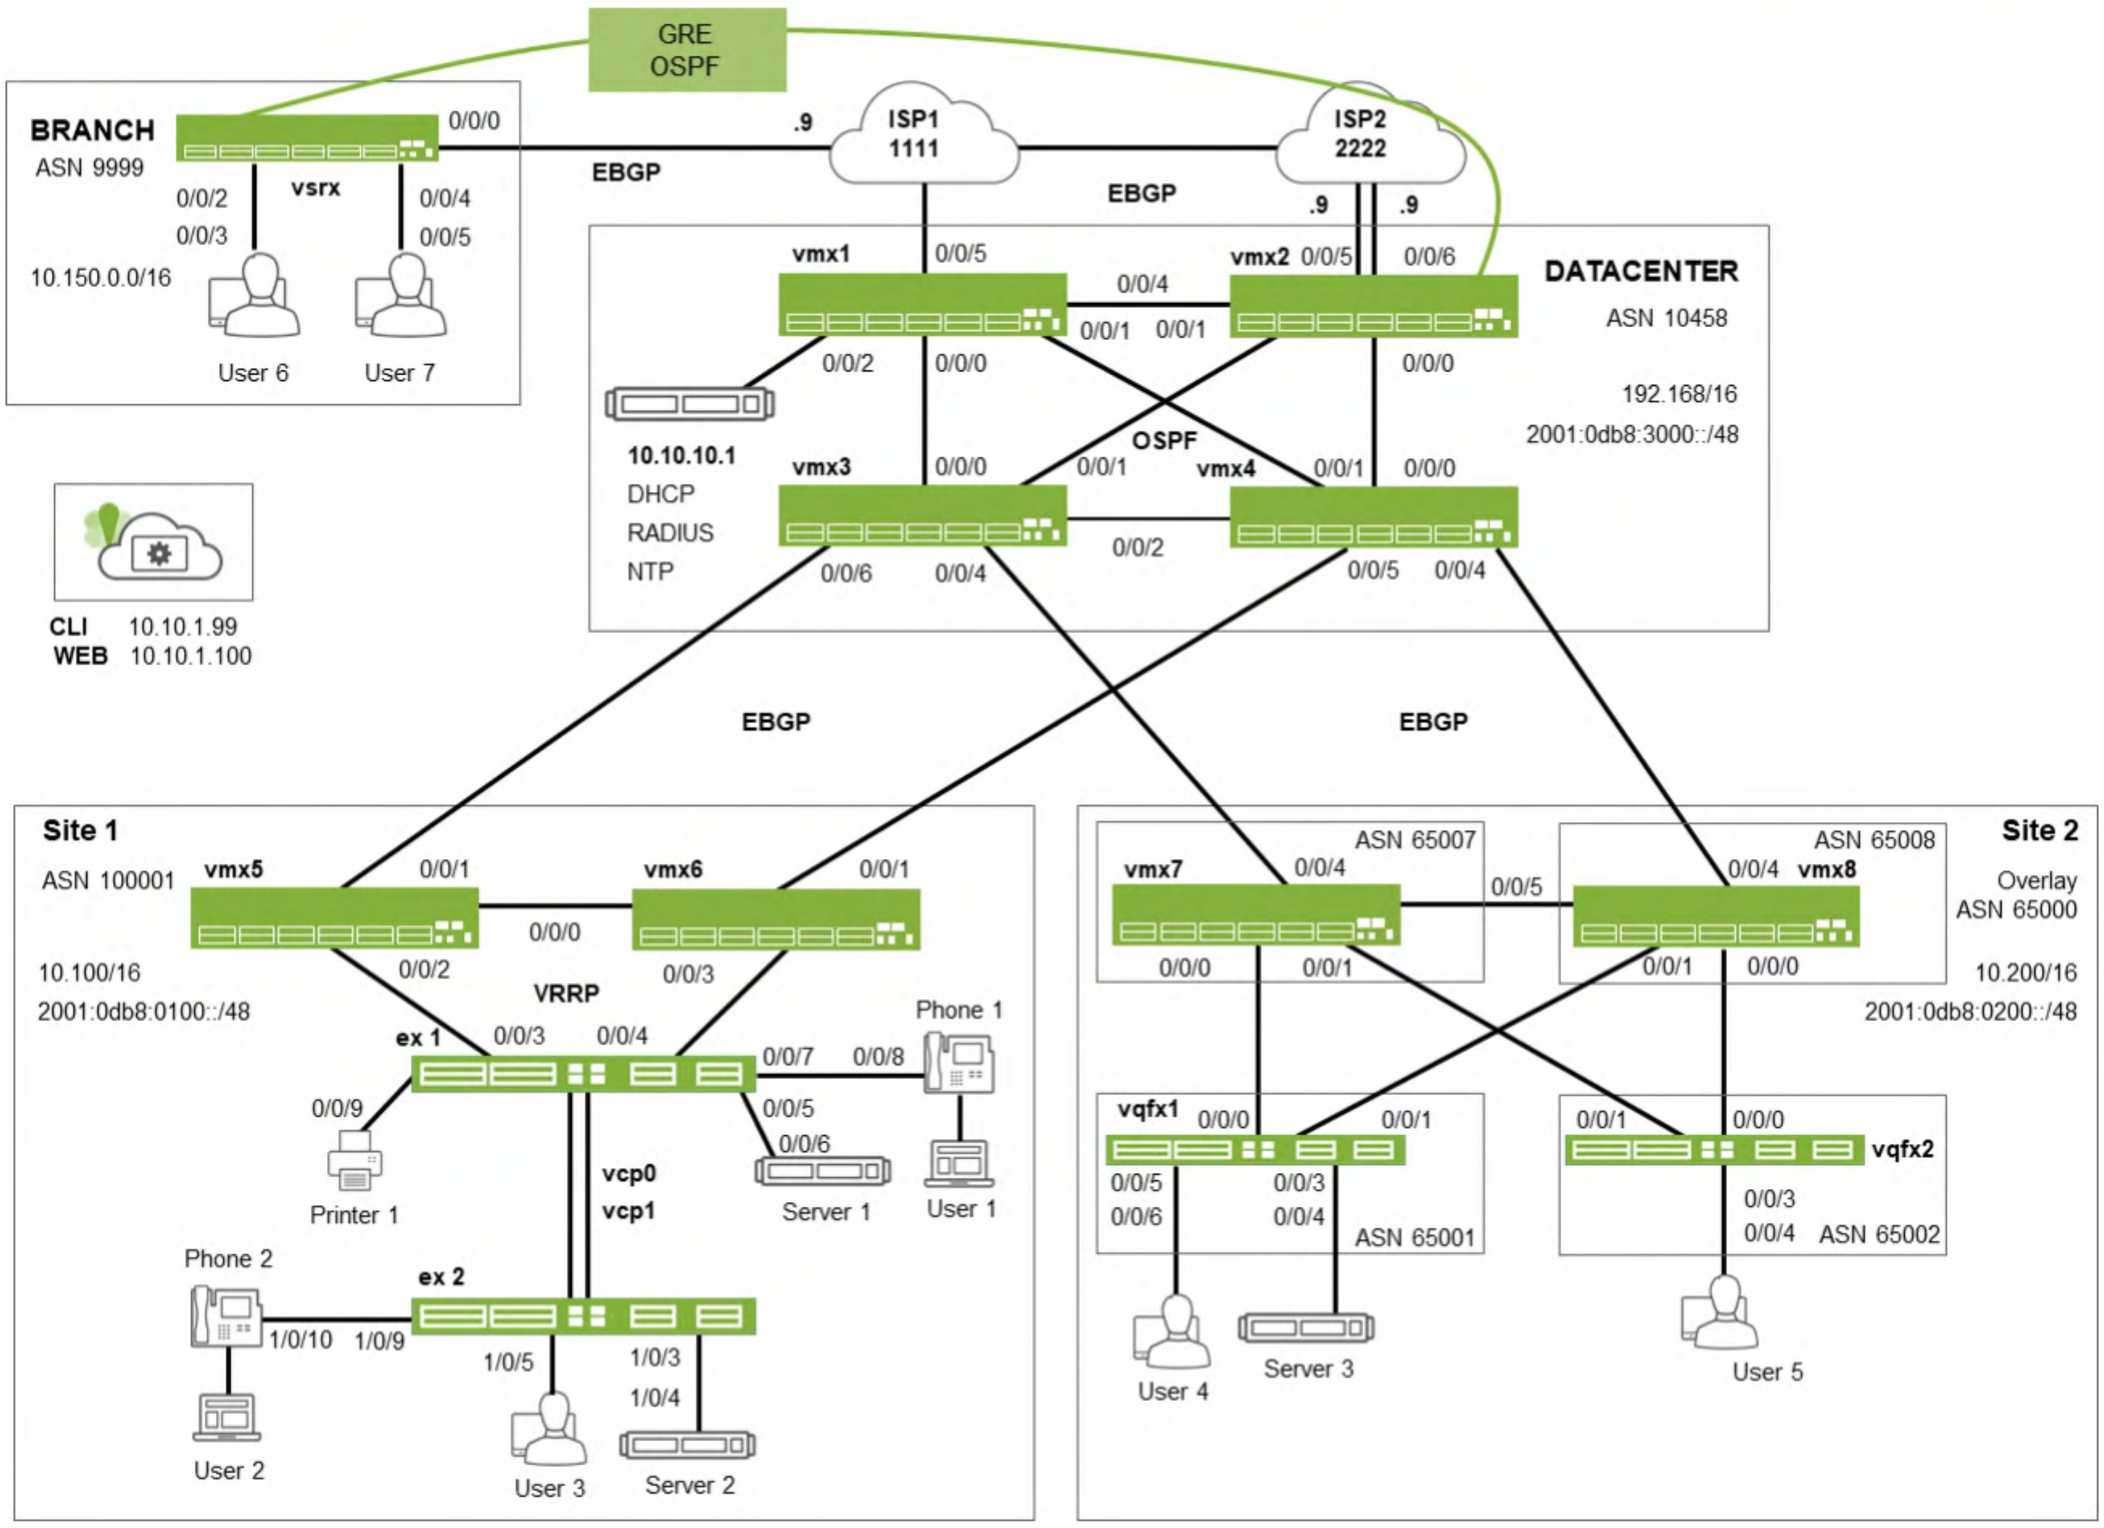

OSPFv2

In this section you will configure an OSPFv2 for IPv4 routing in the data center and over the GRE tunnel between the data center and the branch office.

Device	Interface	Routing instance	area
vMX2	GRE interface	ospf-branch	0.0.0.0
vSRX	GRE interface	default	0.0.0.0

Table 41 OSPF information for GRE tunnel

5) Enable OSPFv2 over the GRE interface as shown in the table above.
6) Ensure the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.

© 2020 Juniper Networks, Inc. All rights reserved.

7) Ensure that vSRX OSPF database contains a default route and the other OSPFv2 destinations received from vMX2. Configure OSPF on vSRX that the default route is the only external LSA that is installed in its routing table.

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
vMX 2	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
vMX 3	lo0	0.0.0.1
	ge-0/0/2	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
vMX 4	lo0	0.0.0.1
	ge-0/0/2	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0

Enable OSPFv2 for IPv4 routing on the vMX devices in the data center. Use the information as shown in the table above.

8) Ensure that OSPFv2 calculates backup paths for (most) of the LSAs found in the OSPF database.
9) Ensure that Gigabit Ethernet interfaces have an OSPF cost of 100. Do not make any changes under the OSPF interface hierarchy.

vMX1:

```
lab@vmx1# show protocols ospf
spf-options {
    holddown 3000;
    rapid-runs 5;
}
reference-bandwidth 100g;
area 0.0.0.0 {
    interface ge-0/0/4.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
interface ge-0/0/2.0 {
    passive;
}
interface ge-0/0/1.0 {
    interface-type p2p;
    link-protection;
    hello-interval 1;
}
interface ge-0/0/0.0 {
    interface-type p2p;
    link-protection;
    hello-interval 1;
        flood-reduction;
    }
}
area 0.0.0.1 {
    interface lo0.0;
}
```

vMX2:

```
lab@vmx2# show routing-instances
ospf-branch {
    interface gr-0/0/10.0;
    routing-options {
        interface-routes {
            rib-group inet branch-to-inet;
        }
        static {
            route 0.0.0.0/0 next-table inet.0;
        }
    }
    protocols {
        ospf {
            rib-group branch-to-inet;
            export 2ospfbranch;
            area 0.0.0.0 {
                interface gr-0/0/10.0 {
                    hello-interval 1;
                }
            }
        }
    }
}

lab@vmx2# show policy-options policy-statement 2ospfbranch
term 1 {
    from protocol [ ospf direct static ];
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    then accept;
}
term 2 {
    then reject;
}


lab@vmx2# show routing-options
interface-routes {
    rib-group inet interfaces;
}
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 ospf-branch.inet.0 ];
    }
    branch-to-inet {
        import-rib [ ospf-branch.inet.0 inet.0 ];
    }
    inet-to-branch {
        import-rib [ inet.0 ospf-branch.inet.0 ];
    }
}

lab@vmx2# show protocols ospf
spf-options {
    holddown 3000;
    rapid-runs 5;
}
rib-group inet-to-branch;
reference-bandwidth 100g;
area 0.0.0.0 {
    interface ge-0/0/4.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
        flood-reduction;
    }
}
area 0.0.0.1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    interface lo0.0;
}
```

vMX3:

```
lab@vmx3# show protocols ospf
reference-bandwidth 100g;
area 0.0.0.0 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
}
area 0.0.0.1 {
    interface lo0.0;
}
```

vMX4:

```
lab@vmx4# show protocols ospf
reference-bandwidth 100g;
area 0.0.0.0 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        link-protection;
        hello-interval 1;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
area 0.0.0.1 {
    interface lo0.0;
}
```

vSRX:

```
lab@vsrx# show protocols ospf
export 2ospf;
import ospfimport;
area 0.0.0.0 {
    interface gr-0/0/0.0 {
        hello-interval 1;
    }
}

lab@vsrx# show policy-options
policy-statement 2ospf {
    term 1 {
        from interface irb.99;
        then {
            external {
                type 1;
            }
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
policy-statement ospfimport {
    term 1 {
        from {
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
```

Verifiy OSPF adjacencies:

vMX1:

```
lab@vmx1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
192.168.13.3     ge-0/0/0.0             Full      192.168.0.3      128     3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.14.4     ge-0/0/1.0          Full       192.168.0.4      128     3
192.168.12.2     ge-0/0/4.0          Full       192.168.0.2      128     3
```

vMX2:

```
lab@vmx2# run show ospf neighbor
Address          Interface           State      ID               Pri   Dead
192.168.24.4     ge-0/0/0.0          Full       192.168.0.4      128     3
192.168.23.3     ge-0/0/1.0          Full       192.168.0.3      128     3
192.168.12.1     ge-0/0/4.0          Full       192.168.0.1      128     3
```

vMX3:

```
lab@vmx3# run show ospf neighbor
Address          Interface           State      ID               Pri   Dead
192.168.13.1     ge-0/0/0.0          Full       192.168.0.1      128     3
192.168.23.2     ge-0/0/1.0          Full       192.168.0.2      128     3
192.168.34.4     ge-0/0/2.0          Full       192.168.0.4      128     3
```

vMX4:

```
lab@vmx4# run show ospf neighbor
Address          Interface           State      ID               Pri   Dead
192.168.24.2     ge-0/0/0.0          Full       192.168.0.2      128     3
192.168.14.1     ge-0/0/1.0          Full       192.168.0.1      128     3
192.168.34.3     ge-0/0/2.0          Full       192.168.0.3      128     3
```

Verify routing between inet.0 and ospf-branch routing instance:

Note: You should wait until you have completed the BGP routing tasks in the next section before verifying the information between the data center and the branch site. We do not currently have BGP configured through our ISPs. The following commands were gathered after Task 5 was completed.

vMX2:

```
lab@vmx2# run show route table ospf-branch.inet.0

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

0.0.0.0/0          *[Static/5] 1d 21:49:16
                      to table inet.0
10.10.10.0/24      *[OSPF/10] 02:06:16, metric 200
                    > to 192.168.12.1 via ge-0/0/4.0
                      to 192.168.23.3 via ge-0/0/1.0
10.150.99.0/24     *[OSPF/150] 02:06:15, metric 1, tag 0
                    > via gr-0/0/10.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.16.2.0/24      *[Direct/0] 02:06:16
                    > via ge-0/0/5.0
172.16.2.2/32      *[Local/0] 02:06:16
                      Local via ge-0/0/5.0
172.16.3.0/24      *[Direct/0] 02:06:16
                    > via ge-0/0/6.0
172.16.3.2/32      *[Local/0] 02:06:16
                      Local via ge-0/0/6.0
173.17.99.0/24     *[Direct/0] 02:06:16
                    > via gr-0/0/10.0
                    [OSPF/10] 02:06:16, metric 1
                    > via gr-0/0/10.0
173.17.99.2/32     *[Local/0] 02:06:16
                      Local via gr-0/0/10.0
192.168.0.1/32     *[OSPF/10] 02:06:16, metric 100
                    > to 192.168.12.1 via ge-0/0/4.0
                      to 192.168.23.3 via ge-0/0/1.0
192.168.0.2/32     *[Direct/0] 3d 01:36:51
                    > via lo0.0
192.168.0.3/32     *[OSPF/10] 02:06:16, metric 100
                    > to 192.168.23.3 via ge-0/0/1.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.0.4/32     *[OSPF/10] 02:06:16, metric 100
                    > to 192.168.24.4 via ge-0/0/0.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.12.0/24    *[Direct/0] 02:06:16
                    > via ge-0/0/4.0
192.168.12.2/32    *[Local/0] 02:06:16
                      Local via ge-0/0/4.0
192.168.13.0/24    *[OSPF/10] 02:06:16, metric 200
                    > to 192.168.23.3 via ge-0/0/1.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.14.0/24    *[OSPF/10] 02:06:16, metric 200
                    > to 192.168.24.4 via ge-0/0/0.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.23.0/24    *[Direct/0] 02:06:16
                    > via ge-0/0/1.0
192.168.23.2/32    *[Local/0] 02:06:16
                      Local via ge-0/0/1.0
192.168.24.0/24    *[Direct/0] 02:06:16
                    > via ge-0/0/0.0
192.168.24.2/32    *[Local/0] 02:06:16
                      Local via ge-0/0/0.0
192.168.34.0/24    *[OSPF/10] 02:06:16, metric 200
                      to 192.168.24.4 via ge-0/0/0.0
                    > to 192.168.23.3 via ge-0/0/1.0
224.0.0.5/32       *[OSPF/10] 3d 01:55:49, metric 1
                      MultiRecv
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# run show route protocol ospf

inet.0: 256 destinations, 267 routes (256 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/24      *[OSPF/10] 02:07:54, metric 200
                    > to 192.168.12.1 via ge-0/0/4.0
                      to 192.168.23.3 via ge-0/0/1.0
10.150.99.0/24     *[OSPF/150] 02:07:53, metric 1, tag 0
                    > via gr-0/0/10.0
173.17.99.0/24      [OSPF/10] 02:07:54, metric 1
                    > via gr-0/0/10.0
```

vSRX:

```
lab@vsrx# run show route protocol ospf

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 02:13:53, metric 0, tag 0
                    >  via gr-0/0/0.0
173.17.99.0/24      [OSPF/10] 1w0d 00:59:29, metric 1
                    >  via gr-0/0/0.0
224.0.0.5/32       *[OSPF/10] 1w0d 01:00:07, metric 1
                       MultiRecv

..

lab@vsrx# run show ospf database detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *10.150.0.99      10.150.0.99      0x8000013a  2045  0x22 0x1ef4  48
  bits 0x2, link count 2
  id 173.17.99.2, data 173.17.99.99, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 173.17.99.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 1
  Topology default (ID 0)
    Type: PointToPoint, Node ID: 173.17.99.2
      Metric: 1, Bidirectional
Router   173.17.99.2      173.17.99.2      0x800000a3   447  0x22 0x7477  48
  bits 0x2, link count 2
  id 10.150.0.99, data 173.17.99.2, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 173.17.99.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 1
  Topology default (ID 0)
    Type: PointToPoint, Node ID: 10.150.0.99
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Metric: 1, Bidirectional
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Extern   0.0.0.0          173.17.99.2       0x8000003c  1847  0x22 0x2330  36
  mask 0.0.0.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   10.10.10.0       173.17.99.2       0x80000064   247  0x22 0x4103  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 200, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern  *10.150.99.0      10.150.0.99       0x800000cc  1597  0x22 0xabb3  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 1, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   172.16.2.0       173.17.99.2       0x80000069  1647  0x22 0x2d3a  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   172.16.3.0       173.17.99.2       0x80000069  1447  0x22 0x2244  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.0.1      173.17.99.2       0x80000064    47  0x22 0x359   36
  mask 255.255.255.255
  Topology default (ID 0)
    Type: 2, Metric: 100, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.0.2      173.17.99.2       0x80000067  1247  0x22 0x7b5   36
  mask 255.255.255.255
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.0.3      173.17.99.2       0x8000006b  2847  0x22 0xe072  36
  mask 255.255.255.255
  Topology default (ID 0)
    Type: 2, Metric: 100, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.0.4      173.17.99.2       0x80000007  2647  0x22 0x9f17  36
  mask 255.255.255.255
  Topology default (ID 0)
    Type: 2, Metric: 100, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.12.0     173.17.99.2       0x80000064  1047  0x22 0x9c19  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.13.0     173.17.99.2       0x80000064  2447  0x22 0x6982  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 200, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
Extern   192.168.14.0     173.17.99.2       0x8000005d  2247  0x22 0x6c85  36
  mask 255.255.255.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Topology default (ID 0)
….
```

The OSPF import policy “ospfimport” ensures that only the default route is installed in the routing table of vSRX.

Configurations for Task Completion:

vMX1:

```
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.3 transit-area
0.0.0.1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.4 transit-area
0.0.0.1
set protocols ospf spf-options holddown 3000
set protocols ospf spf-options rapid-runs 5
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
set protocols ospf area 0.0.0.1 interface lo0.0
```

vMX2:

```
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 hello-interval 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.3 transit-area
0.0.0.1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.4 transit-area
0.0.0.1
set routing-instances ospf-branch interface gr-0/0/10.0
set routing-instances ospf-branch routing-options interface-routes rib-group inet
branch-to-inet
set routing-instances ospf-branch routing-options static route 0.0.0.0/0 next-table
inet.0
set routing-instances ospf-branch protocols ospf rib-group branch-to-inet
set routing-instances ospf-branch protocols ospf export 2ospfbranch
set routing-instances ospf-branch protocols ospf area 0.0.0.0 interface gr-0/0/10.0
hello-interval 1
set policy-options policy-statement 2ospfbranch term 1 from protocol [ ospf direct
static ]
set policy-options policy-statement 2ospfbranch term 1 then accept
set policy-options policy-statement 2ospfbranch term 2 then reject
set routing-options rib-groups interfaces import-rib ospf-branch.inet.0
set routing-options rib-groups branch-to-inet import-rib [ ospf-branch.inet.0 inet.0 ]
set routing-options rib-groups inet-to-branch import-rib [ inet.0 ospf-branch.inet.0 ]
set protocols ospf spf-options holddown 3000
set protocols ospf spf-options rapid-runs 5
set protocols ospf rib-group inet-to-branch
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 flood-reduction
set protocols ospf area 0.0.0.1 interface lo0.0
```

vMX3:

```
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.1 transit-area
0.0.0.1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.2 transit-area
0.0.0.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.4 transit-area
0.0.0.1
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.1 interface lo0.0
```

vMX4:

```
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/2.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.1 transit-area
0.0.0.1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.2 transit-area
0.0.0.1
set protocols ospf3 area 0.0.0.0 virtual-link neighbor-id 192.168.0.3 transit-area
0.0.0.1
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 link-protection
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.1 interface lo0.0
```

vSRX:

```
set protocols ospf export 2ospf
set protocols ospf import ospfimport
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 hello-interval 1
set policy-options policy-statement 2ospf term 1 from interface irb.99
set policy-options policy-statement 2ospf term 1 then external type 1
set policy-options policy-statement 2ospf term 1 then accept
set policy-options policy-statement 2ospf term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement ospfimport term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement ospfimport term 1 then accept
set policy-options policy-statement ospfimport term 2 then reject
```

Task 5: BGP

This task focuses on the implementation of BGP for routing within the data center and within Site 2. BGP is also used for routing between the data center, ISP1, ISP2, and between the data center, Site 1, and Site 2. At various points in the topology you need to perform route redistribution between and IGP and BGP for end to end connectivity.

Device	Interface	Ipv4 address	Ipv6 address	Description
vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128	IBGP source interface
				v4 EBGP to AS 2222
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64	v6 EBGP to AS 2222
	ge-0/0/6	172.16.3.2/24	Link local	v6 EBGP to AS 2222
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128	IBGP source interface
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64	EBGP to AS 64513
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64	EBGP to AS 100001
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128	IBGP source interface
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64	EBGP to AS 64513
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64	EBGP to AS 100001
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table 42 BGP peering information

© 2020 Juniper Networks, Inc. All rights reserved.

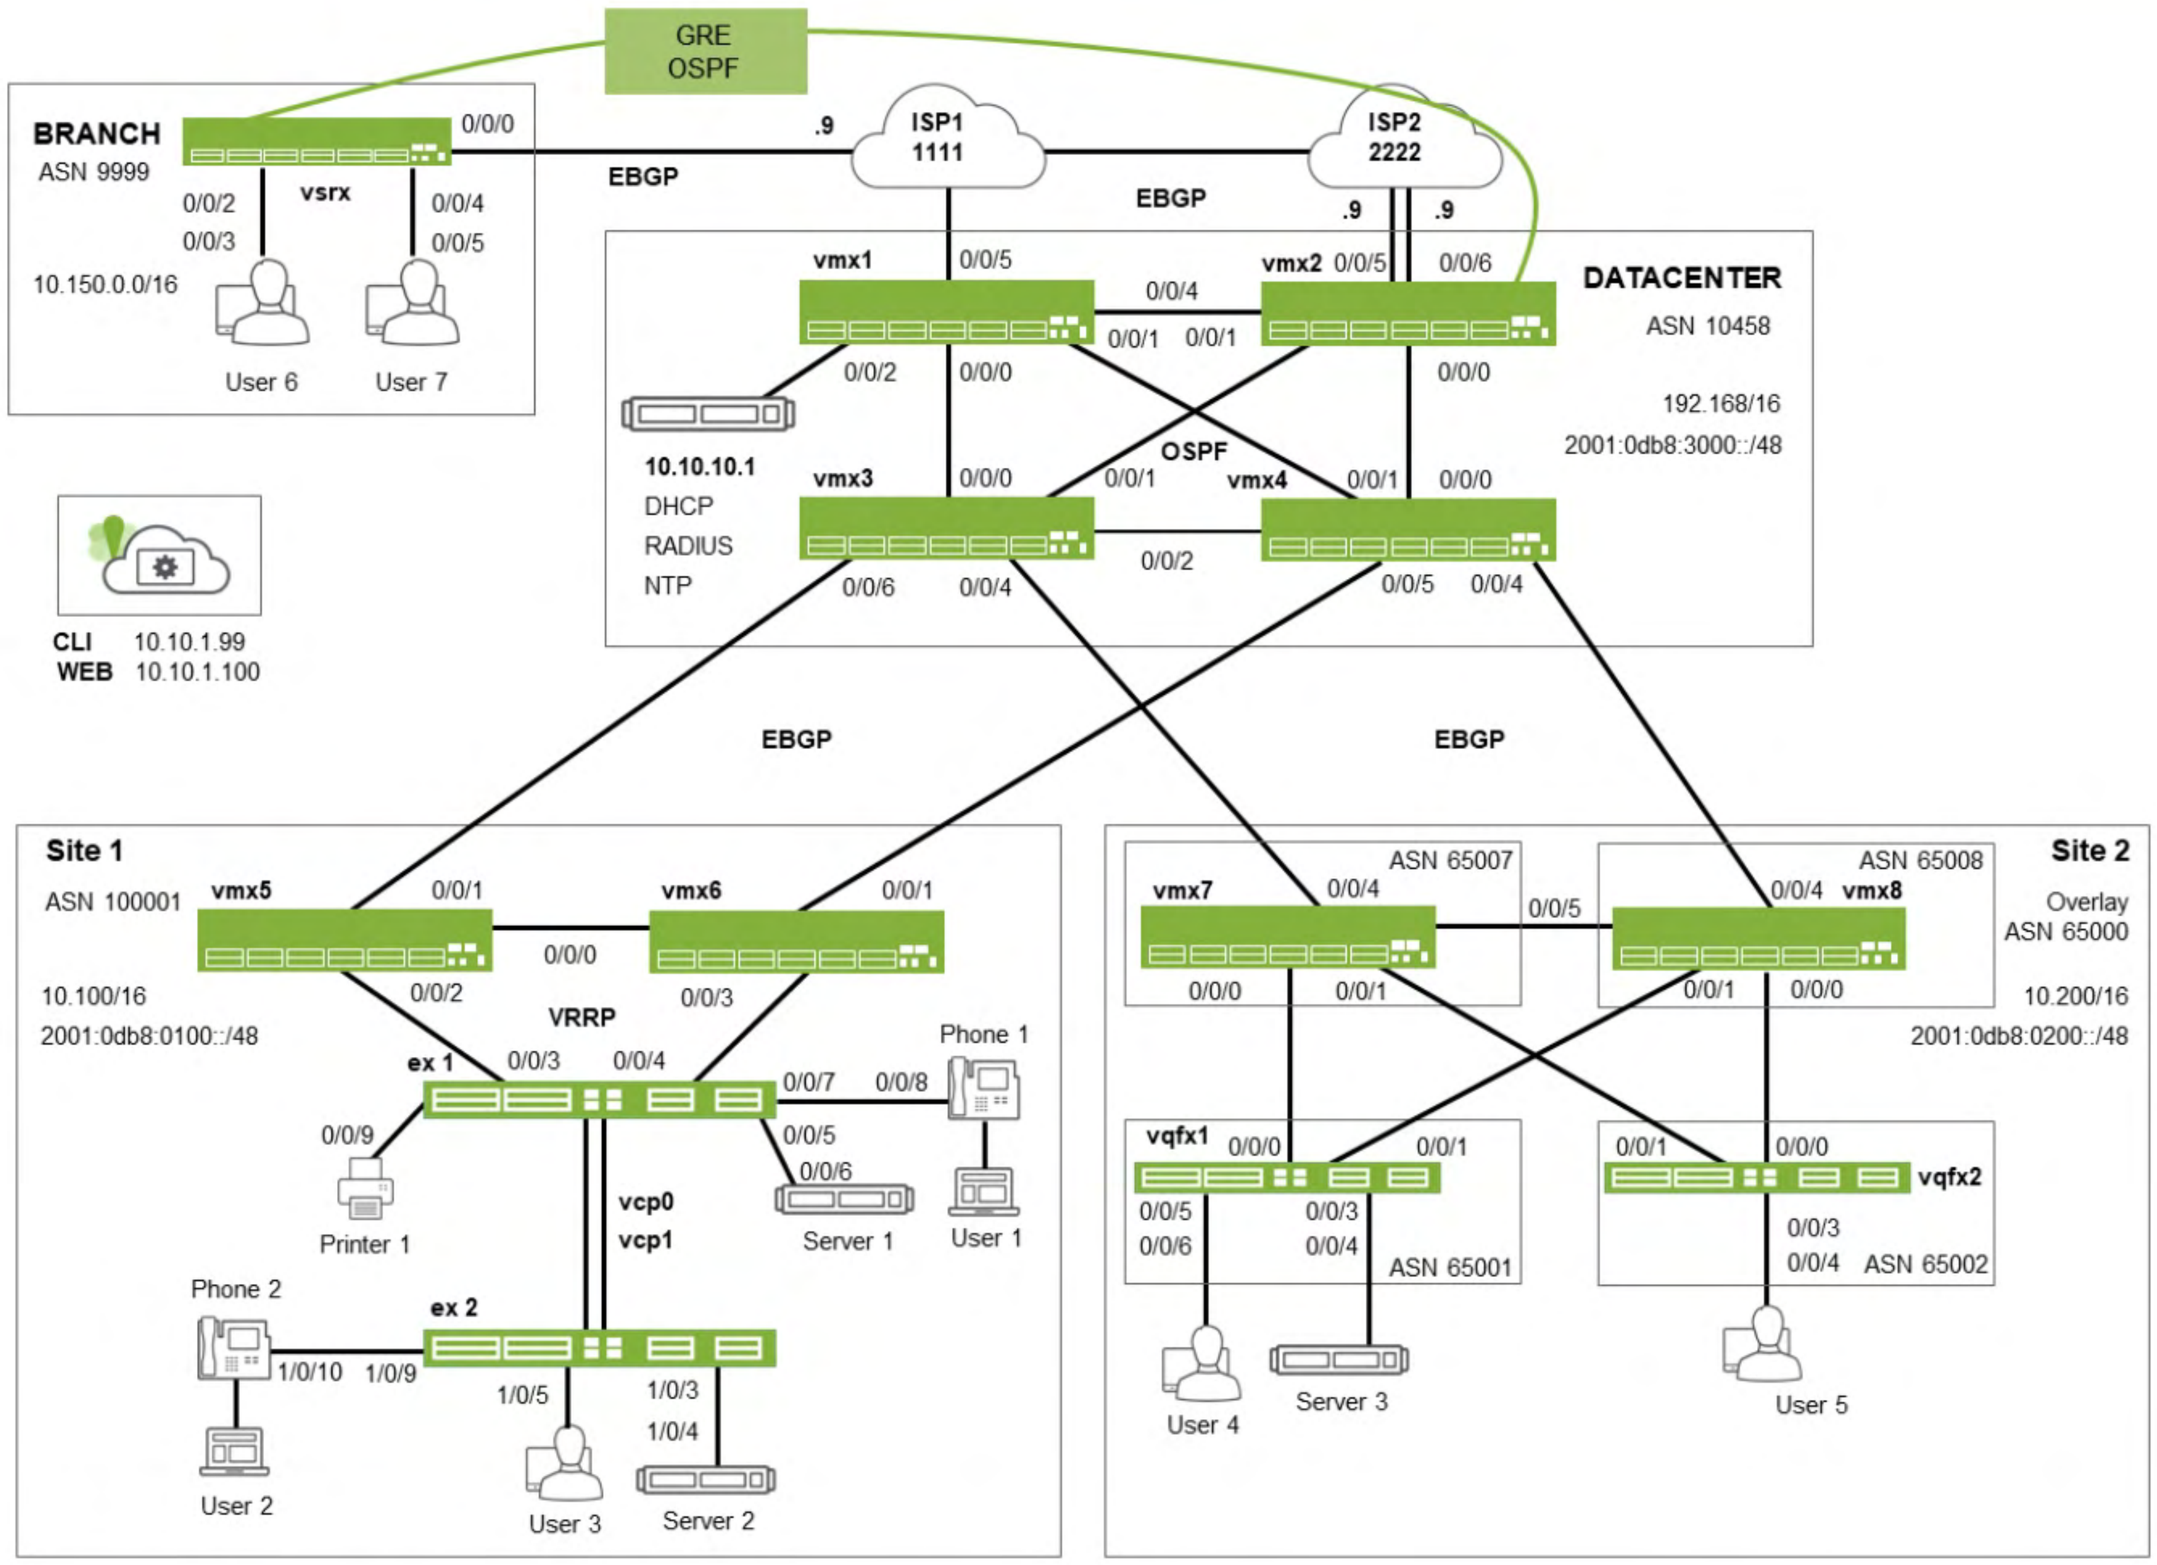

IBGP Data Center

1) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have a IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and also accept routes received from each other.
2) All IBGP sessions on vMX1 and vMX2 should be established on port 179 locally. Do not use the passive command or a firewall filter to accomplish this.
3) Enable BFD with authentication for IBGP sessions in the data center. Allow multihop sessions. Do not configure BFD between the Route Reflectors.

vMX1:

```
lab@vmx1# show routing-options autonomous-system
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
10458;

lab@vmx1# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.1;
    cluster 192.168.0.1;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.1;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::2;
}

lab@vmx1# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

vMX2:

```
lab@vmx2# show routing-options autonomous-system
10458;

lab@vmx2# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.2;
    cluster 192.168.0.2;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.2;
    neighbor 192.168.0.1;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.2;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::1;
}


lab@vmx2# show security
authentication-key-chains {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

vMX3:

```
lab@vmx3# show routing-options autonomous-system
10458;

lab@vmx3# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.3;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}

lab@vmx3# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

vMX4:

```
lab@vmx4# show routing-options autonomous-system
10458;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx4# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.4;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
lab@vmx4# show security
authentication-key-chains {
    key-chain bfd {
        key 1 {
            secret "$9$Q9T23/t1RSM87uO87-V4oz36"; ## SECRET-DATA
            start-time "2019-3-18.17:00:00 -0700";
        }
    }
}
```

Verification:

```
lab@vmx1# run show bgp summary | match 10458
192.168.0.2           10458       1075       1514       0       0     2:01:25 0/2/2/0
0/0/0/0
192.168.0.3           10458        268        271       0       0     2:00:11 0/0/0/0
0/0/0/0
192.168.0.4           10458        266        269       0       0     1:59:23 0/0/0/0
0/0/0/0
2001:db8:3000::2       10458        418        573       0       0      2:01:17 Establ
2001:db8:3000::3       10458        266        266       0       0      1:59:55 Establ
2001:db8:3000::4       10458        265        265       0       0      1:59:07 Establ
```

Note: vMX2, vMX3, and vMX4 should provide similar outputs

© 2020 Juniper Networks, Inc. All rights reserved.

4) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS and the Branch Office
5) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down. vMX2 should not advertise the default route if the BGP peering with ISP2 goes down

vMX1:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}

lab@vmx1# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.1.9;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vmx1# show policy-options policy-statement 2ibgp
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        accept;
    }
}
term 4 {
    then reject;
}

lab@vmx1# show protocols bgp group ibgp export
export 2ibgp;
```

vMX2:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.200.176;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx2# show policy-options policy-statement 2ibgp
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        accept;
    }
}
term 4 {
    then reject;
}
```

```
lab@vmx2# show protocols bgp
group ibgp {
    export 2ibgp;
}
```

Verification:

Note: The EBGP sessions with ISP1 and ISP2 must be configured before you verify this behavior. The following outputs were taken after the EBGP sessions to the ISPs were established which is carried out in steps 9 and 10.

```
lab@vmx1# run show route

inet.0: 245 destinations, 470 routes (245 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 00:00:08
                    > to 172.16.1.9 via ge-0/0/5.0
```

```
lab@vmx1# run show bgp summary | match 172.16.1.9
172.16.1.9             1111        417        216       0       0     1:36:47
222/222/222/0         0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# run show route advertising-protocol bgp 192.168.0.3

inet.0: 245 destinations, 470 routes (245 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               172.16.1.9                   100        I

lab@vmx1# set interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete
```

```
lab@vmx1# run show bgp summary | match 172.16.1.9
172.16.1.9              1111        419        218       0       1          30 Idle

lab@vmx1# run show route 0/0 exact
```

Since the EBGP peering with ISP1 is down, the generate route is not installed in the routing-table due to the applied policy.

```
lab@vmx1# delete interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete
```

6) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the AS

vMX1 & vMX2:

```
lab@vmx1# show routing-options
rib inet6.0 {
    static {
        route 0::/0 {
            discard;
            no-install;
        }
    }
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    then {
        accept;
    }
}
term 3 {
    then reject;
}
```

```
lab@vmx1# show protocols bgp group ibgp-v6 export
export 2ibgp-v6;
```

Verification:

```
lab@vmx1# run show route advertising-protocol bgp 2001:db8:3000::3

inet6.0: 102 destinations, 182 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                         150        I
```

IBGP Site 1

7) Configure IBGP in Site1 between vMX5 and vMX6. Use the information as shown in the table above.

vMX5:

```
lab@vmx5# show routing-options
autonomous-system 100001;

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    neighbor 2001:0db8:0100:0000::6;
}
```

vMX6:

```
lab@vmx6# show routing-options
autonomous-system 100001;

lab@vmx6# show protocols bgp
group ibgp {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    type internal;
    neighbor 10.100.0.5 {
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    neighbor 2001:0db8:0100:0000::5;
}
```

Verification:

```
lab@vmx5# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.100.0.6            100001       2064       2064       0       0    15:31:15 0/0/0/0
0/0/0/0
2001:db8:100::6       100001       2052       2051       0       0    15:31:15 Establ
  inet6.0: 0/0/0/0
```

EBGP Data Center

8) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
9) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.

vMX1:

```
lab@vmx1# show protocols bgp
group isp1 {
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
group isp1-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0db9:4000:0001::9 {
        peer-as 1111;
   }
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

```
lab@vmx1# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Unconfigured peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                        80         79          0          0          0          0
inet.0
                       226        224          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.9             1111        253         50       0       2       20:45
222/222/222/0          0/0/0/0
...
2001:db9:4000:1::9         1111        120         47       0       2       20:31 Establ
  inet6.0: 78/78/78/0
```

vMX2:

The task states to establish one IPv4 peering with ISP2 and we need to provide load balancing using both links towards ISP2. As ISP2 is outside of our control it must be preconfigured. This means that if we can debug the BGP packets received from ISP2 we might be able to find out valuable information to configure the BGP peering. Note that the BGP packets received from ISP2 can come in on interface ge-0/0/5 or interface ge-0/0/6. Keep this in mind when debugging the BGP packets.

In this example we are going to use the “monitor traffic” command to verify incoming BGP packets destinated to the control plane of vMX2.

```
lab@vmx2# run monitor traffic interface ge-0/0/5 extensive matching tcp
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on ge-0/0/5, capture size 1514 bytes

...

        PFE proto 2 (ipv4): (tos 0xc0, ttl  64, id 53971, offset 0, flags [none], proto:
TCP (6), length: 64) 172.16.200.176.55517 > 192.168.0.2.bgp: S 3936139201:3936139201(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 3639125731 0,sackOK,eol>
23:50:28.618250 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
```

The output above shows that we received an incoming tcp packet from 172.16.200.176 destined to the loopback address of vMX2 (192.168.0.2), port BGP.

© 2020 Juniper Networks, Inc. All rights reserved.

We now know the source and destination address for the IPv4 BGP peering with ISP2. Due to the load balancing requirement we need to make sure that IP 172.16.200.176 can be reached using both interfaces towards ISP2. For this we need to configure a static route.

vMX2:

```
lab@vmx2# show routing-options
static {
    route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ];
}
```

Now configure the BGP peering with ISP2. Keep in mind that the "multihop" statement is needed as EBGP sessions by default are single-hop sessions (TTL checking).

vMX2:

```
lab@vmx2# show protocols bgp
group isp2 {
    neighbor 172.16.200.176 {
        multihop;
        peer-as 2222;
    }
}
```

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called per-packet load balancing. The naming may be counter-intuitive. However, Junos per-packet load balancing is functionally equivalent to what other vendors may term per-flow load balancing. Configure load balancing to spread traffic across both links towards ISP2.

vMX2:

```
lab@vmx2# show policy-options policy-statement lbpp
term 1 {
    then {
        load-balance per-packet;
    }
}

lab@vmx2# show routing-options forwarding-table
export lbpp;
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.200.176         2222        238         36       0       0        14:39
222/222/222/0         0/0/0/0
...

lab@vmx2# run show route receive-protocol bgp 172.16.200.176

inet.0: 250 destinations, 255 routes (250 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 1.84.160.0/20           172.16.200.176                          2222 1111 33112 I
...

[edit]
lab@vmx2# run show route forwarding-table destination 1.84.160.1
...

Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
1.84.160.0/20      user     0                    indr  1048575   225
                                                 ulst  1048574     4

                              172.16.2.9         ucst      660     3 ge-0/0/5.0
                              172.16.3.9         ucst      659     4 ge-0/0/6.0
...
```

10) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.

First configure the IPv6 peering with ISP2 for the ge-0/0/5 link.

vMX2:

```
lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

The IPv6 peering over the ge-0/0/6 interface must be established using the IPv6 link local address. We need to find out the peering addresses using the “monitor traffic interface” command.

```
lab@vmx2# run monitor traffic interface ge-0/0/6 no-resolve matching tcp
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on ge-0/0/6, capture size 96 bytes

01:19:33.147860 Out IP6 truncated-ip6 - 31 bytes missing!fe80::250:56ff:fea9:43ac.64092
> fe80::250:56ff:fea9:7fec.179: P 3276364454:3276364473(19) ack 2827393527 win 16384
<[|tcp]>
01:19:33.257758  In IP6 fe80::250:56ff:fea9:7fec.179 > fe80::250:56ff:fea9:43ac.64092: .
ack 19 win 16384 <nop,nop,timestamp 3644476929[|tcp]>
^C

lab@vmx2# run show interfaces ge-0/0/6 | match fe80
        Destination: fe80::/64, Local: fe80::250:56ff:fea9:43ac

lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
    neighbor fe80::250:56ff:fea9:7fec {
        local-interface ge-0/0/6.0;
        peer-as 2222;
    }
}
```

As link-local addresses are only unique per interface, it is required to specify the “local-interface” command for the peering so that Junos known on which interface the peering must be established.

11) Advertise the IPv4 summary routes received from Site 1, Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs.

vMX1:

```
lab@vmx1# show routing-options aggregate
route 192.168.0.0/16;

lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        accept;
    }
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp
group isp1 {
    export 2ebgp;
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
```

Verification:

Note: You will not currently have all the summary routes from the other sites. You will need to complete the reset of this task before verifying this task. The sample output is provided after completing all configurations in this task.

```
lab@vmx1# run show route advertising-protocol bgp 172.16.1.9

inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.0.0/16           Self                                    100001 I
* 10.200.0.0/16           Self                                    I
* 192.168.0.0/16          Self                                    I
```

vMX2:

```
lab@vmx2# show routing-options aggregate
route 192.168.0.0/16;

[edit]
lab@vmx2# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then accept;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
term 2 {
    then reject;
}

lab@vmx2# show protocols bgp group isp2
export 2ebgp;
neighbor 172.16.200.176 {
    multihop;
    peer-as 2222;
}
```

Verification:

Note: You will not currently have all the summary routes from the other sites. You will need to complete the reset of this task before verifying this task. The sample output is provided after completing all configurations in this task.

```
lab@vmx2# run show route advertising-protocol bgp 172.16.200.176

inet.0: 251 destinations, 255 routes (251 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.0.0/16           Self                                    100001 I
* 10.200.0.0/16           Self                                    I
* 192.168.0.0/16          Self                                    I
```

12) Advertise a single /32 route to both ISPs that encompasses the IPv6 network of the data center, Site 1 and Site 2.

vMX1:

```
lab@vmx1# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;

lab@vmx1# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp group isp1-v6
family inet6 {
    unicast;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
export 2ebgp-v6;
neighbor 2001:0db9:4000:0001::9 {
    peer-as 1111;
}
```

vMX2:

```
lab@vmx2# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;

lab@vmx2# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}


lab@vmx2# show protocols bgp group isp2-v6
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0dba:4000:0002::9 {
    peer-as 2222;
}
neighbor fe80::250:56ff:fea9:7fec {
    local-interface ge-0/0/6.0;
    peer-as 2222;
}
```

Verification:

vMX1:

```
lab@vmx1# run show route advertising-protocol bgp 2001:0db9:4000:0001::9

inet6.0: 102 destinations, 181 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I
```

vMX2:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx2# run show route advertising-protocol bgp 2001:0dba:4000:0002::9

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I

[edit]
lab@vmx2# run show route advertising-protocol bgp fe80::250:56ff:fea9:7fec

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                    {100001} I
```

13) Ensure that ISP2 is the primary path to all IPv4 address blocks announced by your AS. Both ISPs apply route policy based on received communities. Use the information in the table below to meet the task requirement

ISP	COMMUNITY	ISP ACTION
ISP1	1111:1	Prepend 1x
	1111:2	Prepend 2x
	1111:3	Prepend 3x
ISP2	2222:1	Prepend 1x
	2222:2	Prepend 2x
	2222:3	Prepend 1x

vMX1:

```
lab@vmx1# show policy-options community prepend
members 1111:2;

lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        community add prepend;
        accept;
    }
}
term 2 {
    then reject;
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

Note: You can review the summary route on the vr-device using the *lab/lab123* credentials.

VR-Device:

```
lab@vr-device> show route 192.168.0.0

ISP1.inet.0: 232 destinations, 235 routes (232 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 01:07:01, localpref 100
                      AS path: 2222 10458 I, validation-state: unverified
                    >  to 10.20.20.2 via lt-0/0/0.1
                    [BGP/170] 02:36:50, localpref 100
...
                      AS path: 1111 1111 10458 I, validation-state: unverified
                    >  to 172.16.1.1 via ge-0/0/2.0
```

14) Ensure the ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.

vMX1:

```
lab@vmx1# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}
lab@vmx1# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
```

vMX2:

```
lab@vmx2# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        accept;
    }
}

lab@vmx2# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs

lab@vmx2# show protocols bgp | display set | match fromisp2
set protocols bgp group isp2 import fromisp2
set protocols bgp group isp2-v6 import fromisp2
```

Verification:

vMX1:

```
lab@vmx1# run show route protocol bgp

inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.84.160.0/20       *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 1111 33112 I, validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 33112 I, validation-state: unverified
                    > to 172.16.1.9 via ge-0/0/5.0
1.176.0.0/12        *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 172.16.1.9 via ge-0/0/5.0
3.189.0.0/16        *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    AS path: 1111 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                  > to 172.16.1.9 via ge-0/0/5.0

...
```

15) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2. Prefixes from Site1 that are advertised to both ISPs should contain the ASN of Site 1.

vMX1:

```
lab@vmx1# show protocols bgp | display set | match remove
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
```

vMX2:

```
lab@vmx2# show protocols bgp | display set | match remove-private
set protocols bgp group isp2 remove-private
set protocols bgp group isp2-v6 remove-private
```

16) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Site 1 and Site 2.

vMX1:

```
lab@vmx1# show policy-options community no-export
members no-export;

lab@vmx1# show policy-options policy-statement fromisp1
term 1 {
    then {
        community set no-export;
        accept;
    }
}

lab@vmx1# show protocols bgp group isp1
import fromisp1;
export 2ebgp;
remove-private;
neighbor 172.16.1.9 {
    peer-as 1111;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

```
lab@vmx2# show policy-options community no-export
members no-export;


lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        community set no-export;
        accept;
    }
}
```

Verification:

vMX3:

```
lab@vmx3# run show route protocol bgp detail

inet.0: 247 destinations, 249 routes (247 active, 0 holddown, 0 hidden)
...
1.84.160.0/20 (1 entry, 1 announced)
        *BGP    Preference: 170/-151
                Next hop type: Indirect, Next hop index: 0
                Address: 0xce3dd50
                Next-hop reference count: 450
                Source: 192.168.0.2
                Next hop type: Router, Next hop index: 643
                Next hop: 192.168.23.2 via ge-0/0/1.0, selected
                Session Id: 0x238
                Protocol next hop: 192.168.0.2
                Indirect next hop: 0xba3f580 1048574 INH Session ID: 0x256
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 1:18:42    Metric2: 10
                Validation State: unverified
                ORR Generation-ID: 0
                Task: BGP_10458.192.168.0.2
                Announcement bits (2): 0-KRT 4-Resolve tree 4
                AS path: 2222 1111 33112 I
                Communities: no-export
                Accepted
                Localpref: 150
                Router ID: 192.168.0.2
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: The next verification command was taken after the EBGP session was created between data center and Site 1 which is carried out in the next step.

vMX4:

```
lab@vmx4# run show route advertising-protocol bgp 192.168.46.6

inet.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    I
  10.10.10.0/24           Self                                    I
* 10.150.99.0/24          Self                                    I
* 10.200.0.0/16           Self                                    65008 I
* 192.168.0.0/16          Self                                    I
```

17) Site 1 and Site 2 should learn the summarized prefixes of the other site, data center, and branch office. Site 1 and Site 2 should also learn the IPv4 and IPv6 default route and the management network 10.10.10/24 that originated in the data center.
18) Ensure the device in Site1, Site2 and Branch Office have full reachability to each other and the Internet.

vMX1:

```
lab@vmx1# show policy-options policy-statement 2ibgp
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol direct;
        route-filter 10.10.10.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        next-hop self;
        accept;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
} term 4 {
    then reject;
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        next-hop self;
        accept;
    }
}
term 3 {
    then reject;
}
```

vMX2:

```
lab@vmx2# show policy-options policy-statement 2ibgp
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol ospf;
        route-filter 10.150.99.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
term 4 {
    then reject;
}

[edit]
lab@vmx2# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
term 3 {
    then reject;
}
```

vMX3:

```
lab@vmx3# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}

lab@vmx3# show protocols bgp
advertise-inactive;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
group ibgp {
    type internal;
    local-address 192.168.0.3;
    export nhs;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site1 {
    neighbor 192.168.35.5 {
        peer-as 100001;
    }
}
group site2 {
    neighbor 192.168.37.7 {
        peer-as 65007;
    }
}
group site1-v6 {
    neighbor 2001:0db8:3000:0035::5 {
        family inet6 {
            unicast;
        }
        peer-as 100001;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0037::7 {
        family inet6 {
            unicast;
        }
        peer-as 65007;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX4:

```
lab@vmx4# show policy-options
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx4# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.4;
    export nhs;
    bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 5;
        authentication {
            key-chain bfd;
            algorithm keyed-sha-1;
        }
        session-mode multihop;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site1 {
    neighbor 192.168.46.6 {
        peer-as 100001;
    }
}
group site2 {
    neighbor 192.168.48.8 {
        peer-as 65008;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
group site1-v6 {
    neighbor 2001:0db8:3000:0046::6 {
        family inet6 {
            unicast;
        }
        peer-as 100001;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0048::8 {
        family inet6 {
            unicast;
        }
        peer-as 65008;
    }
}
```

vMX5:

```
lab@vmx5# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol static;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol static;
            route-filter 2001:0db8:0100::/48 exact;
        }
        then accept;
    }
}
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
    export nhs;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    export nhs;
    neighbor 2001:0db8:0100:0000::6;
}
group dc {
    export 2bgp;
    neighbor 192.168.35.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0035::3 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vMX6:

```
lab@vmx6# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol static;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol static;
            route-filter 2001:0db8:0100::/48 exact;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
        then accept;
    }
}
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}
```

```
lab@vmx6# show protocols bgp
group ibgp {
    type internal;
    export nhs;
    neighbor 10.100.0.5 {
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    export nhs;
    neighbor 2001:0db8:0100:0000::5;
}
group dc {
    export 2bgp;
    neighbor 192.168.46.4 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0046::4 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vSRX1:

```
lab@vsrx# show routing-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
aggregate {
    route 10.150.0.0/16;
}

lab@vsrx# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            route-filter 10.150.0.0/16 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement frombgp {
    term 1 {
        from {
            route-filter 192.168.0.0/16 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
lab@vsrx# show protocols bgp
group ebgp {
    import frombgp;
    export 2bgp;
    neighbor 10.1.1.1 {
        peer-as 1111;
        local-as 9999;
    }
}
```

Verification Branch Office:

vSRX

```
lab@vsrx# 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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    225          1          0          0          0          0
Peer                         AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.1                   1111        313        113       0       0       49:37 Establ
  inet.0: 1/225/1/0

[edit]
lab@vsrx# run show route protocol bgp

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 00:49:52, localpref 100
                      AS path: 1111 2222 10458 I, validation-state: unverified
                    >  to 10.1.1.1 via ge-0/0/0.0
...

lab@vsrx# run show route advertising-protocol bgp 10.1.1.1

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.150.0.0/16           Self                                    I
```

Verification Site1:

vMX6:

```
lab@vmx6# run show route protocol bgp

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

0.0.0.0/0          *[BGP/170] 00:31:29, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 00:31:29, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
                    > to 10.100.56.5 via ge-0/0/0.0
10.10.10.0/24      *[BGP/170] 01:54:16, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 01:54:16, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
                    > to 10.100.56.5 via ge-0/0/0.0
10.150.99.0/24     *[BGP/170] 00:01:52, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.46.4 via ge-0/0/1.0
                    [BGP/170] 00:01:52, localpref 100, from 10.100.0.5
                      AS path: 10458 I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to 10.100.56.5 via ge-0/0/0.0
192.168.0.0/16     *[BGP/170] 01:54:16, localpref 100
                       AS path: 10458 I, validation-state: unverified
                     > to 192.168.46.4 via ge-0/0/1.0
                     [BGP/170] 01:54:16, localpref 100, from 10.100.0.5
                       AS path: 10458 I, validation-state: unverified
                     > to 10.100.56.5 via ge-0/0/0.0

inet6.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                 *[BGP/170] 02:16:44, localpref 100
                       AS path: 10458 I, validation-state: unverified
                     > to 2001:db8:3000:46::4 via ge-0/0/1.0
                     [BGP/170] 02:06:53, localpref 100, from 2001:db8:100::5
                       AS path: 10458 I, validation-state: unverified
                     > to 2001:db8:100:56::5 via ge-0/0/0.0
```

It seems we are missing the 2001:db8::/32 aggregate network from the data center. Verify if vMX4 announces the route to Site1.

vMX4:

```
lab@vmx4# run show route advertising-protocol bgp 2001:db8:3000:46::6
inet6.0: 27 destinations, 29 routes (27 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    I
* 2001:db8::/32           Self                                    {100001} I
```

The IPv6 aggregate route is announced by vMX4 to Site1, but not accepted by Site1. It seems that vMX1 and vMX2 are aggregating the 2001:db8::/32 network using AS 100001 contributing routes. This is the reason that Site1 does not accept the route. To fix this we need to ensure that vMX1 and vMX2 create the aggregate route 2001:db8::/32 based on local contributing routes. We can do this with route policy as shown below

vMX1:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
}
```

vMX2:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
}
```

Verification:

```
lab@vmx4# run show route advertising-protocol bgp 2001:db8:3000:46::6

inet6.0: 31 destinations, 33 routes (31 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    I
* 2001:db8::/32           Self                                    I
```

```
lab@vmx6# run show route receive-protocol bgp 2001:db8:3000:46::4

...
inet6.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    2001:db8:3000:46::4                     10458 I
* 2001:db8::/32           2001:db8:3000:46::4                     10458 I
```

vMX7:

```
lab@vmx7# show protocols bgp
group dc {
    export 2bgp;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    local-as 65007;
    neighbor 192.168.37.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    local-as 65007;
    neighbor 2001:0db8:3000:0037::3 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}

lab@vmx7# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}

lab@vmx7# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
```

vMX8:

```
lab@vmx8# show protocols bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
group dc {
    export 2bgp;
    local-as 65008;
    neighbor 192.168.48.4 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    local-as 65008;
    neighbor 2001:0db8:3000:0048::4 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
lab@vmx8# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}

lab@vmx8# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
```

Verification Site 2:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx7# run show route protocol bgp

inet.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 02:41:54, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 02:41:53, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.10.10.0/24      *[BGP/170] 04:29:15, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 19:18:58, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.100.0.0/16      *[BGP/170] 02:41:00, localpref 100
                      AS path: 10458 100001 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 19:18:58, localpref 100, from 10.200.0.8
                      AS path: 10458 100001 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
10.150.99.0/24     *[BGP/170] 00:34:31, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 00:34:31, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0
192.168.0.0/16     *[BGP/170] 19:20:18, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
                    [BGP/170] 19:18:58, localpref 100, from 10.200.0.8
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.78.8 via ge-0/0/5.0

inet6.0: 15 destinations, 17 routes (14 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

::/0               *[BGP/170] 19:20:10, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
                    [BGP/170] 19:19:02, localpref 100, from 2001:db8:3000::8
                      AS path: 10458 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
2001:db8::/32      *[BGP/170] 00:03:48, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
                    [BGP/170] 00:03:48, localpref 100, from 2001:db8:3000::8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      AS path: 10458 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
2001:db8:100::/48   *[BGP/170] 19:19:02, localpref 100, from 2001:db8:3000::8
                      AS path: 10458 100001 I, validation-state: unverified
                    > to fe80::250:56ff:fea9:4f3c via ge-0/0/5.0
```

Configurations for Task Completion:

vMX1:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.1
set protocols bgp group ibgp cluster 192.168.0.1
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.1
set protocols bgp group rr neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.1
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::1
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.1.9
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set policy-options policy-statement 2ibgp term 3 from protocol aggregate
set policy-options policy-statement 2ibgp term 3 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2ibgp term 3 then accept
set policy-options policy-statement 2ibgp term 4 then reject
set protocols bgp group ibgp export 2ibgp
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
set protocols bgp group isp1 neighbor 172.16.1.9 peer-as 1111
set protocols bgp group isp1-v6 family inet6 unicast
set protocols bgp group isp1-v6 neighbor 2001:0db9:4000:0001::9 peer-as 1111
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp1 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp1-v6 export 2ebgp-v6
set policy-options community prepend members 1111:2
set policy-options policy-statement 2ebgp term 1 then community add prepend
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
set policy-options community no-export members no-export
set policy-options policy-statement fromisp1 term 1 then community set no-export
set policy-options policy-statement fromisp1 term 1 then accept
set protocols bgp group isp1 import fromisp1
set policy-options policy-statement 2ibgp term 1 from protocol aggregate
set policy-options policy-statement 2ibgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ibgp term 1 then accept
set policy-options policy-statement 2ibgp term 2 from protocol direct
set policy-options policy-statement 2ibgp term 2 from route-filter 10.10.10.0/24 exact
set policy-options policy-statement 2ibgp term 2 then accept
insert policy-options policy-statement 2ibgp term 1 before term 3
insert policy-options policy-statement 2ibgp term 2 before term 3
set policy-options policy-statement 2ibgp term 3 then next-hop self
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
set policy-options policy-statement 2ibgp-v6 term 2 then next-hop self
set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::/48 orlonger
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX2:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.2
set protocols bgp group ibgp cluster 192.168.0.2
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.2
set protocols bgp group rr neighbor 192.168.0.1
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.2
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::1
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.200.176
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set policy-options policy-statement 2ibgp term 3 from protocol aggregate
set policy-options policy-statement 2ibgp term 3 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2ibgp term 3 then accept
set policy-options policy-statement 2ibgp term 4 then reject
set protocols bgp group ibgp export 2ibgp
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
set routing-options static route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ]
set protocols bgp group isp2 neighbor 172.16.200.176 multihop
set protocols bgp group isp2 neighbor 172.16.200.176 peer-as 2222
set policy-options policy-statement lbpp term 1 then load-balance per-packet
set routing-options forwarding-table export lbpp
set protocols bgp group isp2-v6 family inet6 unicast
set protocols bgp group isp2-v6 neighbor 2001:0dba:4000:0002::9 peer-as 2222
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec local-interface ge-
0/0/6.0
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec peer-as 2222
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp2 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp2-v6 export 2ebgp-v6
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set policy-options policy-statement fromisp2 term 1 then local-preference 150
set policy-options policy-statement fromisp2 term 1 then accept
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set protocols bgp group isp2 import fromisp2
set protocols bgp group isp2-v6 import fromisp2
set protocols bgp group isp2 remove-private
set protocols bgp group isp2-v6 remove-private
set policy-options community no-export members no-export
set policy-options policy-statement fromisp2 term 1 then community set no-export
set policy-options policy-statement 2ibgp term 1 from protocol aggregate
set policy-options policy-statement 2ibgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ibgp term 1 then accept
set policy-options policy-statement 2ibgp term 2 from protocol ospf
set policy-options policy-statement 2ibgp term 2 from route-filter 10.150.99.0/24 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ibgp term 2 then accept
insert policy-options policy-statement 2ibgp term 1 before term 3
insert policy-options policy-statement 2ibgp term 2 before term 3
set policy-options policy-statement 2ibgp term 3 then local-preference 150
set policy-options policy-statement 2ibgp term 3 then next-hop self
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
set policy-options policy-statement 2ibgp-v6 term 2 then next-hop self
set policy-options policy-statement 2ibgp-v6 term 2 then local-preference 150
```

vMX3:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.3
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::3
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site1 neighbor 192.168.35.5 peer-as 100001
set protocols bgp group site2 neighbor 192.168.37.7 peer-as 65007
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0035::5 family inet6 unicast
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0035::5 peer-as 100001
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 peer-as 65007
set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::2/128 exact
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX4:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.4
set protocols bgp group ibgp bfd-liveness-detection minimum-interval 1000
set protocols bgp group ibgp bfd-liveness-detection multiplier 5
set protocols bgp group ibgp bfd-liveness-detection authentication key-chain bfd
set protocols bgp group ibgp bfd-liveness-detection authentication algorithm keyed-sha-1
set protocols bgp group ibgp bfd-liveness-detection session-mode multihop
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::4
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set security authentication-key-chains key-chain bfd key 1 secret "$9$Q9T23/t1RSM87uO87-
V4oz36"
set security authentication-key-chains key-chain bfd key 1 start-time "2019-3-
18.17:00:00 -0700"
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site1 neighbor 192.168.46.6 peer-as 100001
set protocols bgp group site2 neighbor 192.168.48.8 peer-as 65008
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0046::6 family inet6 unicast
set protocols bgp group site1-v6 neighbor 2001:0db8:3000:0046::6 peer-as 100001
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 peer-as 65008
```

vMX5:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 10.100.0.6 local-address 10.100.0.5
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::5
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::6
set policy-options policy-statement 2bgp term 1 from protocol static
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol static
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.35.3 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 peer-as 10458
```

vMX6:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 10.100.0.5 local-address 10.100.0.6
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::6
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::5
set policy-options policy-statement 2bgp term 1 from protocol static
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol static
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.46.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 peer-as 10458
```

vMX7:

```
set protocols bgp group dc export 2bgp
set protocols bgp group dc local-as 65007
set protocols bgp group dc neighbor 192.168.37.3 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 local-as 65007
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 peer-as 10458
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
```

vMX8:

```
set protocols bgp group dc export 2bgp
set protocols bgp group dc local-as 65008
set protocols bgp group dc neighbor 192.168.48.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 local-as 65008
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 peer-as 10458
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
```

vSRX:

```
set routing-options aggregate route 10.150.0.0/16
set policy-options policy-statement 2bgp term 1 from route-filter 10.150.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp term 2 then reject
set policy-options policy-statement frombgp term 1 from route-filter 192.168.0.0/16
exact
set policy-options policy-statement frombgp term 1 then accept
set policy-options policy-statement frombgp term 2 then reject
set protocols bgp group ebgp import frombgp
set protocols bgp group ebgp export 2bgp
set protocols bgp group ebgp neighbor 10.1.1.1 peer-as 1111
set protocols bgp group ebgp neighbor 10.1.1.1 local-as 9999
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 6: EVPN/VXLAN (Spine and Leaf)

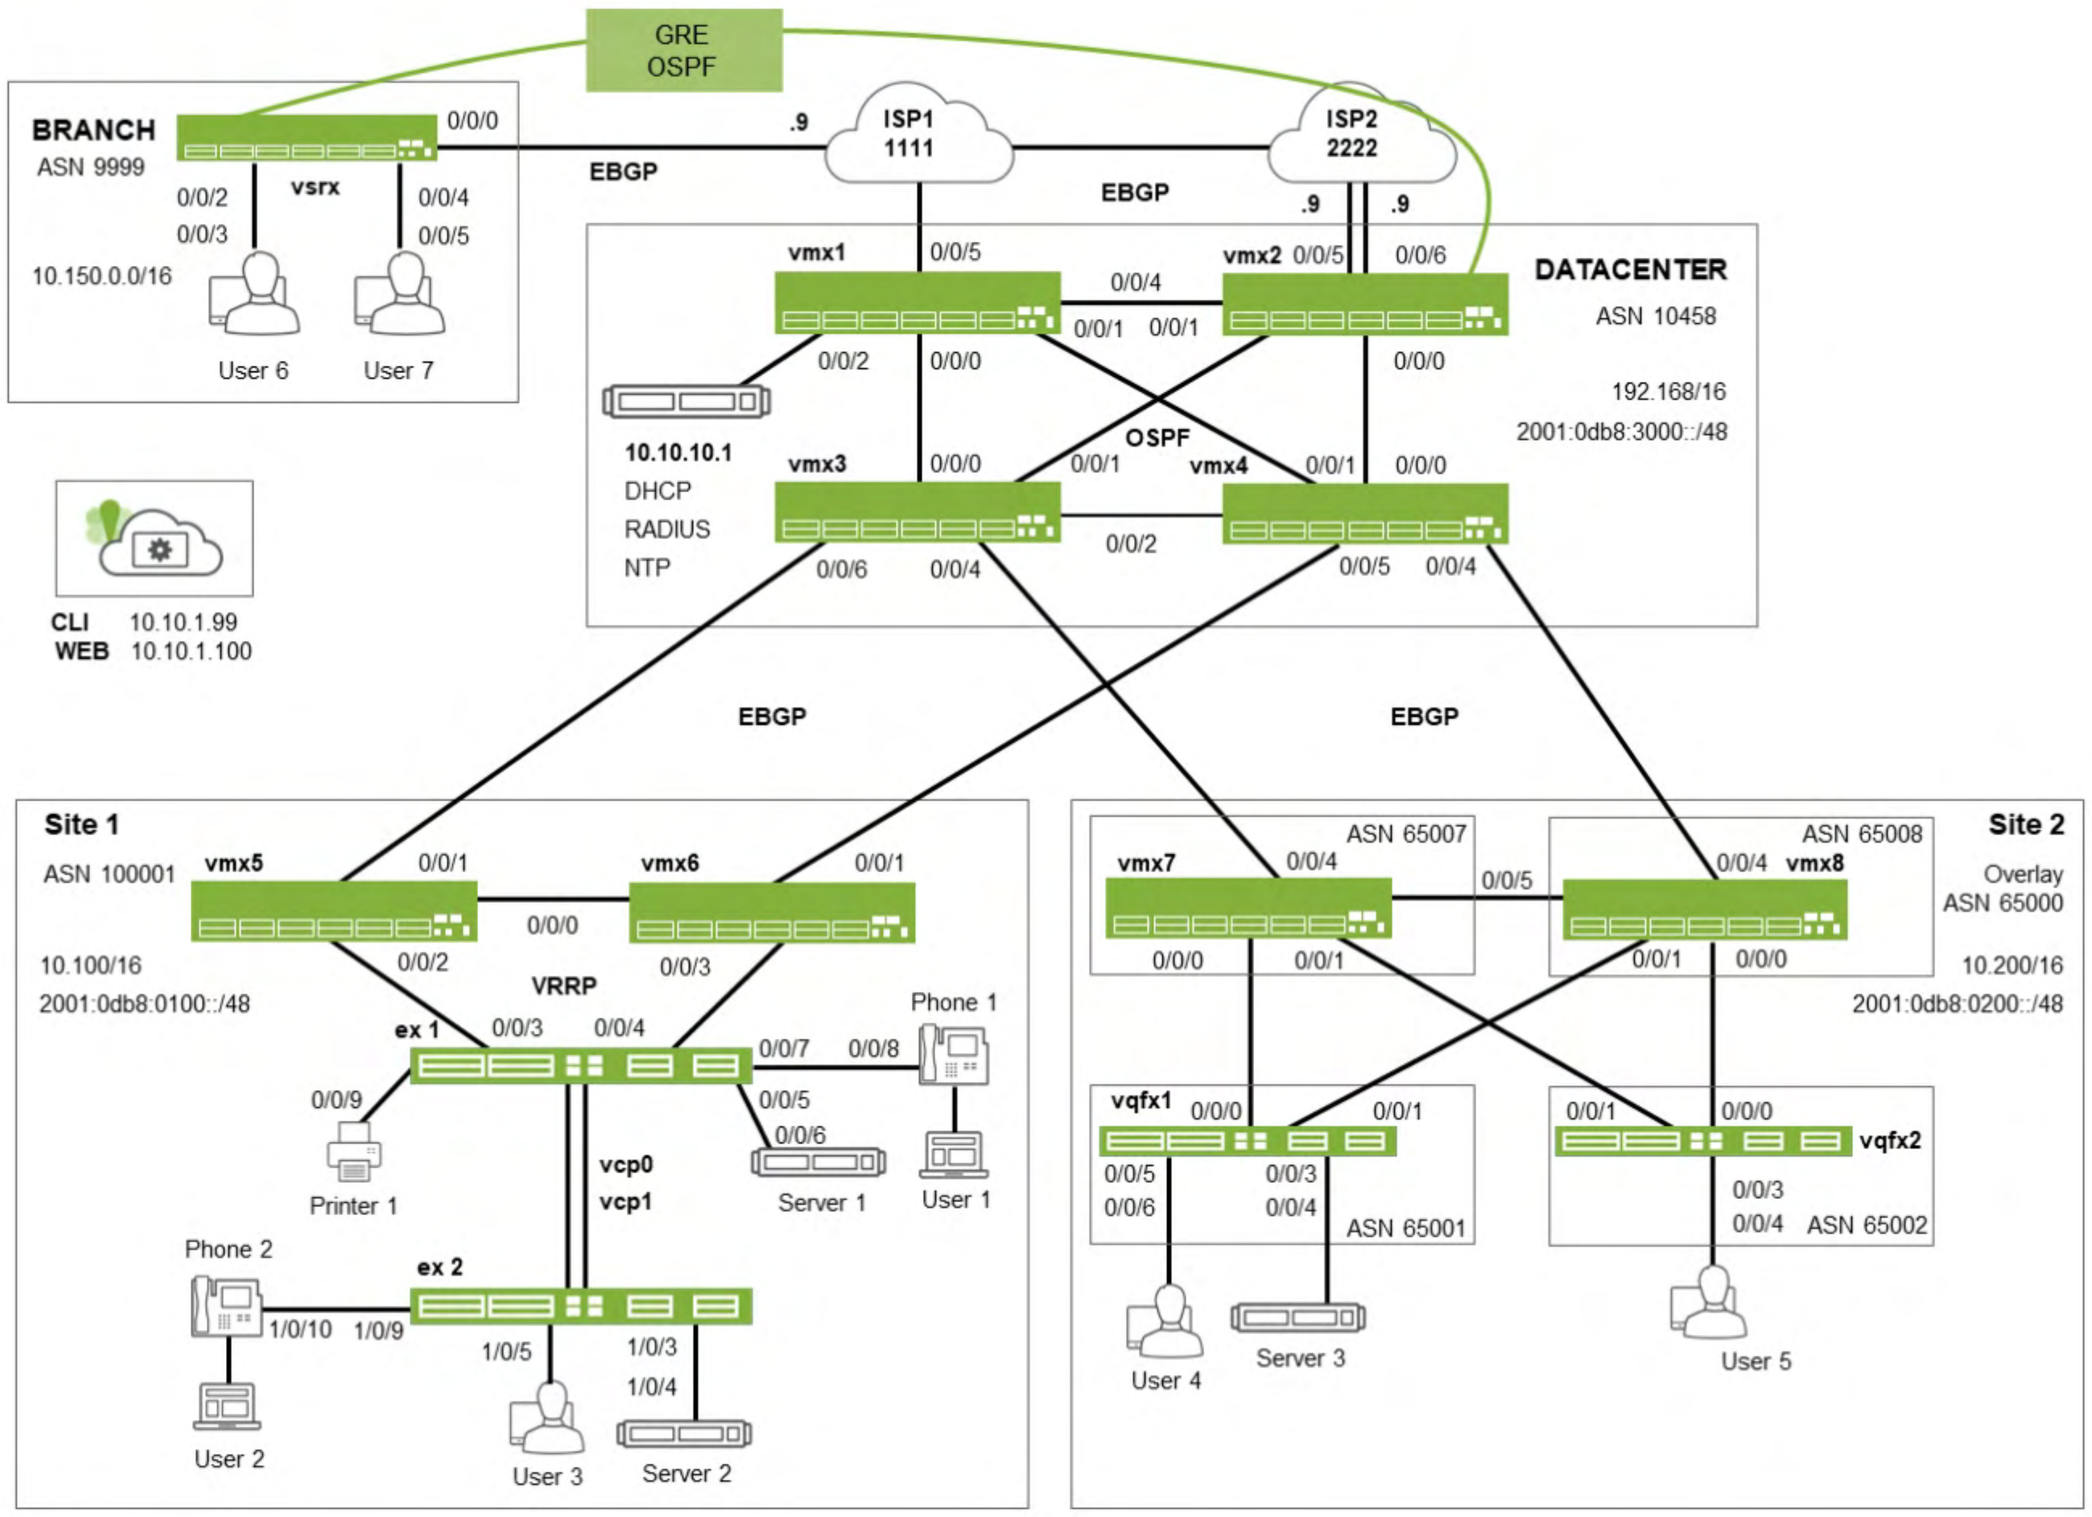

EVPN VXLAN

In this task you will configure the layer 2 connectivity in Site 2 using the EVPN/VXLAN technology.

1) Configure the EBGP underlay in Site 2 to exchange device loopback reachability. Use the AS numbers as shown in the topology. Use the interface addresses to establish the EBGP peerings.

vMX7:

```
lab@vmx7# show protocols bgp group underlay
export export-loopback;
local-as 65007;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
multipath {
    multiple-as;
}
neighbor 10.200.17.1 {
    peer-as 65001;
}
neighbor 10.200.27.2 {
    peer-as 65002;
}
lab@vmx7# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

vMX8:

```
lab@vmx8# show protocols bgp group underlay
export export-loopback;
local-as 65008;
multipath {
    multiple-as;
}
neighbor 10.200.18.1 {
    peer-as 65001;
}
neighbor 10.200.28.2 {
    peer-as 65002;
}

lab@vmx8# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

vQFX1:

```
lab@vQFX1# show protocols bgp group underlay
export export-loopback;
local-as 65001;
multipath {
    multiple-as;
}
neighbor 10.200.17.7 {
    peer-as 65007;
}
neighbor 10.200.18.8 {
    peer-as 65008;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vQFX1# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

vQFX2:

```
lab@vQFX2# show protocols bgp group underlay
export export-loopback;
local-as 65002;
multipath {
    multiple-as;
}
neighbor 10.200.27.7 {
    peer-as 65007;
}
neighbor 10.200.28.8 {
    peer-as 65008;
}

lab@vQFX2# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

Verification:

vMX7:

```
lab@vmx7# run show bgp summary group underlay
Groups: 4 Peers: 7 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.17.1           65001         45         50       0       0       17:57 Establ
  inet.0: 2/3/3/0
10.200.27.2           65002         44         50       0       0       17:14 Establ
  inet.0: 2/3/3/0
```

vMX8:

```
lab@vmx8# run show bgp summary group underlay
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Groups: 4 Peers: 7 Down peers: 0
Table           Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.18.1          65001         52         51       0       0       18:28 Establ
  inet.0: 2/7/7/0
10.200.28.2          65002         53         52       0       0       17:48 Establ
  inet.0: 2/7/7/0
```

vQFX1:

```
lab@vQFX1# run show bgp summary group underlay
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table           Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.17.7          65007         54         48       0       0       19:09 7/7/7/0
0/0/0/0
10.200.18.8          65008         53         52       0       0       19:01 7/7/7/0
0/0/0/0
```

vQFX2:

```
lab@vQFX2# run show bgp summary group underlay
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table           Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.27.7          65007         54         46       0       0       18:28 7/8/8/0
0/0/0/0
10.200.28.8          65008         53         52       0       0       18:24 7/8/8/0
0/0/0/0
```

2) Configure the IBGP EVPN overlay for Site 2. vMX7 and vMX8 are router reflectors for the EVPN overlay. Leaf devices should only peer with Spine devices.

vMX7:

```
lab@vmx7# show protocols bgp group overlay
type internal;
local-address 10.200.0.7;
family evpn {
    signaling;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
cluster 10.200.0.7;
local-as 65000;
multipath;
neighbor 10.200.0.1;
neighbor 10.200.0.2;
neighbor 10.200.0.8;
```

vMX8:

```
lab@vmx8# show protocols bgp group overlay
type internal;
local-address 10.200.0.8;
family evpn {
    signaling;
}
cluster 10.200.0.8;
local-as 65000;
multipath;
neighbor 10.200.0.1;
neighbor 10.200.0.2;
neighbor 10.200.0.7;
```

vQFX1:

```
lab@vQFX1# show protocols bgp group overlay
type internal;
local-address 10.200.0.1;
family evpn {
    signaling;
}
local-as 65000;
multipath;
neighbor 10.200.0.7;
neighbor 10.200.0.8;
```

vQFX2:

```
lab@vQFX2# show protocols bgp group overlay
type internal;
local-address 10.200.0.2;
family evpn {
    signaling;
}
local-as 65000;
multipath;
neighbor 10.200.0.7;
neighbor 10.200.0.8;
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

vMX7:

```
lab@vmx7# run show bgp summary group overlay
Groups: 4 Peers: 7 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1            65000          7          5       0       0        2:13 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.2            65000          8          7       0       0        2:39 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.8            65000         11         10       0       0        4:33 Establ
  bgp.evpn.0: 0/0/0/0
```

vMX8:

```
lab@vmx8# run show bgp summary group overlay
Groups: 4 Peers: 7 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1            65000          9          7       0       0        2:43 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.2            65000         10          8       0       0        3:08 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.7            65000         13         12       0       0        5:06 Establ
  bgp.evpn.0: 0/0/0/0
```

vQFX1:

```
lab@vQFX1# run show bgp summary group overlay
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.7            65000          8          8       0       0        2:52 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.8            65000          8          8       0       0        2:48 Establ
  bgp.evpn.0: 0/0/0/0
```

vQFX2:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vQFX2# run show bgp summary group overlay
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.7            65000          9          8       0       0        3:20 Establ
  bgp.evpn.0: 0/0/0/0
10.200.0.8            65000          9          9       0       0        3:16 Establ
  bgp.evpn.0: 0/0/0/0
```

3) Ensure that virtual gateway addresses respond to ICMP packets.
4) Configure one EVPN VLAN aware service on vMX7 and vMX8 for VLAN 100 and 101

vMX7:

```
lab@vmx7# show policy-options
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}
community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;

lab@vmx7# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
    route-distinguisher 10.200.0.7:1;
    vrf-import fabric;
    vrf-target target:65000:1000;
    protocols {
        evpn {
            encapsulation vxlan;
            extended-vni-list [ 100 101 ];
        }
    }
    bridge-domains {
        vlan100 {
            domain-type bridge;
            vlan-id 100;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        routing-interface irb.100;
        vxlan {
            vni 100;
            ingress-node-replication;
        }
    }
    vlan101 {
        domain-type bridge;
        vlan-id 101;
        routing-interface irb.101;
        vxlan {
            vni 101;
            ingress-node-replication;
        }
    }
}

lab@vmx7# show interfaces irb
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.250/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::250/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.250/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::250/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

vMX8:

```
lab@vmx8# show policy-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}

community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;
```

```
lab@vmx8# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
    route-distinguisher 10.200.0.8:1;
    vrf-import fabric;
    vrf-target target:65000:1000;
    protocols {
        evpn {
            encapsulation vxlan;
            extended-vni-list [ 100 101 ];
        }
    }
    bridge-domains {
        vlan100 {
            domain-type bridge;
            vlan-id 100;
            routing-interface irb.100;
            vxlan {
                vni 100;
                ingress-node-replication;
            }
        }
        vlan101 {
            domain-type bridge;
            vlan-id 101;
            routing-interface irb.101;
            vxlan {
                vni 101;
                ingress-node-replication;
            }
        }
    }
}

lab@vmx8# show interfaces irb
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.251/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::251/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.251/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::251/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

vQFX1:

```
lab@vQFX1# show switch-options
vtep-source-interface lo0.0;
route-distinguisher 10.200.0.1:1;
vrf-import fabric;
vrf-target target:65000:1;

lab@vQFX1# show protocols evpn
vni-options {
    vni 100 {
        vrf-target target:65000:100;
    }
    vni 101 {
        vrf-target target:65000:101;
    }
}
encapsulation vxlan;
multicast-mode ingress-replication;
extended-vni-list [ 100 101 ];

lab@vQFX1# show policy-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}
community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;

lab@vQFX1# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
    vxlan {
        vni 101;
    }
}
user {
    vlan-id 100;
    vxlan {
        vni 100;
    }
}
```

vQFX2:

```
lab@vQFX2# show switch-options
vtep-source-interface lo0.0;
route-distinguisher 10.200.0.2:1;
vrf-import fabric;
vrf-target target:65000:1;

lab@vQFX2# show protocols evpn
vni-options {
    vni 100 {
        vrf-target target:65000:100;
    }
    vni 101 {
        vrf-target target:65000:101;
    }
}
encapsulation vxlan;
multicast-mode ingress-replication;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
extended-vni-list [ 100 101 ];

lab@vQFX2# show policy-options
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}
community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;

lab@vQFX2# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
    vxlan {
        vni 101;
    }
}
user {
    vlan-id 100;
    vxlan {
        vni 100;
    }
}
```

5) Ensure the devices in Site 2 can reach each other, the device in Site1 and the Internet

vQFX1:

```
lab@vQFX1# run ping 10.200.100.5 routing-instance user4
PING 10.200.100.5 (10.200.100.5): 56 data bytes
64 bytes from 10.200.100.5: icmp_seq=0 ttl=64 time=578.779 ms
64 bytes from 10.200.100.5: icmp_seq=1 ttl=64 time=502.334 ms
64 bytes from 10.200.100.5: icmp_seq=2 ttl=64 time=504.193 ms
^C
--- 10.200.100.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 502.334/528.435/578.779/35.606 ms

lab@vQFX1# run ping 10.200.101.3 routing-instance user4
PING 10.200.101.3 (10.200.101.3): 56 data bytes
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
64 bytes from 10.200.101.3: icmp_seq=0 ttl=63 time=700.634 ms
64 bytes from 10.200.101.3: icmp_seq=1 ttl=63 time=700.138 ms
^C
--- 10.200.101.3 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 700.138/700.386/700.634/0.248 ms
```

vQFX2:

```
lab@vQFX2# run ping 10.100.100.2 routing-instance user5
PING 10.100.100.2 (10.100.100.2): 56 data bytes
64 bytes from 10.100.100.2: icmp_seq=0 ttl=60 time=300.049 ms
64 bytes from 10.100.100.2: icmp_seq=1 ttl=60 time=300.238 ms
^C
--- 10.100.100.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 300.049/300.144/300.238/0.094 ms

lab@vQFX2# run ping 10.100.100.1 routing-instance user5
PING 10.100.100.1 (10.100.100.1): 56 data bytes
64 bytes from 10.100.100.1: icmp_seq=0 ttl=60 time=300.708 ms
^C
--- 10.100.100.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 300.708/300.708/300.708/0.000 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

EX1:

Note: If you have not previously verified DHCP is working you may need to bounce the ge-1/0/4 and ge-0/0/6 interfaces by disabling and enabling them.

```
lab@EX1# run show dhcp client binding routing-instance server1
IP address        Hardware address   Expires     State      Interface
10.100.101.50     cc:e1:94:d2:6e:49  26005       BOUND      ge-0/0/6.0

lab@EX1# run show dhcp client binding routing-instance server2
IP address        Hardware address   Expires     State      Interface
10.100.101.51     b8:c2:53:97:97:a7  26000       BOUND      ge-1/0/4.0
```

vQFX2:

```
lab@vQFX2# run ping 10.100.101.50 routing-instance user5
PING 10.100.101.50 (10.100.101.50): 56 data bytes
64 bytes from 10.100.101.50: icmp_seq=0 ttl=60 time=380.841 ms
^C
--- 10.100.101.50 ping statistics ---
2 packets transmitted, 1 packets received, 50% packet loss
round-trip min/avg/max/stddev = 380.841/380.841/380.841/0.000 ms

lab@vQFX2# run ping 10.100.101.51 routing-instance user5
PING 10.100.101.51 (10.100.101.51): 56 data bytes
64 bytes from 10.100.101.51: icmp_seq=0 ttl=60 time=301.208 ms
^C
--- 10.100.101.51 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 301.208/301.208/301.208/0.000 ms
```

Configurations for Task Completion:

vMX7:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65007
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.17.1 peer-as 65001
set protocols bgp group underlay neighbor 10.200.27.2 peer-as 65002
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.7
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 10.200.0.7
set protocols bgp group overlay local-as 65000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.1
set protocols bgp group overlay neighbor 10.200.0.2
set protocols bgp group overlay neighbor 10.200.0.8
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 route-distinguisher 10.200.0.7:1
set routing-instances evpn1 vrf-import fabric
set routing-instances evpn1 vrf-target target:65000:1000
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list 100
set routing-instances evpn1 protocols evpn extended-vni-list 101
set routing-instances evpn1 bridge-domains vlan100 domain-type bridge
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan100 vxlan ingress-node-replication
set routing-instances evpn1 bridge-domains vlan101 domain-type bridge
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set routing-instances evpn1 bridge-domains vlan101 vxlan ingress-node-replication
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 100 family inet address 10.200.100.250/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::250/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 virtual-gateway-accept-data
set interfaces irb unit 101 family inet address 10.200.101.250/24 virtual-gateway-
address 10.200.101.254
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::250/64 virtual-
gateway-address 2001:0db8:0200:101::254
set interfaces lo0 unit 0 family inet address 10.200.0.7/32 primary
set interfaces lo0 unit 0 family inet address 10.200.0.7/32 preferred
set interfaces lo0 unit 0 family inet6 address 2001:0db8:3000:0000::7/128
```

vMX8:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65008
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.18.1 peer-as 65001
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group underlay neighbor 10.200.28.2 peer-as 65002
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.8
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 10.200.0.8
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.1
set protocols bgp group overlay neighbor 10.200.0.2
set protocols bgp group overlay neighbor 10.200.0.7
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 route-distinguisher 10.200.0.8:1
set routing-instances evpn1 vrf-import fabric
set routing-instances evpn1 vrf-target target:65000:1000
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list 100
set routing-instances evpn1 protocols evpn extended-vni-list 101
set routing-instances evpn1 bridge-domains vlan100 domain-type bridge
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan100 vxlan ingress-node-replication
set routing-instances evpn1 bridge-domains vlan101 domain-type bridge
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set routing-instances evpn1 bridge-domains vlan101 vxlan ingress-node-replication
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 100 family inet address 10.200.100.251/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::251/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 virtual-gateway-accept-data
set interfaces irb unit 101 family inet address 10.200.101.251/24 virtual-gateway-
address 10.200.101.254
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::251/64 virtual-
gateway-address 2001:0db8:0200:101::254
```


© 2020 Juniper Networks, Inc. All rights reserved.

vQFX1:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65001
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.17.7 peer-as 65007
set protocols bgp group underlay neighbor 10.200.18.8 peer-as 65008
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.7
set protocols bgp group overlay neighbor 10.200.0.8
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 10.200.0.1:1
set switch-options vrf-import fabric
set switch-options vrf-target target:65000:1
set protocols evpn vni-options vni 100 vrf-target target:65000:100
set protocols evpn vni-options vni 101 vrf-target target:65000:101
set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list [ 100 101 ]
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set vlans server vxlan vni 101
set vlans user vxlan vni 100
```

vQFX2:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65002
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.27.7 peer-as 65007
set protocols bgp group underlay neighbor 10.200.28.8 peer-as 65008
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group overlay neighbor 10.200.0.7
set protocols bgp group overlay neighbor 10.200.0.8
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 10.200.0.2:1
set switch-options vrf-import fabric
set switch-options vrf-target target:65000:1
set protocols evpn vni-options vni 100 vrf-target target:65000:100
set protocols evpn vni-options vni 101 vrf-target target:65000:101
set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list [ 100 101 ]
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set vlans server vxlan vni 101
set vlans user vxlan vni 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 7: Class of Service

This task is focused on Class of Services. You will configure and monitor MF and BA Classifiers, Two-color Policers, different Forwarding classes with respected Queues and Schedulers, DSCP Rewrite markers, RED Drop profiles and Class based forwarding.

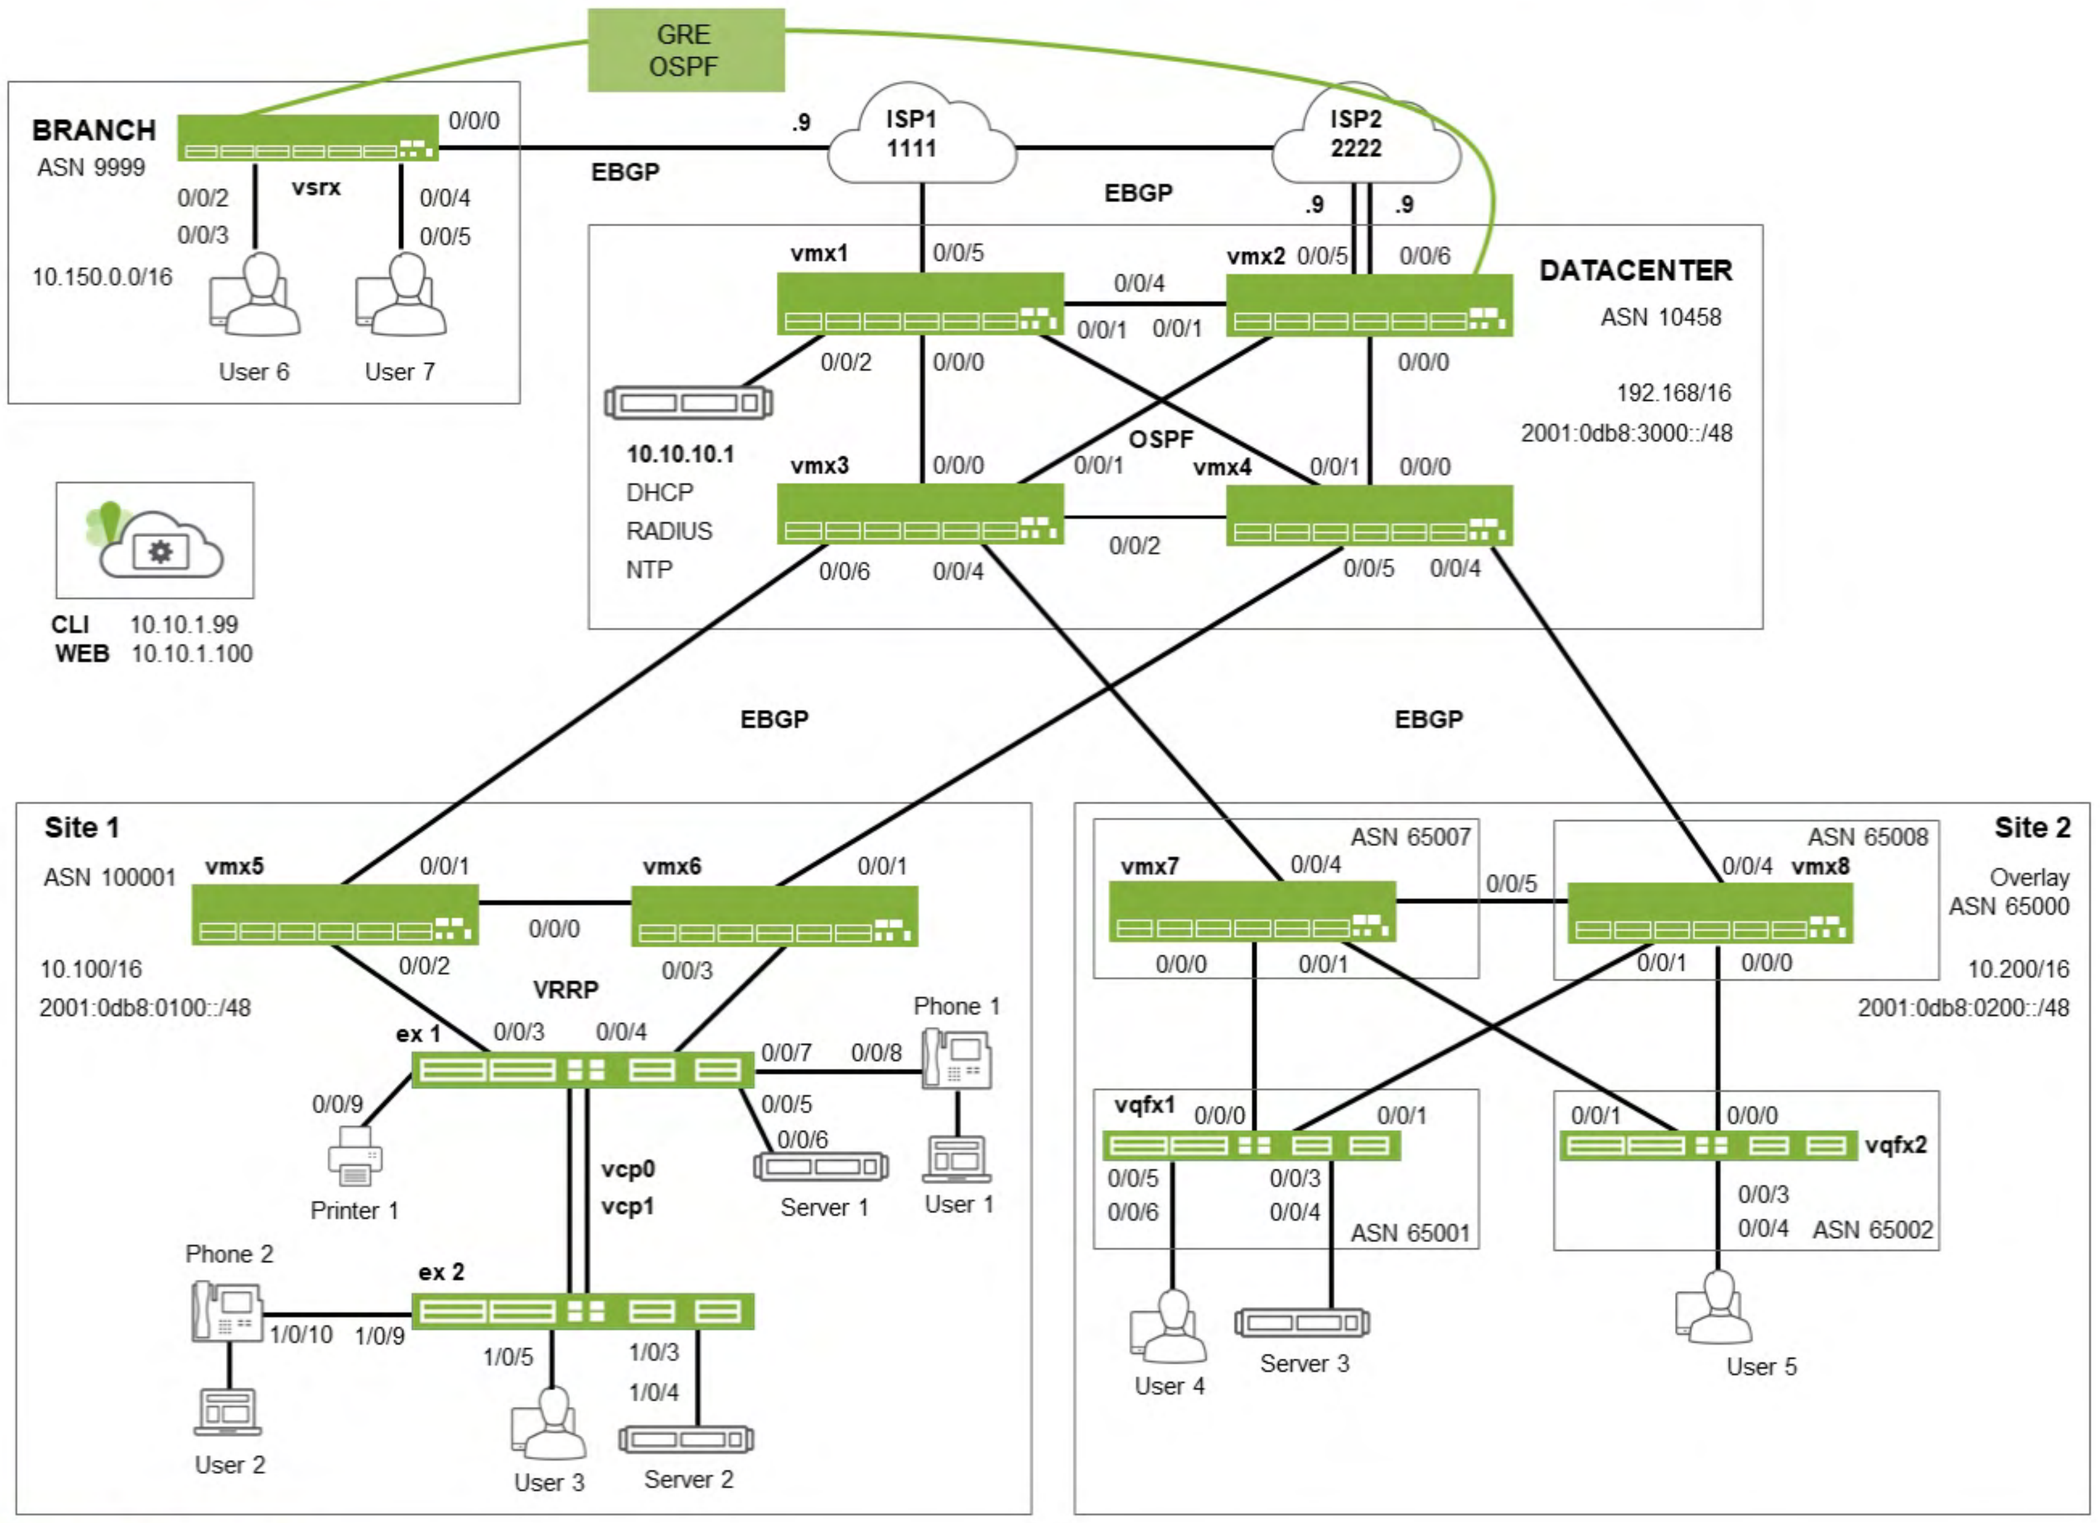

Figure 7 topology

vMX Class of Service

In this section you will configure all vMX in the datacenter to support Class of Service.

1) Configure Forwarding Classes and map them to the output Queues as indicated in the table below.

© 2020 Juniper Networks, Inc. All rights reserved.

Forwarding Class	Queue	Scheduler
control	7	control-sched
voice	5	voice-sched
server-gold	3	server-gold-sched
server-silver	2	server-silver-sched
server-bronze	1	server-bronze-sched
best-effort	0	best-effort-sched

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service forwarding-classes
queue 0 best-effort;
queue 1 server-bronze;
queue 2 server-silver;
queue 3 server-gold;
queue 5 voice;
queue 7 control;
```

2) Configure schedulers with the parameters in Table 7.

Scheduler	Parameter	Value
control-sched	priority	medium-high
	transmit rate	5%
	buffer-size	10%
voice-sched	priority	strict-high
	transmit rate	10%
	buffer-size	20 (milliseconds)
server-gold-sched	priority	high
	transmit rate	25%
	buffer-size	10% exact
server-silver-sched	priority	medium-high
	transmit rate	35%
	buffer-size	50% exact
	drop profile high	relaxed
server-bronze-sched	priority	medium-low
	transmit rate	remainder
	buffer-size	remainder
	drop profile low	relaxed
	drop profile high	aggressive

© 2020 Juniper Networks, Inc. All rights reserved.

best-effort-sched	priority	low
	transmit rate	remainder
	buffer-size	remainder
	drop profile any	aggressive

Table 43

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service schedulers
control-sched {
    transmit-rate percent 5;
    buffer-size percent 10;
    priority medium-high;
}
voice-sched {
    transmit-rate percent 10;
    buffer-size temporal 20k;
    priority strict-high;
}
server-gold-sched {
    transmit-rate percent 25;
    priority high;
}
server-silver-sched {
    transmit-rate percent 35;
    priority medium-high;
    drop-profile-map loss-priority high protocol any drop-profile relaxed;
}
server-bronze-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority medium-low;
    drop-profile-map loss-priority low protocol any drop-profile relaxed;
    drop-profile-map loss-priority high protocol any drop-profile aggressive;
}
best-effort-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile aggressive;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

3) Maps

Forwarding Class	Scheduler
control	control-sched
voice	voice-sched
server-gold	server-gold-sched
server-silver	server-silver-sched
server-bronze	server-bronze-sched
best-effort	best-effort-sched

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service scheduler-maps
cos-sched-map {
    forwarding-class control scheduler control-sched;
    forwarding-class voice scheduler voice-sched;
    forwarding-class server-gold scheduler server-gold-sched;
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
```

4) Configure two interpolated Drop profiles listed in Table 8.

Drop profile	Fill level (%)	Drop probability (%)
relaxed	75	0
	90	75
	100	100
aggressive	50	0
	75	50
	90	75
	100	100

Table 44

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service drop-profiles
relaxed {
    interpolate {
        fill-level [ 75 90 100 ];
        drop-probability [ 0 75 100 ];
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
aggressive {
    interpolate {
        fill-level [ 50 75 90 100 ];
        drop-probability [ 0 50 75 100 ];
    }
}
```

5) Apply the schedulers to all physical gigabit Ethernet interfaces.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
}
```

6) Configure IPv4 and IPv6 Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	low
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service classifiers
dscp ba-classifier {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
dscp-ipv6 ba-classifier-v6 {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
```

7) Apply the classifiers to all (logical) interfaces.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
```

8) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	DSCP value	Loss-priority
control	110000	low
	111000	high
voice	101110	low
server-gold	100010	low
server-silver	010010	low
	010110	high
server-bronze	001010	low
	001110	high
best-effort	000000	high

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show class-of-service rewrite-rules
dscp rewrite-v4 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
dscp-ipv6 rewrite-v6 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
forwarding-class voice {
    loss-priority low code-point ef;
}
forwarding-class server-gold {
    loss-priority low code-point af41;
}
forwarding-class server-silver {
    loss-priority low code-point af21;
    loss-priority high code-point af23;
}
forwarding-class server-bronze {
    loss-priority low code-point af11;
    loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
```

9) Apply the rewrite rules on all logical interfaces.

vMX1, vMX2, vMX3 and vM4:

```
lab@vmx1# show class-of-service interfaces
ge-* {
    scheduler-map cos-sched-map;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
```

Verification:

```
lab@vmx1# run show class-of-service interface ge-0/0/0
Physical interface: ge-0/0/0, Index: 148
Maximum usable queues: 8, Queues in use: 6
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/0.0, Index: 333
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Object                  Name                  Type                  Index
Rewrite-Output          rewrite-v4            dscp                   4683
Rewrite-Output          rewrite-v6            dscp-ipv6              4687
Classifier              ba-classifier         dscp                  19373
Classifier              ba-classifier-v6      dscp-ipv6             29019

[edit]
lab@vmx1# run show class-of-service interface ge-0/0/1
Physical interface: ge-0/0/1, Index: 149
Maximum usable queues: 8, Queues in use: 6
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/1.0, Index: 334
Object                  Name                  Type                  Index
Rewrite-Output          rewrite-v4            dscp                   4683
Rewrite-Output          rewrite-v6            dscp-ipv6              4687
Classifier              ba-classifier         dscp                  19373
Classifier              ba-classifier-v6      dscp-ipv6             29019

[edit]
lab@vmx1# run show class-of-service interface ge-0/0/2
Physical interface: ge-0/0/2, Index: 150
Maximum usable queues: 8, Queues in use: 6
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/2.0, Index: 335
Object                  Name                  Type                  Index
Rewrite-Output          rewrite-v4            dscp                   4683
Classifier              ba-classifier         dscp                  19373

[edit]
lab@vmx1# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 152
Maximum usable queues: 8, Queues in use: 6
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/4.0, Index: 336
Object                  Name                  Type                  Index
Rewrite-Output          rewrite-v4            dscp                   4683
Rewrite-Output          rewrite-v6            dscp-ipv6              4687
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Classifier                ba-classifier          dscp                    19373
Classifier                ba-classifier-v6       dscp-ipv6               29019

[edit]
lab@vmx1# run show class-of-service interface ge-0/0/5
Physical interface: ge-0/0/5, Index: 153
Maximum usable queues: 8, Queues in use: 6
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: cos-sched-map, Index: 8778
  Congestion-notification: Disabled

  Logical interface: ge-0/0/5.0, Index: 337
Object                    Name                   Type                    Index
Rewrite-Output            rewrite-v4             dscp                     4683
Rewrite-Output            rewrite-v6             dscp-ipv6                4687
Classifier                ba-classifier          dscp                    19373
Classifier                ba-classifier-v6       dscp-ipv6               29019
```

Note: vMX2, vMX3, and vMX4 should provide similar CoS outputs.

Configurations for Task Completion:

vMX1, vMX2, vMX3, and vMX4:

```
set class-of-service forwarding-classes queue 0 best-effort
set class-of-service forwarding-classes queue 1 server-bronze
set class-of-service forwarding-classes queue 2 server-silver
set class-of-service forwarding-classes queue 3 server-gold
set class-of-service forwarding-classes queue 5 voice
set class-of-service forwarding-classes queue 7 control
set class-of-service schedulers control-sched transmit-rate percent 5
set class-of-service schedulers control-sched buffer-size percent 10
set class-of-service schedulers control-sched priority medium-high
set class-of-service schedulers voice-sched transmit-rate percent 10
set class-of-service schedulers voice-sched buffer-size temporal 20k
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers server-gold-sched transmit-rate percent 25
set class-of-service schedulers server-gold-sched priority high
set class-of-service schedulers server-silver-sched transmit-rate percent 35
set class-of-service schedulers server-silver-sched priority medium-high
set class-of-service schedulers server-silver-sched drop-profile-map loss-priority high
protocol any drop-profile relaxed
set class-of-service schedulers server-bronze-sched transmit-rate remainder
set class-of-service schedulers server-bronze-sched buffer-size remainder
set class-of-service schedulers server-bronze-sched priority medium-low
set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority low
protocol any drop-profile relaxed
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service schedulers server-bronze-sched drop-profile-map loss-priority high
protocol any drop-profile aggressive
set class-of-service schedulers best-effort-sched transmit-rate remainder
set class-of-service schedulers best-effort-sched buffer-size remainder
set class-of-service schedulers best-effort-sched priority low
set class-of-service schedulers best-effort-sched drop-profile-map loss-priority any
protocol any drop-profile aggressive
set class-of-service scheduler-maps cos-sched-map forwarding-class control scheduler
control-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class voice scheduler
voice-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-gold scheduler
server-gold-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map forwarding-class best-effort scheduler
best-effort-sched
set class-of-service drop-profiles relaxed interpolate fill-level [ 75 90 100 ]
set class-of-service drop-profiles relaxed interpolate drop-probability [ 0 75 100 ]
set class-of-service drop-profiles aggressive interpolate fill-level [ 50 75 90 100 ]
set class-of-service drop-profiles aggressive interpolate drop-probability [ 0 50 75 100
]
set class-of-service interfaces ge-* scheduler-map cos-sched-map
set class-of-service classifiers dscp ba-classifier import default
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc1
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc2
set class-of-service classifiers dscp ba-classifier forwarding-class voice loss-priority
low code-points ef
set class-of-service classifiers dscp ba-classifier forwarding-class server-gold loss-
priority low code-points af41
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority low code-points af21
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority high code-points af23
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority low code-points af11
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority high code-points af13
set class-of-service classifiers dscp ba-classifier forwarding-class best-effort loss-
priority low code-points be
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 import default
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc1
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class voice loss-
priority low code-points ef
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-gold
loss-priority low code-points af41
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority low code-points af21
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority high code-points af23
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority low code-points af11
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority high code-points af13
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class best-effort
loss-priority low code-points be
set class-of-service interfaces ge-* unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-* unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class voice loss-priority
low code-point ef
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-gold loss-
priority low code-point af41
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority low code-point af21
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority high code-point af23
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority low code-point af11
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority high code-point af13
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class best-effort loss-
priority high code-point be
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class voice loss-
priority low code-point ef
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-gold
loss-priority low code-point af41
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority low code-point af21
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority high code-point af23
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority low code-point af11
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority high code-point af13
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class best-effort
loss-priority high code-point be
set class-of-service interfaces ge-* unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-* unit * rewrite-rules dscp-ipv6 rewrite-v6
```

This completes Super Lab 1.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Super Lab 2

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

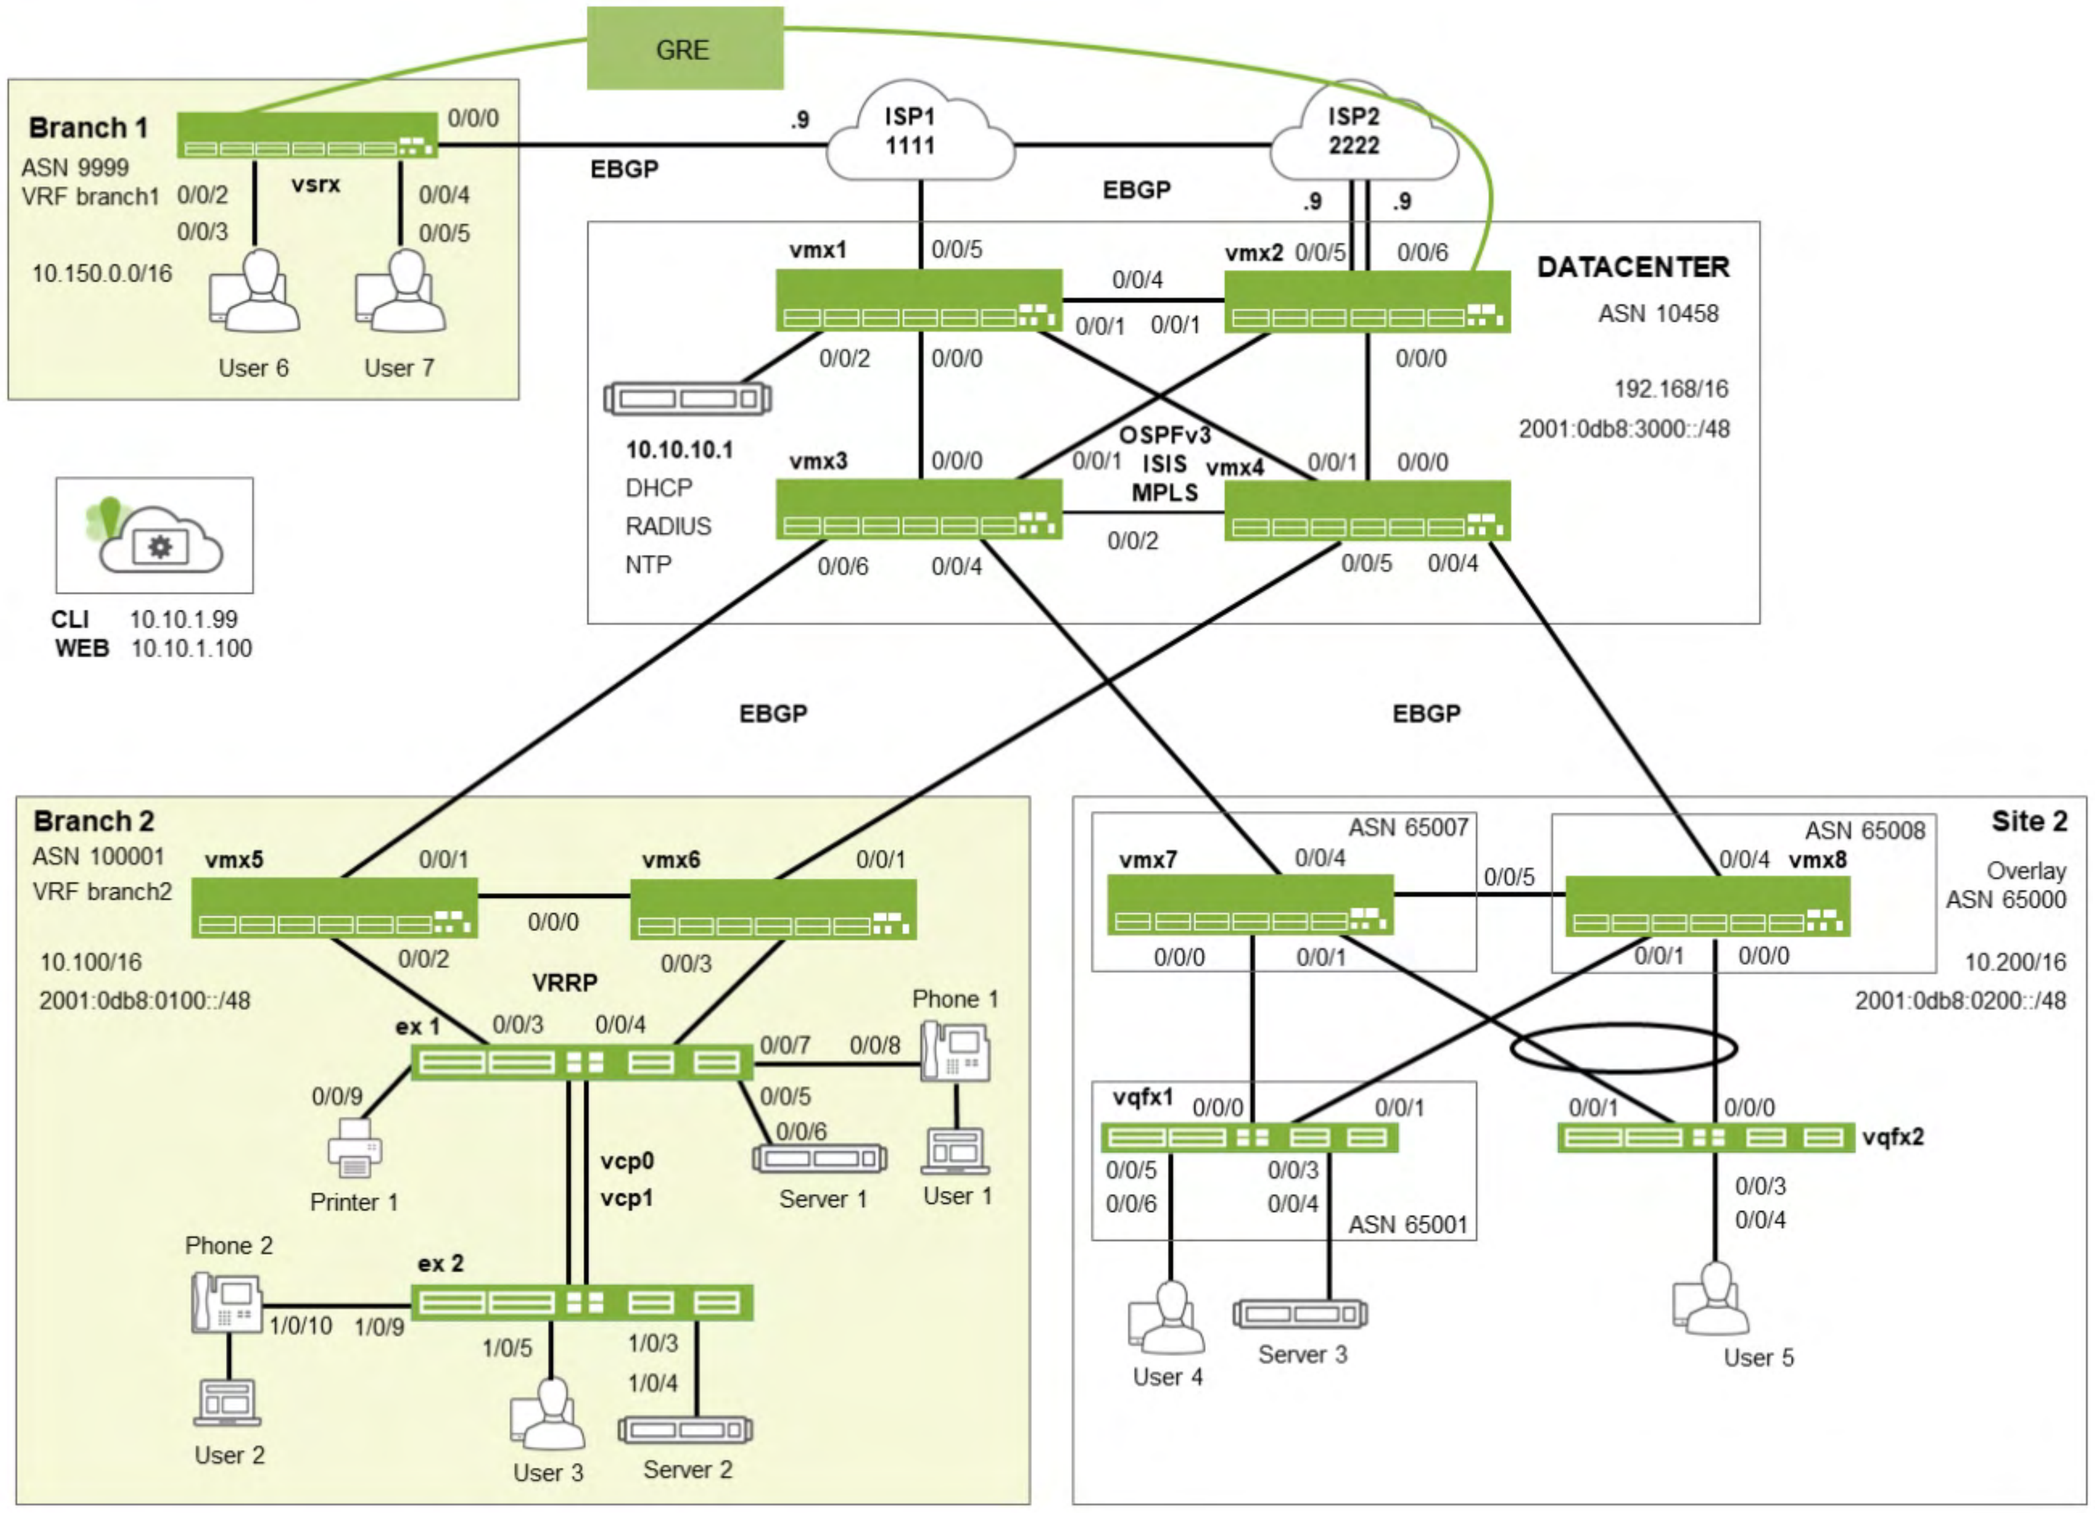

Task 1: General System Features

System Services

© 2020 Juniper Networks, Inc. All rights reserved.

Note: For this section you can verify your solutions from vMX1 as this device is directly connected to the 10.10.10/24 network. The task solutions for the other devices in the topology require that routing to the 10.10.10/24 network is in place.

1) Configure all vMX and EX1 devices to not allow more than 4 simultaneous SSH and Telnet connections and no more than 10 connections per minute. Maximum 32 sessions are allowed per SSH connection.

vMXs and EX1:

```
lab@vmx3# show system services
ssh {
    root-login allow;
    max-sessions-per-connection 32;
    connection-limit 4;
    rate-limit 10;
}
telnet {
    connection-limit 4;
    rate-limit 10;
}
```

2) Enable the FTP service

vMXs and EX1:

```
lab@vmx1# show system services
ftp;
```

3) Set the time zone to *America/Los_Angeles* on all vMX devices, EX1

vMXs and EX1:

```
lab@vmx1# show system time-zone
time-zone America/Los_Angeles;
```

4) Make sure that all vMX devices, EX1 synchronize the time with NTP server 10.10.10.1.

NOTE: *There is a dedicated FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server IP address is 10.10.10.1. Ensure to use the device loopback address as the source address when communicating with server 10.10.10.1*

vMX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show system ntp
server 10.10.10.1;
source-address 192.168.0.1;
```

vMX2:

```
lab@vmx2# show system ntp
server 10.10.10.1;
source-address 192.168.0.2;
```

vMX3:

```
lab@vmx3# show system ntp
server 10.10.10.1;
source-address 192.168.0.3;
```

vMX4:

```
lab@vmx4# show system ntp
server 10.10.10.1;
source-address 192.168.0.4;
```

vMX5:

```
lab@vmx5# show system ntp
server 10.10.10.1;
source-address 10.100.0.5;
```

vMX6:

```
lab@vmx6# show system ntp
server 10.10.10.1;
source-address 10.100.0.6;
```

EX1:

```
lab@ex1# show system ntp
server 10.10.10.1;
source-address 10.100.0.1;
```

vMX7:

```
lab@vmx7# show system ntp
server 10.10.10.1;
source-address 10.200.0.7;
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX8:

```
lab@vmx8# show system ntp
server 10.10.10.1;
source-address 10.200.0.8;
```

User Authentication and Authorization

In this section you will configure new users allowed to access the devices and define their privileges and permissions.

The following tasks should be configured on all vMX devices and EX1

5) Configure the authentication method that first the device tries to authenticate users on the RADIUS server and, if not successful, it uses a local password. Use server 10.10.10.1 as the RADIUS server. Configure the RADIUS server with retry attempts 3 and timeout of 3 seconds. Use *juniper123* as the RADIUS shared secret.

vMX1:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.1;
    }
}
```

vMX2:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.2;
    }
}
```

vMX3:

```
lab@vmx1# show system
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.3;
    }
}
```

vMX4:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 192.168.0.4;
    }
}
```

vMX5:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.5;
    }
}
```

vMX6:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.6;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX7:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.200.0.7;
    }
}
```

vMX8:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.200.0.8;
    }
}
```

EX1:

```
lab@vmx1# show system
authentication-order [ radius password ];
radius-server {
    10.10.10.1 {
        secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"; ## SECRET-DATA
        timeout 3;
        retry 3;
        source-address 10.100.0.1;
    }
}
```

6) User: admin password: admin123 is preconfigured on the RADIUS server. Ensure that user admin can login on all vMX devices and EX1 as a super-user.

Configure the following on all vMXs and EX1:

```
lab@vmx1# show system login user admin
uid 2003;
class super-user;
```


© 2020 Juniper Networks, Inc. All rights reserved.

7) Configure the following additional users on the devices as defined in the table below:

Username	Password	Privileges
ronly	ronly123	Has permissions “view” and “view-configuration”. Additionally, the user can NOT execute the “file delete” command.
admin1	admin123	Has permissions “all”. Can access the configuration mode only using “configure private” command.
restricted	restricted123	Has permissions “clear” and execute only the “show system uptime”, “show system storage” and “show interfaces terse” commands and nothing else.

Configure the following on all vMXs and EX1:

```
lab@vmx1# show system login
class admin-class {
    permissions all;
    allow-commands "configure private";
    deny-commands configure;
}
class restricted-class {
    permissions clear;
    allow-commands "(show system uptime)|(show system storage)|(show interfaces terse)";
}
class ronly-users {
    permissions [ view view-configuration ];
    deny-commands "file delete";
}
user admin {
    uid 2003;
    class super-user;
}
user admin1 {
    uid 2005;
    class admin-class;
    authentication {
        encrypted-password "$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."; ## SECRET-DATA
    }
}
user lab {
    uid 2001;
    class super-user;
    authentication {
        encrypted-password "$1$h3xBz7mh$mMDaYNboS7YubgYbC3KgM1"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
user restricted {
    uid 2006;
    class restricted-class;
    authentication {
        encrypted-password "$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"; ## SECRET-DATA
    }
}
user ronly {
    uid 2002;
    class ronly-users;
    authentication {
        encrypted-password "$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"; ## SECRET-DATA
    }
}
```

SNMP Configuration

In this section you will configure SNMPv2.

The following tasks should be configured on all vMX, vQFX and vSRX devices in the topology.

8) The SNMP read-only community string is jncie. Make sure that only the Junos Space server can poll the devices in the data center, Branch2, Site2, and the branch device.
9) Configure all devices to send SNMP traps to Junos Space. Make sure that traps are only send using the OoB management interface.

vMX1:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.1;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.2;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX3:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.3;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX4:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.4;
}
trap-group traps {
    targets {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        10.10.1.99;
    }
}
```

vMX5:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.5;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX6:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.6;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX7:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
trap-options {
    source-address 10.10.1.7;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vMX8:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.8;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vSRX:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.10;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vQFX1:

```
lab@vmx1# show snmp
client-list list {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.13;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

vQFX2:

```
lab@vmx1# show snmp
client-list list {
    10.10.1.99/32;
}
community jncie {
    client-list-name list;
}
trap-options {
    source-address 10.10.1.14;
}
trap-group traps {
    targets {
        10.10.1.99;
    }
}
```

Firewall Filters

In this section you will configure RE (Routing Engine) protection using a firewall filter for all vMX devices in Branch 2.

10) Configure an IPv6 firewall filter which allows protocol messages from VRRP and BGP. BGP peers that are configured in the future must be automatically allowed in the firewall filter.
11) Configure the devices to accept IPv6 ICMP and traceroute messages. Ensure that the flow of the messages is limited using a policer which has the following parameters: bandwidth limit 100kbps and burst size 25kB. The excess traffic must be dropped.

© 2020 Juniper Networks, Inc. All rights reserved.

12) Configure the IPv6 firewall filter to discard any other traffic and count the packets.
13) Apply the IPv6 firewall filter such as to ensure that it is used for the RE protection.
14) Ensure that the configuration under firewall hierarchy is not displayed when printed out.

vMX5 & vMX6:

```
lab@vmx5# show firewall | display omit
family inet6 {
    filter protect-re-ipv6 {
        apply-flags omit;
        term bgp {
    from {
        prefix-list {
            bgp-peers;
        }
        next-header tcp;
        port bgp;
    }
    then accept;
}
term vrrp {
    from {
        next-header vrrp;
    }
    then accept;
}
term icmp {
    from {
        next-header icmp6;
    }
    then {
        policer re-policer;
        accept;
    }
}
term traceroute {
    from {
        next-header udp;
        destination-port 33434-33534;
    }
    then {
        policer re-policer;
        accept;
    }
}
term last {
    then {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                count dropped-ipv6;
                discard;
            }
        }
    }
}
policer re-policer {
    apply-flags omit;
    if-exceeding {
        bandwidth-limit 100k;
        burst-size-limit 25k;
    }
    then discard;
}
```

```
lab@vmx5# show policy-options
prefix-list bgp-peers {
    apply-path "protocols bgp group <*> neighbor <*:*>";
}
```

```
lab@vmx5# show interfaces lo0
unit 0 {
...
    family inet6 {
        filter {
            input protect-re-ipv6;
        }

    }
}
```

Configurations for Task Completion:

vMX1:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.1
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system radius-server 10.10.10.1 source-address 192.168.0.1
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.1
set snmp trap-group traps targets 10.10.1.99
```

vMX2:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.2
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.2
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.2
set snmp trap-group traps targets 10.10.1.99
```

vMX3:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.3
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.3
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.3
set snmp trap-group traps targets 10.10.1.99
```

vMX4:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 192.168.0.4
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 192.168.0.4
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.4
set snmp trap-group traps targets 10.10.1.99
```

vMX5:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.5
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.5
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.5
set snmp trap-group traps targets 10.10.1.99
set firewall family inet6 filter protect-re-ipv6 apply-flags omit
set firewall family inet6 filter protect-re-ipv6 term bgp from prefix-list bgp-peers
set firewall family inet6 filter protect-re-ipv6 term bgp from next-header tcp
set firewall family inet6 filter protect-re-ipv6 term bgp from port bgp
set firewall family inet6 filter protect-re-ipv6 term bgp then accept
set firewall family inet6 filter protect-re-ipv6 term vrrp from next-header vrrp
set firewall family inet6 filter protect-re-ipv6 term vrrp then accept
set firewall family inet6 filter protect-re-ipv6 term icmp from next-header icmp6
set firewall family inet6 filter protect-re-ipv6 term icmp then policer re-policer
set firewall family inet6 filter protect-re-ipv6 term icmp then accept
set firewall family inet6 filter protect-re-ipv6 term traceroute from next-header udp
set firewall family inet6 filter protect-re-ipv6 term traceroute from destination-port
33434-33534
set firewall family inet6 filter protect-re-ipv6 term traceroute then policer re-policer
set firewall family inet6 filter protect-re-ipv6 term traceroute then accept
set firewall family inet6 filter protect-re-ipv6 term last then count dropped-ipv6
set firewall family inet6 filter protect-re-ipv6 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*:*>"
set interfaces lo0 unit 0 family inet6 filter input protect-re-ipv6
```

vMX6:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.6
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.6
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.6
set snmp trap-group traps targets 10.10.1.99
set firewall family inet6 filter protect-re-ipv6 apply-flags omit
set firewall family inet6 filter protect-re-ipv6 term bgp from prefix-list bgp-peers
set firewall family inet6 filter protect-re-ipv6 term bgp from next-header tcp
set firewall family inet6 filter protect-re-ipv6 term bgp from port bgp
set firewall family inet6 filter protect-re-ipv6 term bgp then accept
set firewall family inet6 filter protect-re-ipv6 term vrrp from next-header vrrp
set firewall family inet6 filter protect-re-ipv6 term vrrp then accept
set firewall family inet6 filter protect-re-ipv6 term icmp from next-header icmp6
set firewall family inet6 filter protect-re-ipv6 term icmp then policer re-policer
set firewall family inet6 filter protect-re-ipv6 term icmp then accept
set firewall family inet6 filter protect-re-ipv6 term traceroute from next-header udp
set firewall family inet6 filter protect-re-ipv6 term traceroute from destination-port
33434-33534
set firewall family inet6 filter protect-re-ipv6 term traceroute then policer re-policer
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set firewall family inet6 filter protect-re-ipv6 term traceroute then accept
set firewall family inet6 filter protect-re-ipv6 term last then count dropped-ipv6
set firewall family inet6 filter protect-re-ipv6 term last then discard
set firewall policer re-policer apply-flags omit
set firewall policer re-policer if-exceeding bandwidth-limit 100k
set firewall policer re-policer if-exceeding burst-size-limit 25k
set firewall policer re-policer then discard
set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*> neighbor
<*:*>"
set interfaces lo0 unit 0 family inet6 filter input protect-re-ipv6
```

vMX7:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.7
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.7
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.7
set snmp trap-group traps targets 10.10.1.99
```

vMX8:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.200.0.8
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.200.0.8
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set snmp trap-options source-address 10.10.1.8
set snmp trap-group traps targets 10.10.1.99
```

EX1:

```
set system services ssh root-login allow
set system services ssh max-sessions-per-connection 32
set system services ssh connection-limit 4
set system services ssh rate-limit 10
set system services telnet connection-limit 4
set system services telnet rate-limit 10
set system services ftp
set system time-zone America/Los_Angeles
set system ntp server 10.10.10.1
set system ntp source-address 10.100.0.1
set system authentication-order [ radius password ]
set system radius-server 10.10.10.1 secret "$9$qPT3ApBSrv69rvWLVb.P5Q69tuORcy"
set system radius-server 10.10.10.1 timeout 3
set system radius-server 10.10.10.1 retry 3
set system radius-server 10.10.10.1 source-address 10.100.0.1
set system login user admin uid 2003
set system login user admin class super-user
set system login class admin-class permissions all
set system login class admin-class allow-commands "configure private"
set system login class admin-class deny-commands configure
set system login class restricted-class permissions clear
set system login class restricted-class allow-commands "(show system uptime)|(show
system storage)|(show interfaces terse)"
set system login class ronly-users permissions view
set system login class ronly-users permissions view-configuration
set system login class ronly-users deny-commands "file delete"
set system login user admin1 uid 2005
set system login user admin1 class admin-class
set system login user admin1 authentication encrypted-password
"$1$uE9mPnHr$ezot4lR9QaUyjzItj.p.5."
set system login user restricted uid 2006
set system login user restricted class restricted-class
set system login user restricted authentication encrypted-password
"$1$XeNFKAWG$nAZBvetrHZk31/9bU6b/11"
set system login user ronly uid 2002
set system login user ronly class ronly-users
set system login user ronly authentication encrypted-password
"$1$cuRZtU4h$cnMqTHxNea5FSWIXcHbCq0"
```

vQFX1:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.13
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set snmp trap-group traps targets 10.10.1.99
```

vQFX2:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.14
set snmp trap-group traps targets 10.10.1.99
```

vSRX:

```
set snmp client-list list 10.10.1.99/32
set snmp community jncie client-list-name list
set snmp trap-options source-address 10.10.1.10
set snmp trap-group traps targets 10.10.1.99
```

Task 2: Layer 2 Configuration and Features

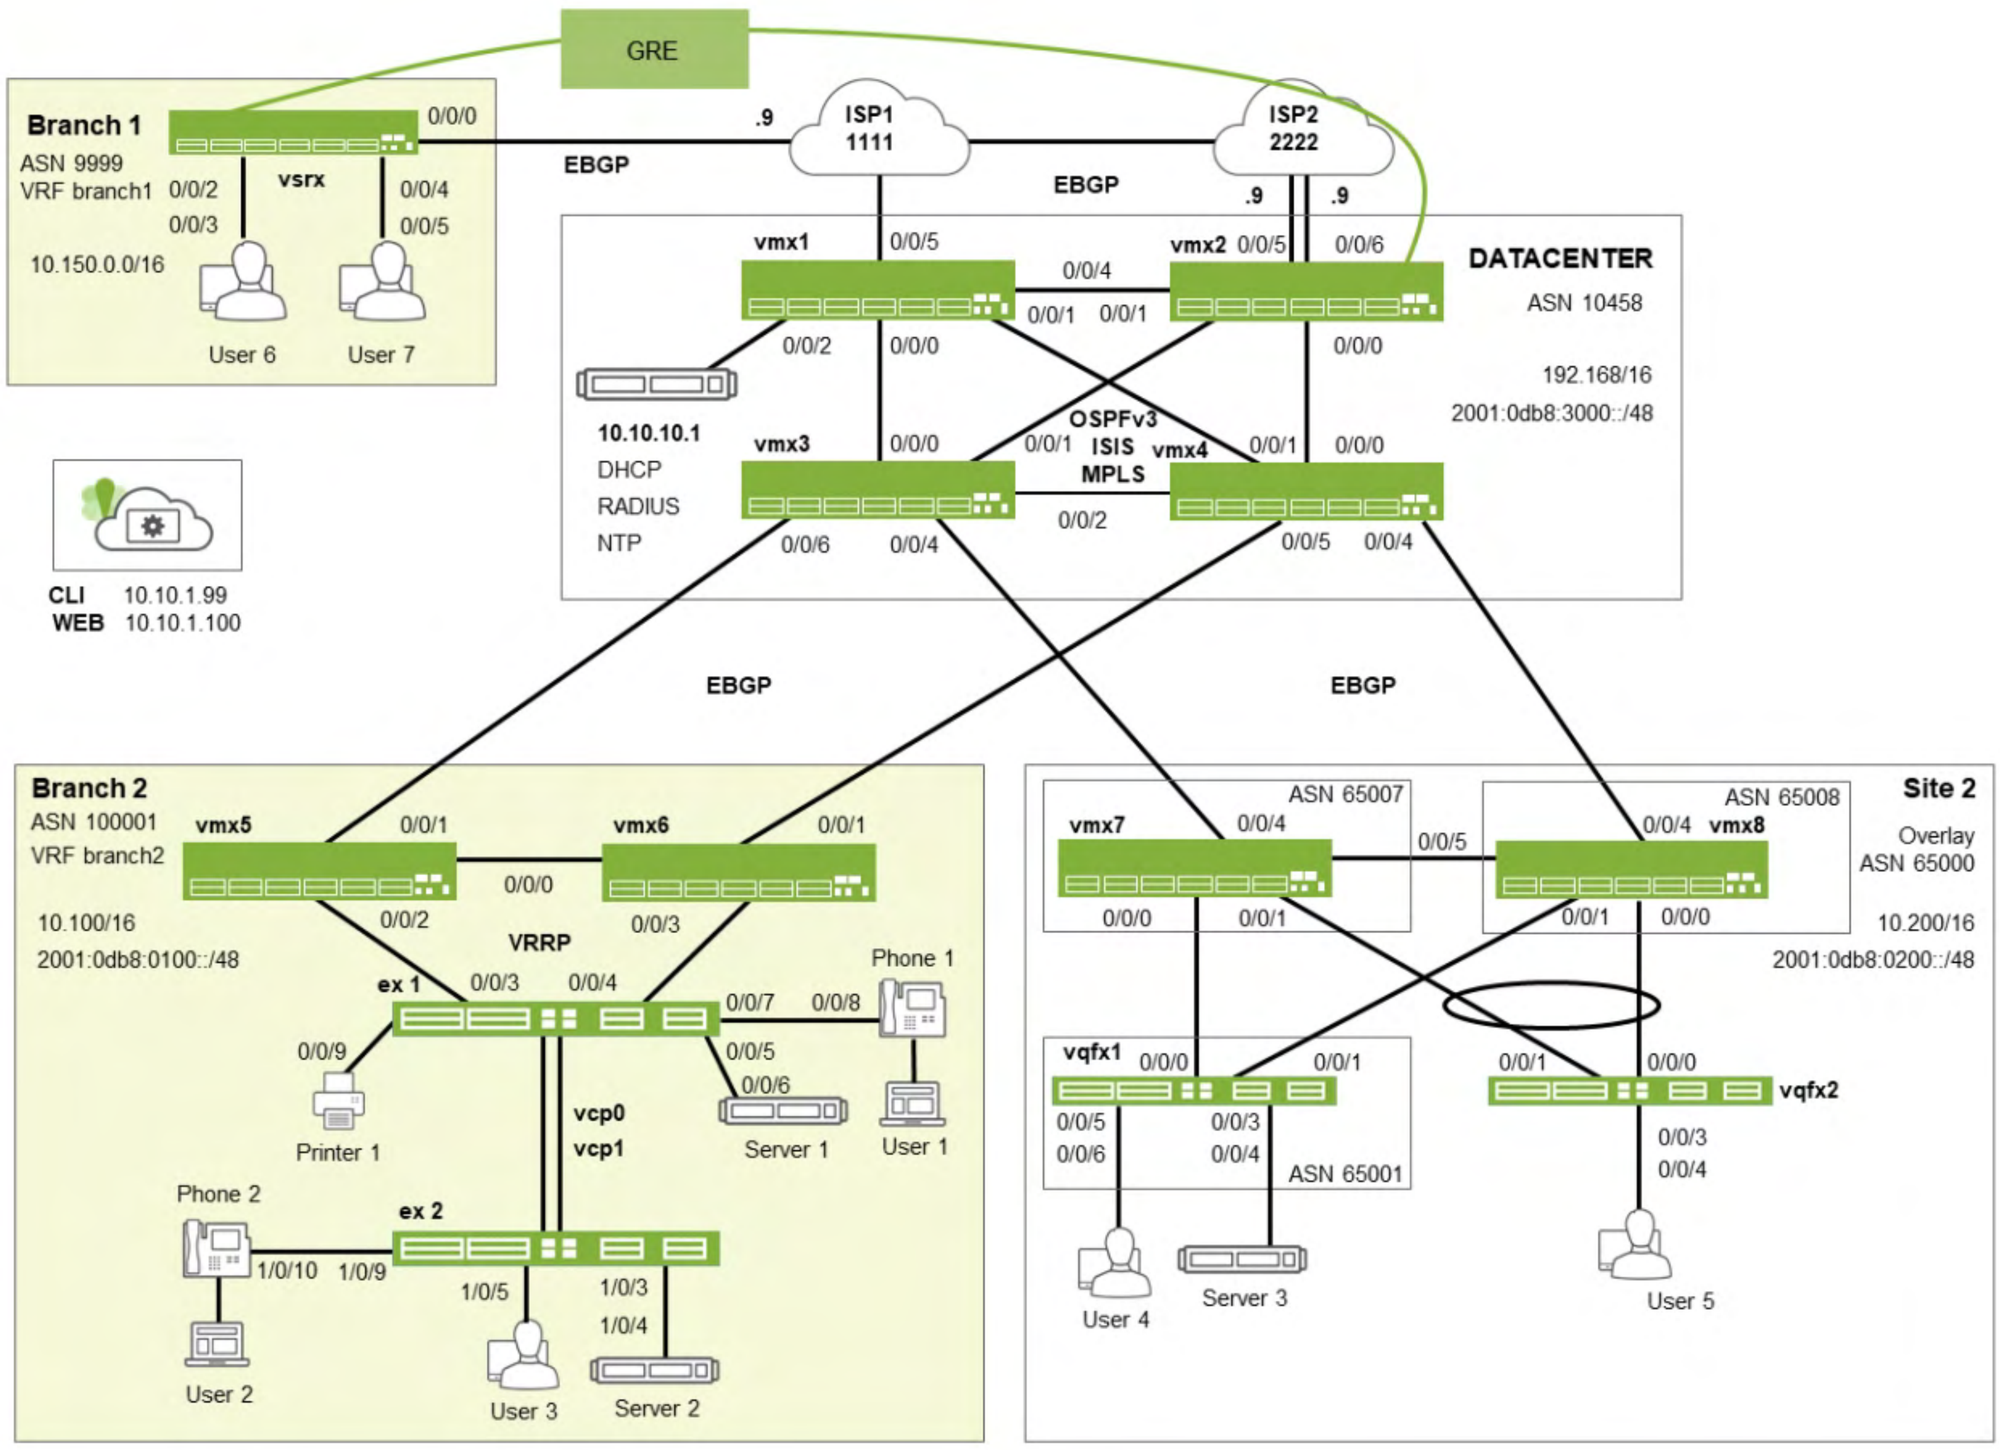

© 2020 Juniper Networks, Inc. All rights reserved.

Virtual Chassis

In this section you will configure a Virtual chassis consisting of EX1 and EX2.

1) Configure EX1 and EX2 as a virtual chassis. Ensure EX1 is FP0 and EX2 is FPC1. Ensure the hostname for the virtual-chassis is EX1
2) Configure Juniper recommended priority settings for a two-member VC.
3) Ensure that the master device retains mastership by not taking the linecard role in case the backup device fails.

EX1:

```
lab@ex1> request virtual-chassis vc-port set interface vcp-255/1/0 member 0

lab@ex1> request virtual-chassis vc-port set interface vcp-255/1/1 member 0


lab@ex1# show virtual-chassis
no-split-detection;
traceoptions {
    file me;
    flag me detail;
}
member 0 {
    mastership-priority 255;
}
member 1 {
    mastership-priority 255;
}
```

EX2:

```
lab@ex2> request virtual-chassis vc-port set interface vcp-255/1/0

lab@ex2> request virtual-chassis vc-port set interface vcp-255/1/1
```

Verify on EX1:

```
lab@ex1# run show virtual-chassis vc-port
fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
1/0         Configured          5    Up           40000        1   vcp-255/1/0
1/1         Configured          5    Up           40000        1   vcp-255/1/1
1/3         Configured               Absent
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
1/2          Configured                Absent

fpc1:
--------------------------------------------------------------------------
Interface    Type                 Trunk  Status       Speed        Neighbor
or                                 ID                 (mbps)       ID  Interface
PIC / Port
1/0          Configured             5    Up           40000        0   vcp-255/1/0
1/1          Configured             5    Up           40000        0   vcp-255/1/1
1/3          Configured                  Absent
1/2          Configured                  Absent
```

```
lab@ex1# run show virtual-chassis status

Virtual Chassis ID: 8f14.33f4.8988
Virtual Chassis Mode: Enabled
                                                  Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model            prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    PG3718460071 ex4300-24t       255   Master*      N  VC   1  vcp-
255/1/0
                                                                          1  vcp-
255/1/1
1 (FPC 1)  Prsnt    PG3718470055 ex4300-24t       255   Backup       N  VC   0  vcp-
255/1/0
vcp-255/1/1
```

Now that your Virtual Chassis has been created you need to add the interface configuration and routing instances for the interfaces that are connected to EX2 on the Virtual Chassis master (EX1).

```
lab@ex1# show interfaces | find ge-1
ge-1/0/3 {
    description "to server2";
}
ge-1/0/4 {
    description "server 2";
}
ge-1/0/5 {
    description "to user3";
}
ge-1/0/9 {
    description "to phone2 user2";
}
ge-1/0/10 {
    description "phone2 user2";
    unit 0 {
        family inet {
            address 10.100.100.2/24;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
        family inet6 {
            address 2001:0db8:0100:100::2/64;
        }
    }
}
...

lab@ex1# show routing-instances
server1 {
    instance-type virtual-router;
    interface ge-0/0/6.0;
    routing-options {
        rib server1.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:101::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.101.254;
        }
    }
}
server2 {
    instance-type virtual-router;
    interface ge-1/0/4.0;
    routing-options {
        rib server2.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:101::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.101.254;
        }
    }
}
user1 {
    instance-type virtual-router;
    interface ge-0/0/8.0;
    routing-options {
        rib user1.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:100::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.100.254;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
user2 {
    instance-type virtual-router;
    interface ge-1/0/10.0;
    routing-options {
        rib user2.inet6.0 {
            static {
                route 0::/0 next-hop 2001:0db8:0100:100::254;
            }
        }
        static {
            route 0.0.0.0/0 next-hop 10.100.100.254;
        }
    }
}
```

VLAN Configuration

In this section you will configure (VLAN) interfaces on the switches.

VLAN Name	VLAN ID	Devices	Interface
user	100	EX1	ge-0/0/7 untagged
			ge-0/0/5 untagged
			ge-1/0/9 untagged
		vQFX1	xe-0/0/5 untagged
		vQFX2	xe-0/0/3 untagged
server	101	EX1	ge-0/0/5 untagged
			ge-1/0/3 untagged
		vQFX1	xe-0/0/3 untagged
printer	200	EX1	ge-0/0/9 untagged
phone	201	EX1	ge-0/0/7 tagged
			ge-1/0/9 tagged
guest	666	EX1	Dot1x guest vlan
vmx5vmx6ex1	156	EX1	ge-0/0/3
			ge-0/0/4

Table VLANs Branch2 and Site 2

4) Configure the VLANs and corresponding interfaces in Branch2 and Site 2 as shown in the table above. Protect the vlans configuration hierarchy so that a warning message is displayed when deleted.

© 2020 Juniper Networks, Inc. All rights reserved.

EX1:

```
lab@ex1# show vlans
default {
    vlan-id 1;
}
guest {
    vlan-id 666;
}
phone {
    vlan-id 201;
    l3-interface irb.201;
}
printer {
    vlan-id 200;
    l3-interface irb.200;
}
server {
    vlan-id 101;
    l3-interface irb.101;
}
user {
    vlan-id 100;
    l3-interface irb.100;
    switch-options {
        interface ge-0/0/7.0;
        interface ge-1/0/9.0;
    }
}
vmx5vmx6ex1 {
    vlan-id 156;
    l3-interface irb.156;
}

lab@ex1# show interfaces
ge-0/0/3 {
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members vmx5vmx6ex1;
            }
        }
    }
}
ge-0/0/4 {
    unit 0 {
        family ethernet-switching {
            interface-mode access;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            vlan {
                members vmx5vmx6ex1;
            }
        }
    }
}
ge-0/0/5 {
    description "to server1";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members server;
            }
        }
    }
}
ge-0/0/7 {
    description "to phone1 user1";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members user;
            }
        }
    }
}
ge-0/0/9 {
    description "to printer1";
    unit 0 {
        family ethernet-switching {
            vlan {
                members printer;
            }
        }
    }
}
ge-1/0/3 {
    description "to server2";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members server;
            }
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-1/0/5 {
    description user3;
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members user;
            }
        }
    }
}
ge-1/0/9 {
    description "phone2 user2";
    unit 0 {
        family ethernet-switching {
            interface-mode access;
            vlan {
                members user;
            }
        }
    }
}
```

```
lab@ex1# protect vlans
```

vQFX1:

```
lab@vQFX1# show interfaces xe-0/0/3
description "to server3";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members server;
        }
    }
}

lab@vQFX1# show interfaces xe-0/0/5
description "to user4";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vQFX1# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
}
user {
    vlan-id 100;
}
```

vQFX2:

```
lab@vQFX2# show vlans
default {
    vlan-id 1;
}
user {
    vlan-id 100;
}

lab@vQFX2# show interfaces xe-0/0/3
description "to user5";
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members user;
        }
    }
}
```

Verification:

```
lab@ex1# show vlans
##
## protect: vlans
##
default {
    vlan-id 1;
}
guest {
    vlan-id 666;
}
phone {
    vlan-id 201;
    l3-interface irb.201;
}
printer {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    vlan-id 200;
    l3-interface irb.200;
}
server {
    vlan-id 101;
    l3-interface irb.101;
}
user {
    vlan-id 100;
    l3-interface irb.100;
    switch-options {
        interface ge-0/0/7.0;
        interface ge-1/0/9.0;
    }
}
vmx5vmx6ex1 {
    vlan-id 156;
    l3-interface irb.156;
}

lab@ex1# delete vlans
warning: [vlans] is protected, 'vlans' cannot be deleted
```

5) Configure interface ge-0/0/7 and interface ge-1/0/9 on virtual chassis EX1 for Voice VLAN support. Make sure that tagged frames go to VLAN201 whereas untagged frames go to VLAN100.

```
lab@ex1# show switch-options
voip {
    interface ge-0/0/7.0 {
        vlan phone;
    }
    interface ge-1/0/9.0 {
        vlan phone;
    }
}
```

6) Configure aggregate ethernet interface (AE15) on vMX7 and vMX8 towards vQFX2. Enable LACP on the AE15 interface. All Site 2 VLANs must be allowed on the AE interfaces. Configure interface AE0 on vQFX2 towards vMX7 and vMX8. Interface AE0 on vQFX2 should not initiate LACP.

vMX7:

```
lab@vmx7# show chassis
aggregated-devices {
    ethernet {
        device-count 16;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}

lab@vmx7# show interfaces ge-0/0/1
gigether-options {
    802.3ad ae15;
}

lab@vmx7# show interfaces ae15
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```

vMX8:

```
lab@vmx8# show chassis
aggregated-devices {
    ethernet {
        device-count 16;
    }
}

lab@vmx8# show interfaces ge-0/0/0
gigether-options {
    802.3ad ae15;
}

lab@vmx8# show interfaces ae15
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```

vQFX2:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vQFX2# show chassis
aggregated-devices {
    ethernet {
        device-count 1;
    }
}

lab@vQFX2# show interfaces xe-0/0/0
gigether-options {
    802.3ad ae0;
}

lab@vQFX2# show interfaces xe-0/0/1
gigether-options {
    802.3ad ae0;
}

lab@vQFX2# show interfaces ae0
aggregated-ether-options {
    lacp {
        passive;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members all;
        }
    }
}
```

Verification:

```
lab@vQFX2# run show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 670, SNMP ifIndex: 559
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, Ethernet-
Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 02:05:86:72:72:f0, Hardware address: 02:05:86:72:72:f0
  Last flapped   : 2019-04-15 13:23:18 UTC (2d 16:46 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 2224 bps (0 pps)

  Logical interface ae0.0 (Index 558) (SNMP ifIndex 560)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes           bps
    Bundle:
        Input :          6354          0        524724             0
        Output:        697393          2      79672550          2264
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol eth-switch, MTU: 1514
```

```
lab@vQFX2# run show lacp statistics interfaces
Aggregated interface: ae0
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    xe-0/0/0              232769      223179            0            0
    xe-0/0/1              232768      223179            0            0
```

7) Ensure the following MAC addresses are in the Ethernet switching table of EX1 for vlan vmx5vmx6ex1:

00:00:5e:00:01:0d
00:00:5e:00:02:63
00:56:56:56:56:05
00:56:56:56:56:06

Note: This task can only be verified after the implementation of the VRRP task.

vMX5:

```
lab@vmx5# show interfaces ge-0/0/2
mac 00:56:56:56:56:05;
unit 0 {
    family inet {
        address 10.100.156.5/24 {
            vrrp-group 13 {
                ...
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::5/64 {
            vrrp-inet6-group 99 {
                ...
            }
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3
mac 00:56:56:56:56:06;
unit 0 {
    family inet {
        address 10.100.156.6/24 {
            vrrp-group 13 {
...
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 99 {
...
            }
        }
    }
}
```

Note: *Keep in mind that the VRRP virtual gateway MAC address is based on the VRRP group ID. For this reason, the IPv4 VRRP group ID must be 13 and the IPv6 VRRP group ID must be 99 to meet the task requirement*

Verification EX1:

Note: This task can only be verified after the implementation of the VRRP task.

```
lab@ex1# run clear ethernet-switching table vlan-name vmx5vmx6ex1

{master:0}[edit]
lab@ex1# run show ethernet-switching table vlan-name vmx5vmx6ex1

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static,
C - Control MAC
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O -
ovsdb MAC)


Ethernet switching table : 6 entries, 6 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                 NH
RTR
    name                address             flags              interface
Index     ID
    vmx5vmx6ex1         00:00:5e:00:01:0d   D             -    ge-0/0/3.0              0
0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    vmx5vmx6ex1          00:00:5e:00:02:63   D              -   ge-0/0/4.0            0
0
    vmx5vmx6ex1          00:50:56:a9:39:44   D              -   ge-0/0/4.0            0
0
    vmx5vmx6ex1          00:50:56:a9:51:c6   D              -   ge-0/0/4.0            0
0
    vmx5vmx6ex1          00:56:56:56:56:05   D              -   ge-0/0/3.0            0
0
    vmx5vmx6ex1          00:56:56:56:56:06   D              -   ge-0/0/4.0            0
0
```

8) Ensure that only traffic received from the MAC address of server2 is allowed on port ge-1/0/3.

Note: Depending on when you run this command you may not have the MAC entry in the table. You can disable and enable the ge-1/0/4 interface to force DHCP to attempt to get an IP. This will cause traffic to be sent so the MAC address can be seen by the Ethernet switching process.

```
lab@ex1# run show ethernet-switching table interface ge-1/0/3.0

...
    name                 address             flags              interface
Index     ID
    server               cc:e1:94:d2:6e:47   D              -   ge-1/0/3.0            0
0

lab@ex1# show firewall
family ethernet-switching {
    filter allowmac {
        term 1 {
            from {
                source-mac-address {
                    cc:e1:94:d2:6e:47/48;
                }
            }
            then accept;
        }
        term 2 {
            then discard;
        }
    }
}

lab@ex1# show interfaces ge-1/0/3
description "to server2";
unit 0 {
    family ethernet-switching {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        interface-mode access;
        vlan {
            members server;
        }
        filter {
            input allowmac;
        }
    }
}
```

Layer 2 Features

9) Ensure Server1 and server 2 receive their IPv4 address from the DHCP server.

EX1:

```
lab@ex1# show forwarding-options
dhcp-relay {
    overrides;
    server-group {
        dhcp-server {
            10.10.10.1;
        }
    }
    group dhcp-server {
        active-server-group dhcp-server;
        interface irb.101;
    }
}

lab@ex1# show interfaces ge-0/0/6
description server1;
unit 0 {
    family inet {
        dhcp;
    }
    family inet6 {
        address 2001:0db8:0100:101::1/64;
    }
}

{master:0}[edit]
lab@ex1# show interfaces ge-1/0/4
description "server 2";
unit 0 {
    family inet {
        dhcp;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@ex1# show interfaces irb.101
family inet {
    address 10.100.101.254/24;
}
family inet6 {
    address 2001:0db8:0100:101::254/64;
}
```

Verification:

Note: You will not be able to verify this task until after you have enabled routing between Site 1 and the data center. The following outputs were taken after completing the BGP configurations in Task 5.

Note: You might need to deactivate and activate the ge-0/0/4 interface to trigger DHCP requests.

```
lab@ex1# run show dhcp client binding routing-instance server1
IP address        Hardware address   Expires     State      Interface
10.100.101.50     cc:e1:94:d2:6e:49  24930       BOUND      ge-0/0/6.0

lab@ex1# run show dhcp client binding routing-instance server2
IP address        Hardware address   Expires     State      Interface
10.100.101.51     b8:c2:53:97:97:a7  24931       BOUND      ge-1/0/4.0
```

10) User3 must be authenticated dynamically and ensure that only an authenticated device is allowed to communicate to the network. If authentication fails or the server is unreachable, the user must be given guest access.

EX1:

```
lab@ex1# show protocols dot1x
authenticator {
    authentication-profile-name jncie;
    interface {
        ge-1/0/5.0 {
            supplicant single-secure;
            guest-vlan guest;
            server-fail vlan-name guest;
        }
    }
}

lab@ex1# show access
radius-server {
    10.10.10.1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"; ## SECRET-DATA
        source-address 10.100.100.254;
    }
}
profile jncie {
    authentication-order radius;
    radius {
        authentication-server 10.10.10.1;
    }
}
```

Verification

Note: You must complete the routing tasks in order for the EX device to be able to reach the RADIUS server. The following verification commands were taken after completing the BGP configuration in Task 6.

Note: You may need to log into the User3 device (lab/lab123) and run the script “restart-supplicant.sh“ designed to restart the supplicant authentication attempts.

```
[lab@user1 ~]$ ./restart-supplicant.sh
```

EX1:

```
lab@ex1# run show dot1x interface
802.1X Information:
Interface     Role           State           MAC address          User
ge-1/0/5.0    Authenticator  Authenticated   00:50:56:00:00:03    jncie
```

11) Configure interface ge-0/0/7 and interface ge-1/0/9 on EX1. No more than 20 WATT should be allocated the POE ports.

EX1:

```
lab@ex1# show poe
interface ge-0/0/7 {
    maximum-power 20;
}
interface ge-1/0/9 {
    maximum-power 20;
}
```

12) Enable Sticky MAC for EX1 interface ge-1/0/9.

EX1:

```
lab@ex1# show switch-options
interface ge-1/0/9.0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    persistent-learning;
}
```

13) Configure EX1 to block all data packets on ports if Unknown Unicast packets consume more than 10% of the interface bandwidth. Broadcast and multicast traffic should not be affected by your solution. Control-plane traffic should never be blocked.

EX1:

```
lab@ex1# show switch-options
...
interface-shutdown-action soft-shutdown;

lab@ex1# show forwarding-options
storm-control-profiles default {
    all {
        bandwidth-percentage 10;
        no-broadcast;
        no-multicast;
    }
    action-shutdown;
}
```

Configurations for Task Completion:

EX1:

```
run request virtual-chassis vc-port set interface vcp-255/1/0 member 0
run request virtual-chassis vc-port set interface vcp-255/1/0 member 0
set virtual-chassis no-split-detection
set virtual-chassis member 0 mastership-priority 255
set virtual-chassis member 1 mastership-priority 255
set interfaces ge-1/0/3 description "to server2"
set interfaces ge-1/0/4 description "server 2"
set interfaces ge-1/0/5 description "to user3"
set interfaces ge-1/0/9 description "to phone2 user2"
set interfaces ge-1/0/10 description "phone2 user2"
set interfaces ge-1/0/10 unit 0 family inet address 10.100.100.2/24
set interfaces ge-1/0/10 unit 0 family inet6 address 2001:0db8:0100:100::2/64
set routing-instances server2 instance-type virtual-router
set routing-instances server2 interface ge-1/0/4.0
set routing-instances server2 routing-options rib server2.inet6.0 static route 0::/0
next-hop 2001:0db8:0100:101::254
set routing-instances server2 routing-options static route 0.0.0.0/0 next-hop
10.100.101.254
set routing-instances user2 instance-type virtual-router
set routing-instances user2 interface ge-1/0/10.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-instances user2 routing-options rib user2.inet6.0 static route 0::/0 next-
hop 2001:0db8:0100:100::254
set routing-instances user2 routing-options static route 0.0.0.0/0 next-hop
10.100.100.254
set vlans guest vlan-id 666
set vlans phone vlan-id 201
set vlans phone l3-interface irb.201
set vlans printer vlan-id 200
set vlans printer l3-interface irb.200
set vlans server vlan-id 101
set vlans server l3-interface irb.101
set vlans user vlan-id 100
set vlans user l3-interface irb.100
set vlans user switch-options interface ge-0/0/7.0
set vlans user switch-options interface ge-1/0/9.0
set vlans vmx5vmx6ex1 vlan-id 156
set vlans vmx5vmx6ex1 l3-interface irb.156
set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members vmx5vmx6ex1
set interfaces ge-0/0/4 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members vmx5vmx6ex1
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members server
set interfaces ge-0/0/7 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/7 unit 0 family ethernet-switching vlan members user
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members printer
set interfaces ge-1/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/3 unit 0 family ethernet-switching vlan members server
set interfaces ge-1/0/5 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/5 unit 0 family ethernet-switching vlan members user
set interfaces ge-1/0/9 unit 0 family ethernet-switching interface-mode access
set interfaces ge-1/0/9 unit 0 family ethernet-switching vlan members user
protect vlans
set switch-options voip interface ge-0/0/7.0 vlan phone
set switch-options voip interface ge-1/0/9.0 vlan phone
set firewall family ethernet-switching filter allowmac term 1 from source-mac-address
cc:e1:94:d2:6e:47/48
set firewall family ethernet-switching filter allowmac term 1 then accept
set firewall family ethernet-switching filter allowmac term 2 then discard
set interfaces ge-1/0/3 unit 0 family ethernet-switching filter input allowmac
set forwarding-options dhcp-relay overrides
set forwarding-options dhcp-relay server-group dhcp-server 10.10.10.1
set forwarding-options dhcp-relay group dhcp-server active-server-group dhcp-server
set forwarding-options dhcp-relay group dhcp-server interface irb.101
set interfaces ge-0/0/6 unit 0 family inet dhcp
set interfaces ge-1/0/4 unit 0 family inet dhcp
set access radius-server 10.10.10.1 secret "$9$/RW3Au1Srv7-wRh-wYgUD9Ap0RhylK8xN"
set access radius-server 10.10.10.1 source-address 10.100.100.254
set access profile jncie authentication-order radius
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set access profile jncie radius authentication-server 10.10.10.1
set protocols dot1x authenticator authentication-profile-name jncie
set protocols dot1x authenticator interface ge-1/0/5.0 supplicant single-secure
set protocols dot1x authenticator interface ge-1/0/5.0 guest-vlan guest
set protocols dot1x authenticator interface ge-1/0/5.0 server-fail vlan-name guest
set poe interface ge-0/0/7 maximum-power 20
set poe interface ge-1/0/9 maximum-power 20
set switch-options interface ge-1/0/9.0 persistent-learning
set switch-options interface-shutdown-action soft-shutdown
set forwarding-options storm-control-profiles default all bandwidth-percentage 10
set forwarding-options storm-control-profiles default all no-broadcast
set forwarding-options storm-control-profiles default all no-multicast
set forwarding-options storm-control-profiles default action-shutdown
```

EX2:

```
run request virtual-chassis vc-port set interface vcp-255/1/0
run request virtual-chassis vc-port set interface vcp-255/1/1
```

vMX5:

```
set interfaces ge-0/0/2 mac 00:56:56:56:56:05
```

vMX6:

```
set interfaces ge-0/0/3 mac 00:56:56:56:56:06
```

vMX7:

```
set chassis aggregated-devices ethernet device-count 16
set interfaces ge-0/0/1 gigether-options 802.3ad ae15
set interfaces ae15 aggregated-ether-options lacp active
set interfaces ae15 unit 0 family bridge interface-mode trunk
set interfaces ae15 unit 0 family bridge vlan-id-list [ 100 101 ]
```

vMX8:

```
set chassis aggregated-devices ethernet device-count 16
set interfaces ge-0/0/0 gigether-options 802.3ad ae15
set interfaces ae15 aggregated-ether-options lacp active
set interfaces ae15 unit 0 family bridge interface-mode trunk
set interfaces ae15 unit 0 family bridge vlan-id-list [ 100 101 ]
```

vQFX1:

```
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members server
set interfaces xe-0/0/5 unit 0 family ethernet-switching interface-mode access
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces xe-0/0/5 unit 0 family ethernet-switching vlan members user
set vlans server vlan-id 101
set vlans user vlan-id 100
```

vQFX2:

```
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members user
set vlans user vlan-id 100
set chassis aggregated-devices ethernet device-count 1
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members all
```

Task 3: Protocol Independent Features

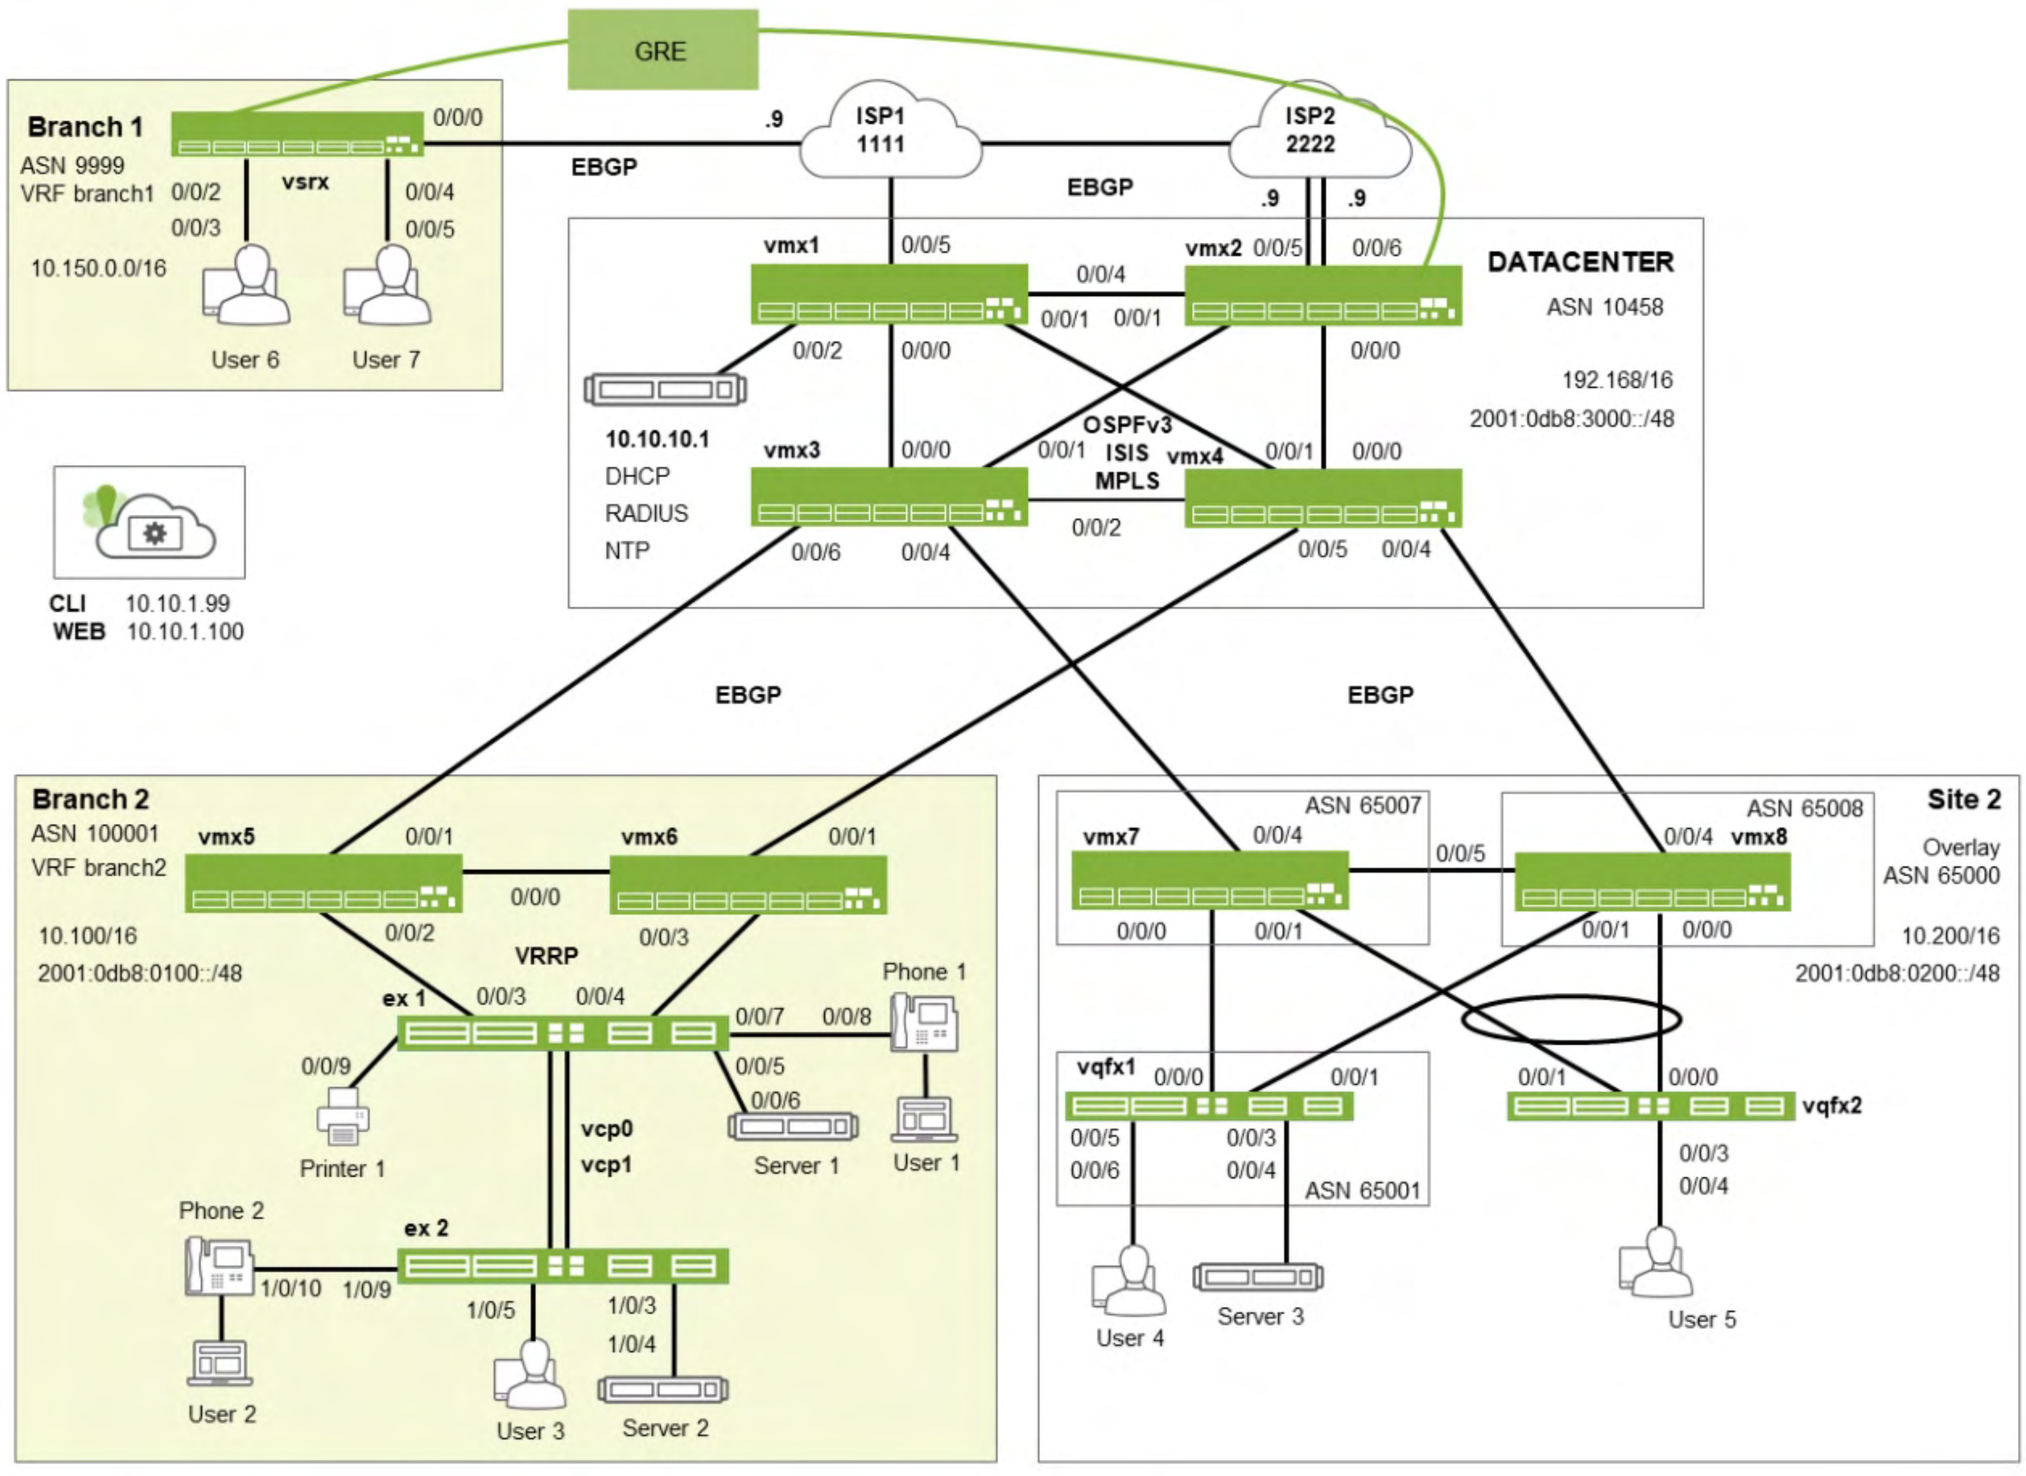

© 2020 Juniper Networks, Inc. All rights reserved.

GRE Tunnel

1) Configure a GRE tunnel between Branch 1 and vMX2 in the data center. Use the IP addressing as shown in the lab setup Information chapter. Use the loopback interface as GRE tunnel source address.

vMX2:

```
lab@vmx2# show chassis
fpc 0 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
    }
    lite-mode;
}


lab@vmx2# run show interfaces terse | match gr
gr-0/0/10               up    up


lab@vmx2# show interfaces gr-0/0/10
unit 0 {
    tunnel {
        source 192.168.0.2;
        destination 10.150.0.99;
    }
    family inet {
        address 173.17.99.2/24;
    }
}
```

vSRX:

```
lab@vsrx# show interfaces gr-0/0/0
unit 0 {
    tunnel {
        source 10.150.0.99;
        destination 192.168.0.2;
    }
    family inet {
        address 173.17.99.99/24;
    }
}

lab@vsrx# show routing-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
```

Filter-Based Forwarding

2) vMX2 should forward ping reply packets received from USER3 to the 28.184.130.64/27 using interface ge-0/0/5. If interface ge-0/0/5 goes down, then these packets should be dropped.

vMX2:

```
lab@vmx2# show routing-options
interface-routes {
    rib-group inet interfaces;
}
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 ];
    }
}

lab@vmx2# show firewall
filter fbf-v4 {
    term 1 {
        from {
            source-address {
                10.100.100.3/32;
            }
            destination-address {
                28.184.130.64/27;
            }
            protocol icmp;
            icmp-type echo-reply;
        }
        then {
            routing-instance fbf;
        }
    }
    term 2 {
        then accept;
    }
}

lab@vmx2# show routing-instances
fbf {
    instance-type forwarding;
    routing-options {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        static {
            route 0.0.0.0/0 next-hop 172.16.2.9;
        }
    }
}

lab@vmx2# show interfaces ge-0/0/0
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.24.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0024::2/64;
    }
}

lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.23.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
}

lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
}
```

Verification:

```
lab@vmx2# run show route table fbf.inet
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fbf.inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 21:00:07
                    > to 172.16.2.9 via ge-0/0/5.0
...
```

Commit Script

3) vMX6 has a commit script stored that needs to be enabled. The commit script provides a warning message if vMX6 has IPv6 control plane filtering disabled. The commit script is currently not working and needs to be fixed.

Locate the commit script:

vMX6:

```
lab@vmx6# run start shell
% cd /var/db/scripts/commit/
% ls
loopback
```

There is only one script (loopback) located in the well-known commit script folder.

Activate the commit script on vMX6:

```
lab@vmx6# show system scripts commit
file loopback;

lab@vmx6# commit
warning: [edit system scripts language python] not configured
error: Unable to run Python script 'loopback'
error: 1 error reported by commit scripts
error: commit script failure
```

It seems the “loopback” script is actually a Python script and we need to explicitly configure this in Junos

vMX6:

```
lab@vmx6# set system scripts language python

[edit]
lab@vmx6# commit
warning: Control plane protection disabled
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

It seems the script it executed as we received a warning message that control plane protection is disabled.

Verify if IPv6 control plane protection is actually disabled

vMX6:

```
lab@vmx6# show interfaces lo0
unit 0 {
    ...
    family inet6 {
        filter {
            input protect-re-ipv6;
        }
        address 2001:0db8:0100:0000::6/128;
    }
}
```

It seems there is an issue with the script as it gave us a warning, while IPv6 control plane protection is enabled.

Verify the actual script code:

vMX6:

```
lab@vmx6# run start shell
% cat /var/db/scripts/commit/loopback
#!/usr/bin/python

from junos import Junos_Configuration
import jcs


def main():
    root = Junos_Configuration

    lo0_int = root.find("./interfaces/interface[name='lo0']")
    lo0_int_fw = lo0_int.find("./unit[name='0']/family/inet/filter/input")

    if lo0_int and not(lo0_int_fw):
        jcs.emit_warning("Control plane protection disabled")


if __name__ == '__main__':
    main()
%
```

vMX6:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx6# show interfaces lo0
unit 0 {
    family inet {
        address 10.100.0.6/32 {
            primary;
            preferred;
        }
    }
```

It seems the script is searching for an IPv4 control plane filter as **family inet** is referenced and not family inet6. As there is no IPv4 control plane filter applied as show in the output above the script gives a warning that control plane protection is disabled.

Change the script so **family inet6** is referenced.

```
% cat /var/db/scripts/commit/loopback
#!/usr/bin/python

from junos import Junos_Configuration
import jcs


def main():
    root = Junos_Configuration

    lo0_int = root.find("./interfaces/interface[name='lo0']")
    lo0_int_fw = lo0_int.find("./unit[name='0']/family/inet6/filter/input")

    if lo0_int and not(lo0_int_fw):
        jcs.emit_warning("Control plane protection disabled")


if __name__ == '__main__':
    main()
%
```

Deactivate the IPv6 control plane filter on interface lo0 and commit.

```
[edit]
lab@vmx6# deactivate interfaces lo0.0 family inet6 filter

lab@vmx6# show interfaces lo0
unit 0 {
    family inet {
        address 10.100.0.6/32 {
            primary;
            preferred;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    family inet6 {
        inactive: filter {
            input protect-re-ipv6;
        }
        address 2001:0db8:0100:0000::6/128;
    }
}


[edit]
lab@vmx6# commit
warning: Control plane protection disabled
commit complete
```

Now the script is working as expected. A warning message is given if IPv6 control plane protection is disabled.

```
[edit]
lab@vmx6# rollback 1
load complete

[edit]
lab@vmx6# commit
commit complete
```

Static Routing and VRRP

© 2020 Juniper Networks, Inc. All rights reserved.

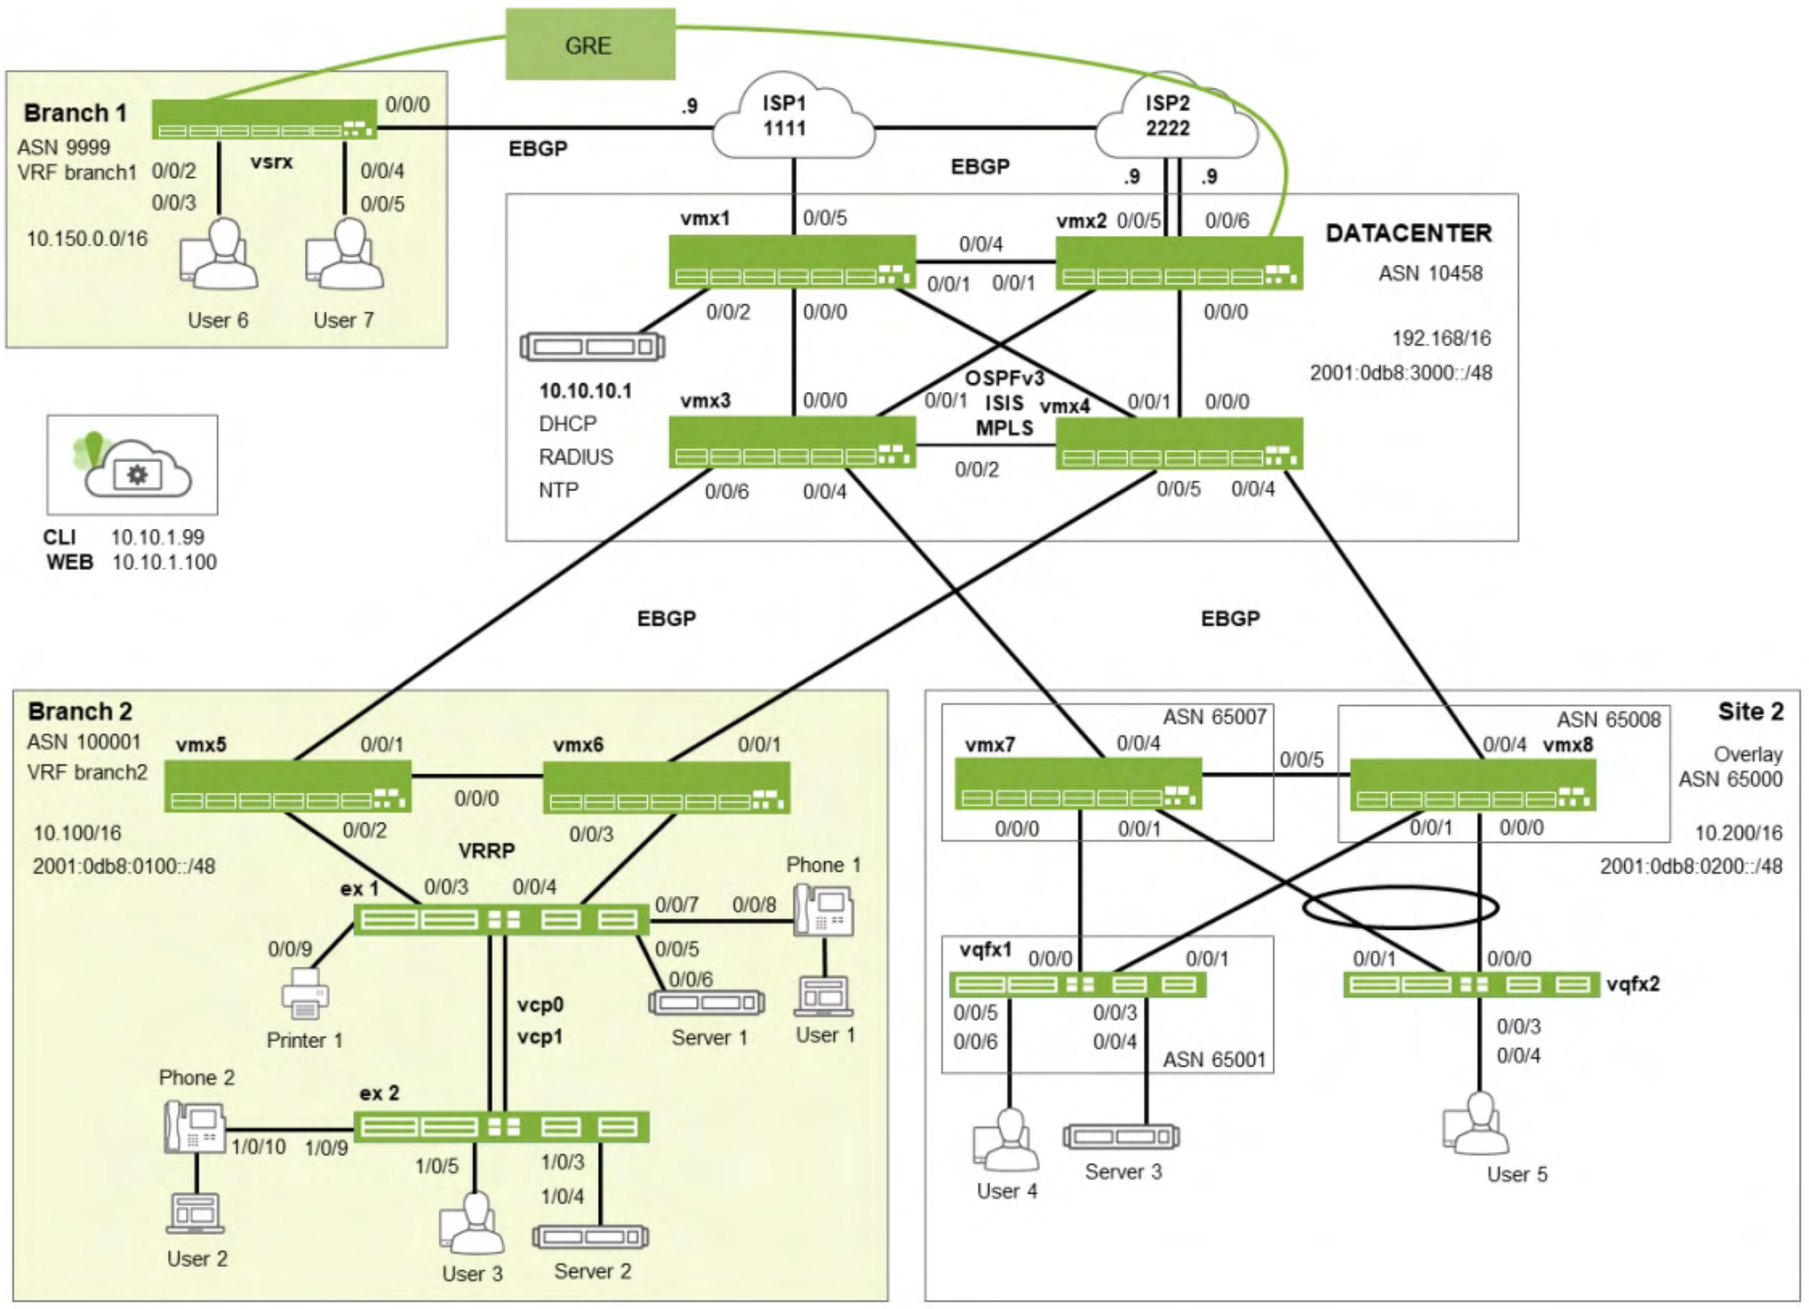

Branch1:

1) Configure static routing over the ge-0/0/0 interface to ensure vMX5 and vMX6 can reach each other loopback address

vMX5:

```
lab@vmx5# show routing-options
rib inet6.0 {
    static {
        route 2001:0db8:0100:0000::6/128 next-hop 2001:0db8:0100:0056::6;
    }
}
static {
    route 10.100.0.6/32 next-hop 10.100.56.6;
}
```

vMX6:

```
lab@vmx6# show routing-options
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
rib inet6.0 {
    static {
        route 2001:0db8:0100:0000::5/128 next-hop 2001:0db8:0100:0056::5;
    }
}
static {
    route 10.100.0.5/32 next-hop 10.100.56.5;
}
```

Verification:

```
lab@vmx5# run show route 2001:0db8:0100:0000::6

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::6/128*[Static/5] 2d 21:36:08
                    > to 2001:db8:100:56::6 via ge-0/0/0.0

[edit]
lab@vmx5# run show route 10.100.0.6

inet.0: 21 destinations, 30 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.6/32      *[Static/5] 2d 21:36:19
                    > to 10.100.56.6 via ge-0/0/0.0
```

2) Create two static summary routes on vMX5 and vMX6 that encompass the EX1 IRB interfaces.

vMX5 and vMX6:

```
lab@vmx5# show routing-options
rib inet6.0 {
    static {
        route 2001:0db8:0100::/48 next-hop 2001:0db8:0100:0156::1;
    }
}
static {
    route 10.100.0.0/16 next-hop 10.100.156.1;
}
```

Verification:

vMX5:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx5# run show route 2001:0db8:0100::/48 exact

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::/48  *[Static/5] 2d 22:28:29
                    > to 2001:db8:100:156::1 via ge-0/0/2.0

[edit]
lab@vmx5# run show route 10.100/16 exact

inet.0: 21 destinations, 30 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/16      *[Static/5] 2d 22:28:51
                    > to 10.100.156.1 via ge-0/0/2.0
```

vMX6:

```
lab@vmx6# run show route 2001:0db8:0100::/48 exact

inet6.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:100::/48  *[Static/5] 2d 22:29:03
                    > to 2001:db8:100:156::1 via ge-0/0/3.0

[edit]
lab@vmx6# run show route 10.100/16 exact

inet.0: 20 destinations, 29 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/16      *[Static/5] 2d 22:29:22
                    > to 10.100.156.1 via ge-0/0/3.0
```

3) Configure VRRP on vMX5 and vMX6 to provide a redundant default gateway for EX1. Make sure that the virtual address responds to icmp packets.

vMX5:

```
lab@vmx5# show interfaces ge-0/0/2
mac 00:56:56:56:56:05;
unit 0 {
    family inet {
        address 10.100.156.5/24 {
            vrrp-group 13 {
                virtual-address 10.100.156.254;
                accept-data;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::5/64 {
            vrrp-inet6-group 99 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                accept-data;
            }
        }
    }
}
```

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3
mac 00:56:56:56:56:06;
unit 0 {
    family inet {
        address 10.100.156.6/24 {
            vrrp-group 13 {
                virtual-address 10.100.156.254;
                accept-data;
            }
        }
    }
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 99 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
            }
        }
    }
}
```

4) vMX5 should be the VRRP master for IPv4 traffic. vMX6 should be the master for IPv6 traffic. After recovery of a failure a VRRP master should revert back to its role.

vMX5:

```
lab@vmx5# show interfaces ge-0/0/2
mac 00:56:56:56:56:05;
unit 0 {
    family inet {
        address 10.100.156.5/24 {
            vrrp-group 13 {
                virtual-address 10.100.156.254;
                priority 150;
                preempt;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                accept-data;
            }
        }
    }
}
```

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3
mac 00:56:56:56:56:06;
unit 0 {
    ...
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 99 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                priority 150;
                preempt;
                accept-data;
                track {
                    interface ge-0/0/1 {
                        priority-cost 51;
                    }
                }
            }
        }
    }
}
```

5) vMX5 should become VRRP master for IPv6 in case vMX6 interface ge-0/0/1 fails.

vMX6:

```
lab@vmx6# show interfaces ge-0/0/3
mac 00:56:56:56:56:06;
unit 0 {
    ...
    family inet6 {
        address 2001:0db8:0100:0156::6/64 {
            vrrp-inet6-group 99 {
                virtual-inet6-address 2001:0db8:0100:0156::254;
                priority 150;
                preempt;
                accept-data;
                track {
                    interface ge-0/0/1 {
                        priority-cost 51;
                    }
                }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            }
        }
    }
```

Verification:

vMX5:

```
lab@vmx5# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/2.0    up             13   master          Active     lcl     10.100.156.5
                                                              vip     10.100.156.254
ge-0/0/2.0    up             99   backup          Active     lcl     2001:db8:100:156::5
                                                              vip
fe80::200:5eff:fe00:263
                                                              vip     2001:db8:100:156::254
```

vMX6:

```
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up             13   backup          Active     lcl     10.100.156.6
                                                              vip     10.100.156.254
ge-0/0/3.0    up             99   master          Active     lcl     2001:db8:100:156::6
                                                              vip
fe80::200:5eff:fe00:263
                                                              vip     2001:db8:100:156::254

lab@vmx6# deactivate interfaces ge-0/0/1

[edit]
lab@vmx6# commit
commit complete

[edit]
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up             13   backup          Active     lcl     10.100.156.6
                                                              vip     10.100.156.254
ge-0/0/3.0    up             99   backup          Active     lcl     2001:db8:100:156::6
                                                              vip
fe80::200:5eff:fe00:263
                                                              vip     2001:db8:100:156::254

[edit]
lab@vmx6# rollback 1
load complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vmx6# commit
commit complete

[edit]
lab@vmx6# run show vrrp summary
Interface     State       Group   VR state       VR Mode     Type    Address
ge-0/0/3.0    up             13   backup          Active     lcl     10.100.156.6
                                                             vip     10.100.156.254
ge-0/0/3.0    up             99   master          Active     lcl     2001:db8:100:156::6
                                                             vip
fe80::200:5eff:fe00:263
                                                             vip     2001:db8:100:156::254
```

If interface ge-0/0/1 on vMX6 is disabled vMX5 becomes the VRRP master for group 99. vMX6 reclaims VRRP mastership when interface ge-0/0/1 is up again. This means VRRP preempt is working as expected.

Create two static default routes on EX1 towards the VRRP virtual address (.254) of vMX5 and vMX6.

EX1:

```
lab@ex1# show routing-options
rib inet6.0 {
    static {
        route ::/0 next-hop 2001:0db8:0100:0156::254;
    }
}
static {
    route 0.0.0.0/0 next-hop 10.100.156.254;
}
```

Verification:

```
lab@ex1# run show route protocol static table inet.0

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

0.0.0.0/0          *[Static/5] 22:00:49
                    >  to 10.100.156.254 via irb.156

lab@ex1# run show route protocol static table inet6.0

inet6.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
::/0                *[Static/5] 22:00:41
                     >  to 2001:db8:100:156::254 via irb.156
```

Configurations for Task Completion:

vMX2:

```
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
set interfaces gr-0/0/10 unit 0 tunnel source 192.168.0.2
set interfaces gr-0/0/10 unit 0 tunnel destination 10.150.0.99
set interfaces gr-0/0/10 unit 0 family inet address 173.17.99.2/24
set routing-options interface-routes rib-group inet interfaces
set routing-options rib-groups interfaces import-rib [ inet.0 fbf.inet.0 ]
set firewall filter fbf-v4 term 1 from source-address 10.100.100.3/32
set firewall filter fbf-v4 term 1 from destination-address 28.184.130.64/27
set firewall filter fbf-v4 term 1 from protocol icmp
set firewall filter fbf-v4 term 1 from icmp-type echo-reply
set firewall filter fbf-v4 term 1 then routing-instance fbf
set firewall filter fbf-v4 term 2 then accept
set routing-instances fbf instance-type forwarding
set routing-instances fbf routing-options static route 0.0.0.0/0 next-hop 172.16.2.9
set interfaces ge-0/0/0 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/1 unit 0 family inet filter input fbf-v4
set interfaces ge-0/0/4 unit 0 family inet filter input fbf-v4
```

vMX5:

```
set routing-options rib inet6.0 static route 2001:0db8:0100:0000::6/128 next-hop
2001:0db8:0100:0056::6
set routing-options static route 10.100.0.6/32 next-hop 10.100.56.6
set routing-options rib inet6.0 static route 2001:0db8:0100::/48 next-hop
2001:0db8:0100:0156::1
set routing-options static route 10.100.0.0/16 next-hop 10.100.156.1
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 13
virtual-address 10.100.156.254
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 13 accept-
data
set interfaces ge-0/0/2 unit 0 family inet6 address 2001:0db8:0100:0156::5/64 vrrp-
inet6-group 99 virtual-inet6-address 2001:0db8:0100:0156::254
set interfaces ge-0/0/2 unit 0 family inet6 address 2001:0db8:0100:0156::5/64 vrrp-
inet6-group 99 accept-data
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 13
priority 150
set interfaces ge-0/0/2 unit 0 family inet address 10.100.156.5/24 vrrp-group 13 preempt
```

vMX6:

```
set system scripts commit file loopback
set system scripts language python
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 static route 2001:0db8:0100:0000::5/128 next-hop
2001:0db8:0100:0056::5
set routing-options static route 10.100.0.5/32 next-hop 10.100.56.5
set routing-options rib inet6.0 static route 2001:0db8:0100::/48 next-hop
2001:0db8:0100:0156::1
set routing-options static route 10.100.0.0/16 next-hop 10.100.156.1
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 13
virtual-address 10.100.156.254
set interfaces ge-0/0/3 unit 0 family inet address 10.100.156.6/24 vrrp-group 13 accept-
data
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 99 virtual-inet6-address 2001:0db8:0100:0156::254
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 99 accept-data
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 99 priority 150
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 99 preempt
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:0db8:0100:0156::6/64 vrrp-
inet6-group 99 track interface ge-0/0/1 priority-cost 51
```

vSRX:

```
set interfaces gr-0/0/0 unit 0 tunnel source 10.150.0.99
set interfaces gr-0/0/0 unit 0 tunnel destination 192.168.0.2
set interfaces gr-0/0/0 unit 0 family inet address 173.17.99.99/24
set routing-options static route 192.168.0.2/32 next-hop 10.1.1.1
```

EX1:

```
set routing-options rib inet6.0 static route ::/0 next-hop 2001:0db8:0100:0156::254
set routing-options static route 0.0.0.0/0 next-hop 10.100.156.254
```

Task 4: IGP

This task focuses on the implementation of various IGPs like OSPFv2, OSPFv3 and static routing in combination with VRRP. In the data center, MPLS is enabled using the Label Distribution Protocol (LDP).

© 2020 Juniper Networks, Inc. All rights reserved.

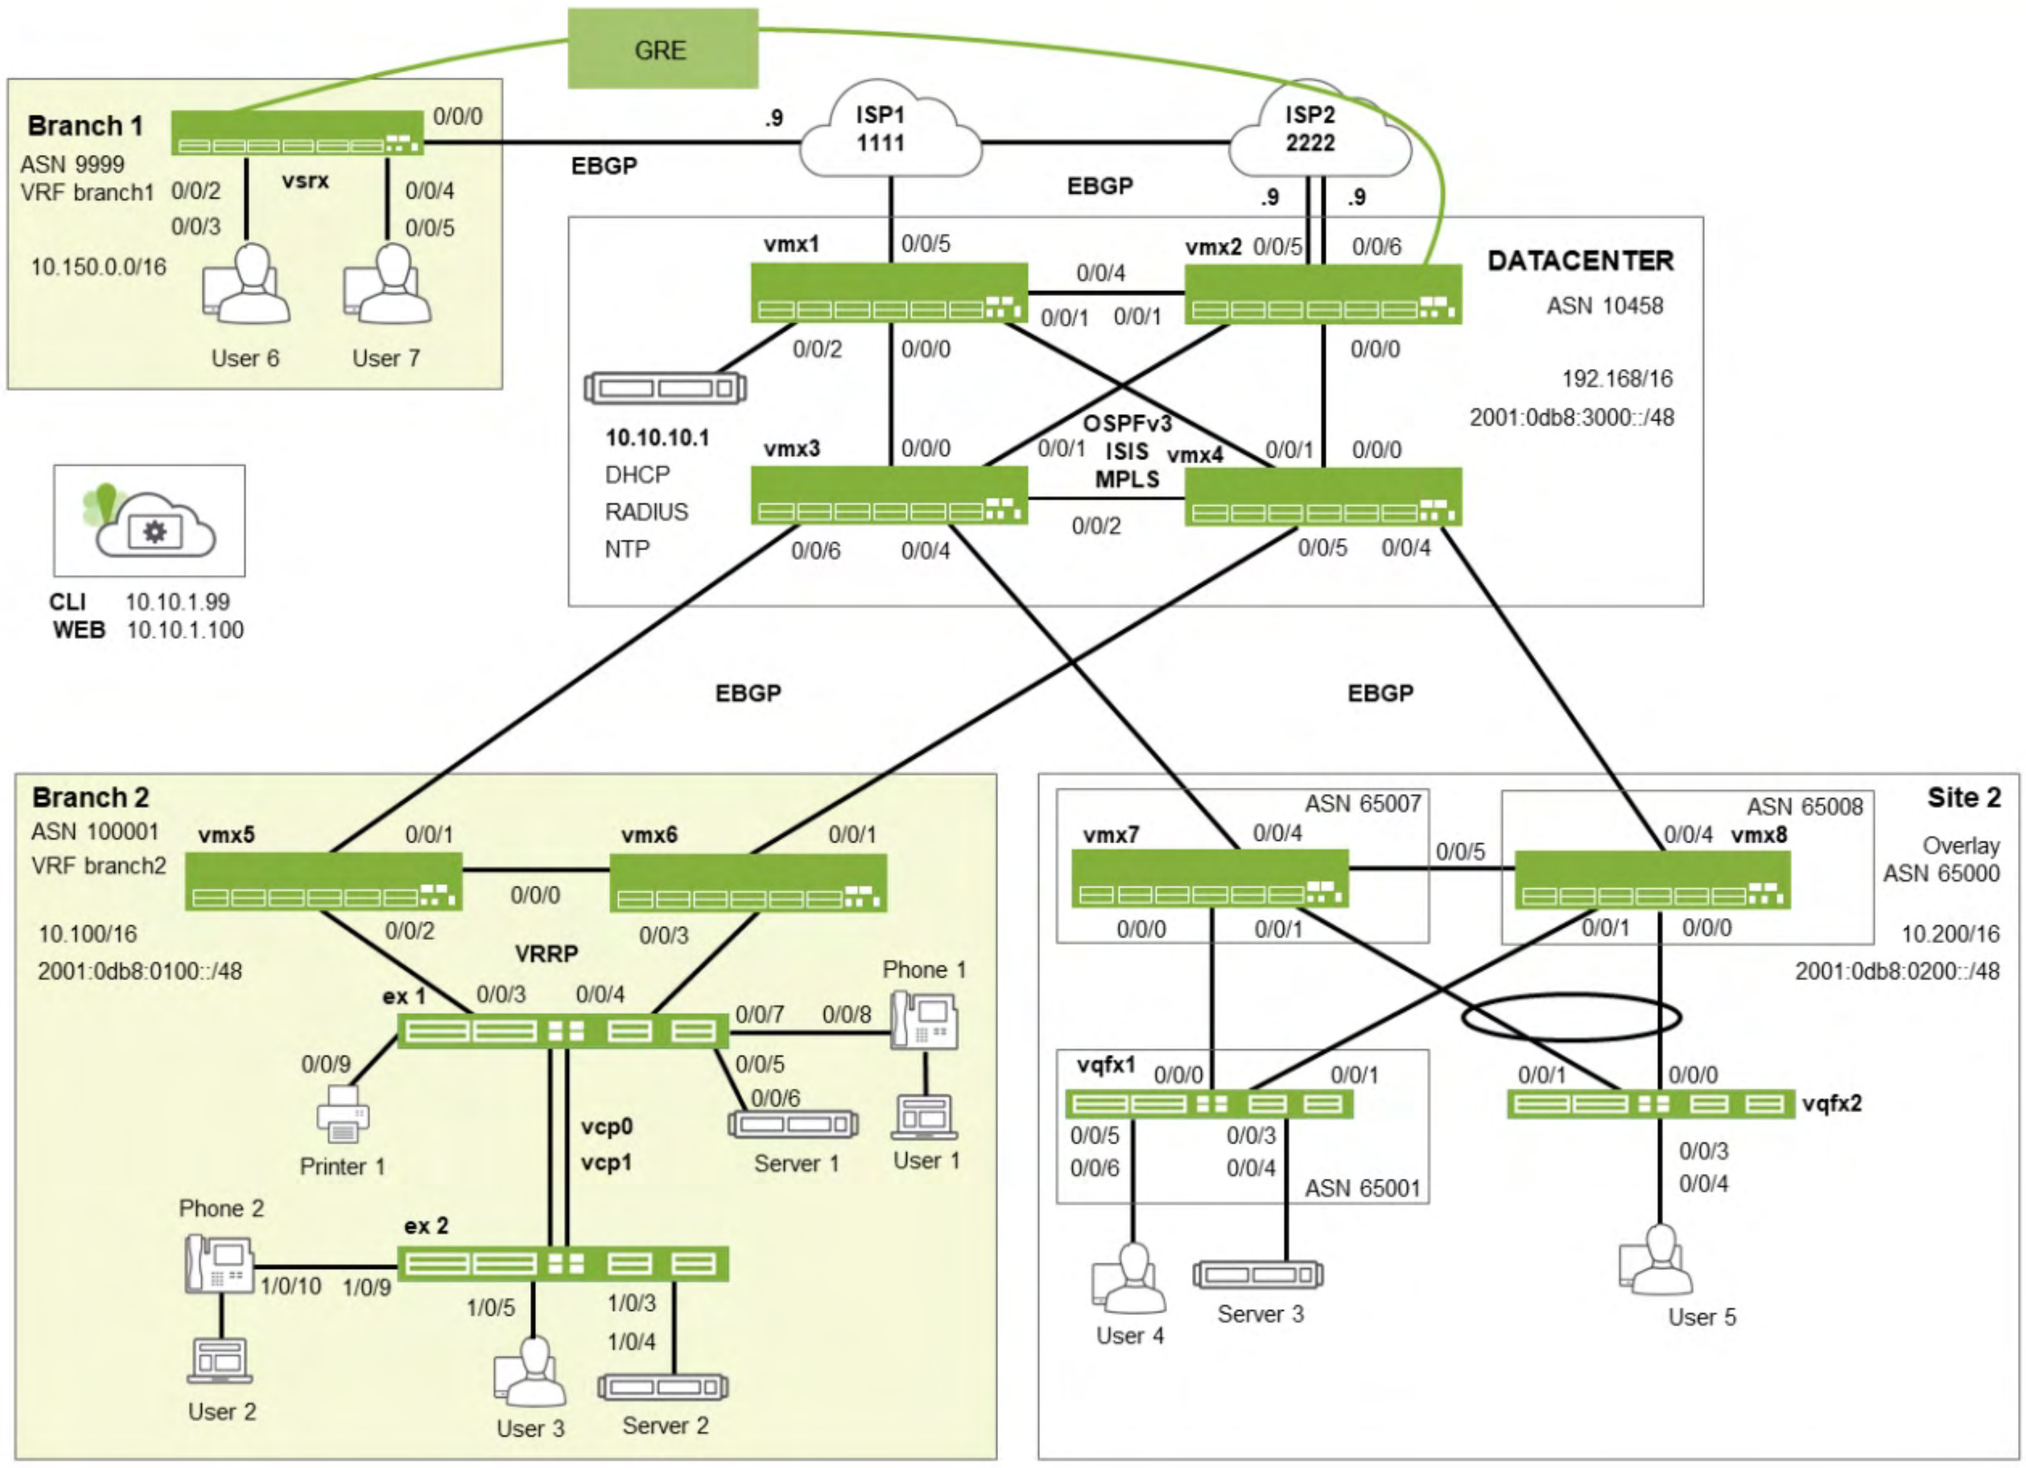

OSPFv3

Data Center:

Device	Interface	area
vMX1	lo0	0.0.0.1
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
vMX2	lo0	0.0.0.2
	ge-0/0/4	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
vMX3	lo0	0.0.0.3

© 2020 Juniper Networks, Inc. All rights reserved.

vMX4	ge-0/0/2	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0
	lo0	0.0.0.4
	ge-0/0/2	0.0.0.0
	ge-0/0/0	0.0.0.0
	ge-0/0/1	0.0.0.0

Table 45 OSPF information Data Center

1) Configure OSPFv3 for IPv6 routing in the data center using the information as shown in the table above.
2) Ensure no LSA type 0x2002 are found in the OSPFv3 database.
3) Configure the OSPF interface timers so that a neighbor is declared dead after four missed hello packets.
4) Ensure that each vMX devices in the data center learn each other's IPv6 loopback address.

vMX1:

```
lab@vmx1# show protocols ospf3
area 0.0.0.0 {
    interface ge-0/0/4.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/2.0 {
        passive;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
area 0.0.0.1 {
    interface lo0.0;
}
```

vMX2:

```
lab@vmx2# show protocols ospf3
area 0.0.0.0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        interface ge-0/0/4.0 {
            interface-type p2p;
            hello-interval 1;
        }
        interface ge-0/0/1.0 {
            interface-type p2p;
            hello-interval 1;
        }
        interface ge-0/0/0.0 {
            interface-type p2p;
            hello-interval 1;
        }
    }
    area 0.0.0.2 {
        interface lo0.0;
    }
```

vMX3:

```
lab@vmx3# show protocols ospf3
area 0.0.0.3 {
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.0 {
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
```

vMX4:

```
lab@vmx4# show protocols ospf3
area 0.0.0.4 {
    interface lo0.0 {
        passive;
    }
}
area 0.0.0.0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    interface ge-0/0/2.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/1.0 {
        interface-type p2p;
        hello-interval 1;
    }
    interface ge-0/0/0.0 {
        interface-type p2p;
        hello-interval 1;
    }
}
```

Verification:

vMX1:

```
lab@vmx1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
192.168.0.3     ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3aaa
192.168.0.4     ge-0/0/1.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:12b5
192.168.0.2     ge-0/0/4.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3b9f

[edit]
lab@vmx1# run show ospf3 database network
```

vMX2:

```
lab@vmx2# run show ospf3 neighbor
run show ospf3 database network
ID              Interface             State     Pri   Dead
192.168.0.4     ge-0/0/0.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:51c6
192.168.0.3     ge-0/0/1.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:32cf
192.168.0.1     ge-0/0/4.0            Full      128      3
  Neighbor-address fe80::250:56ff:fea9:4466

[edit]
lab@vmx2# run show ospf3 database network
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX3:

```
lab@vmx3# run show ospf3 neighbor
run show ospf3 database network    ID                 Interface              State     Pri
Dead
192.168.0.1       ge-0/0/0.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3ed3
192.168.0.2       ge-0/0/1.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:35aa
192.168.0.4       ge-0/0/2.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:36

[edit]
lab@vmx3# run show ospf3 database network
```

vMX4:

```
lab@vmx4# run show ospf3 neighbor
ID                 Interface              State     Pri   Dead
192.168.0.2       ge-0/0/0.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:3944
192.168.0.1       ge-0/0/1.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:427f
192.168.0.3       ge-0/0/2.0              Full      128      3
  Neighbor-address fe80::250:56ff:fea9:6e2a

[edit]
lab@vmx4# run show ospf3 database network
```


© 2020 Juniper Networks, Inc. All rights reserved.

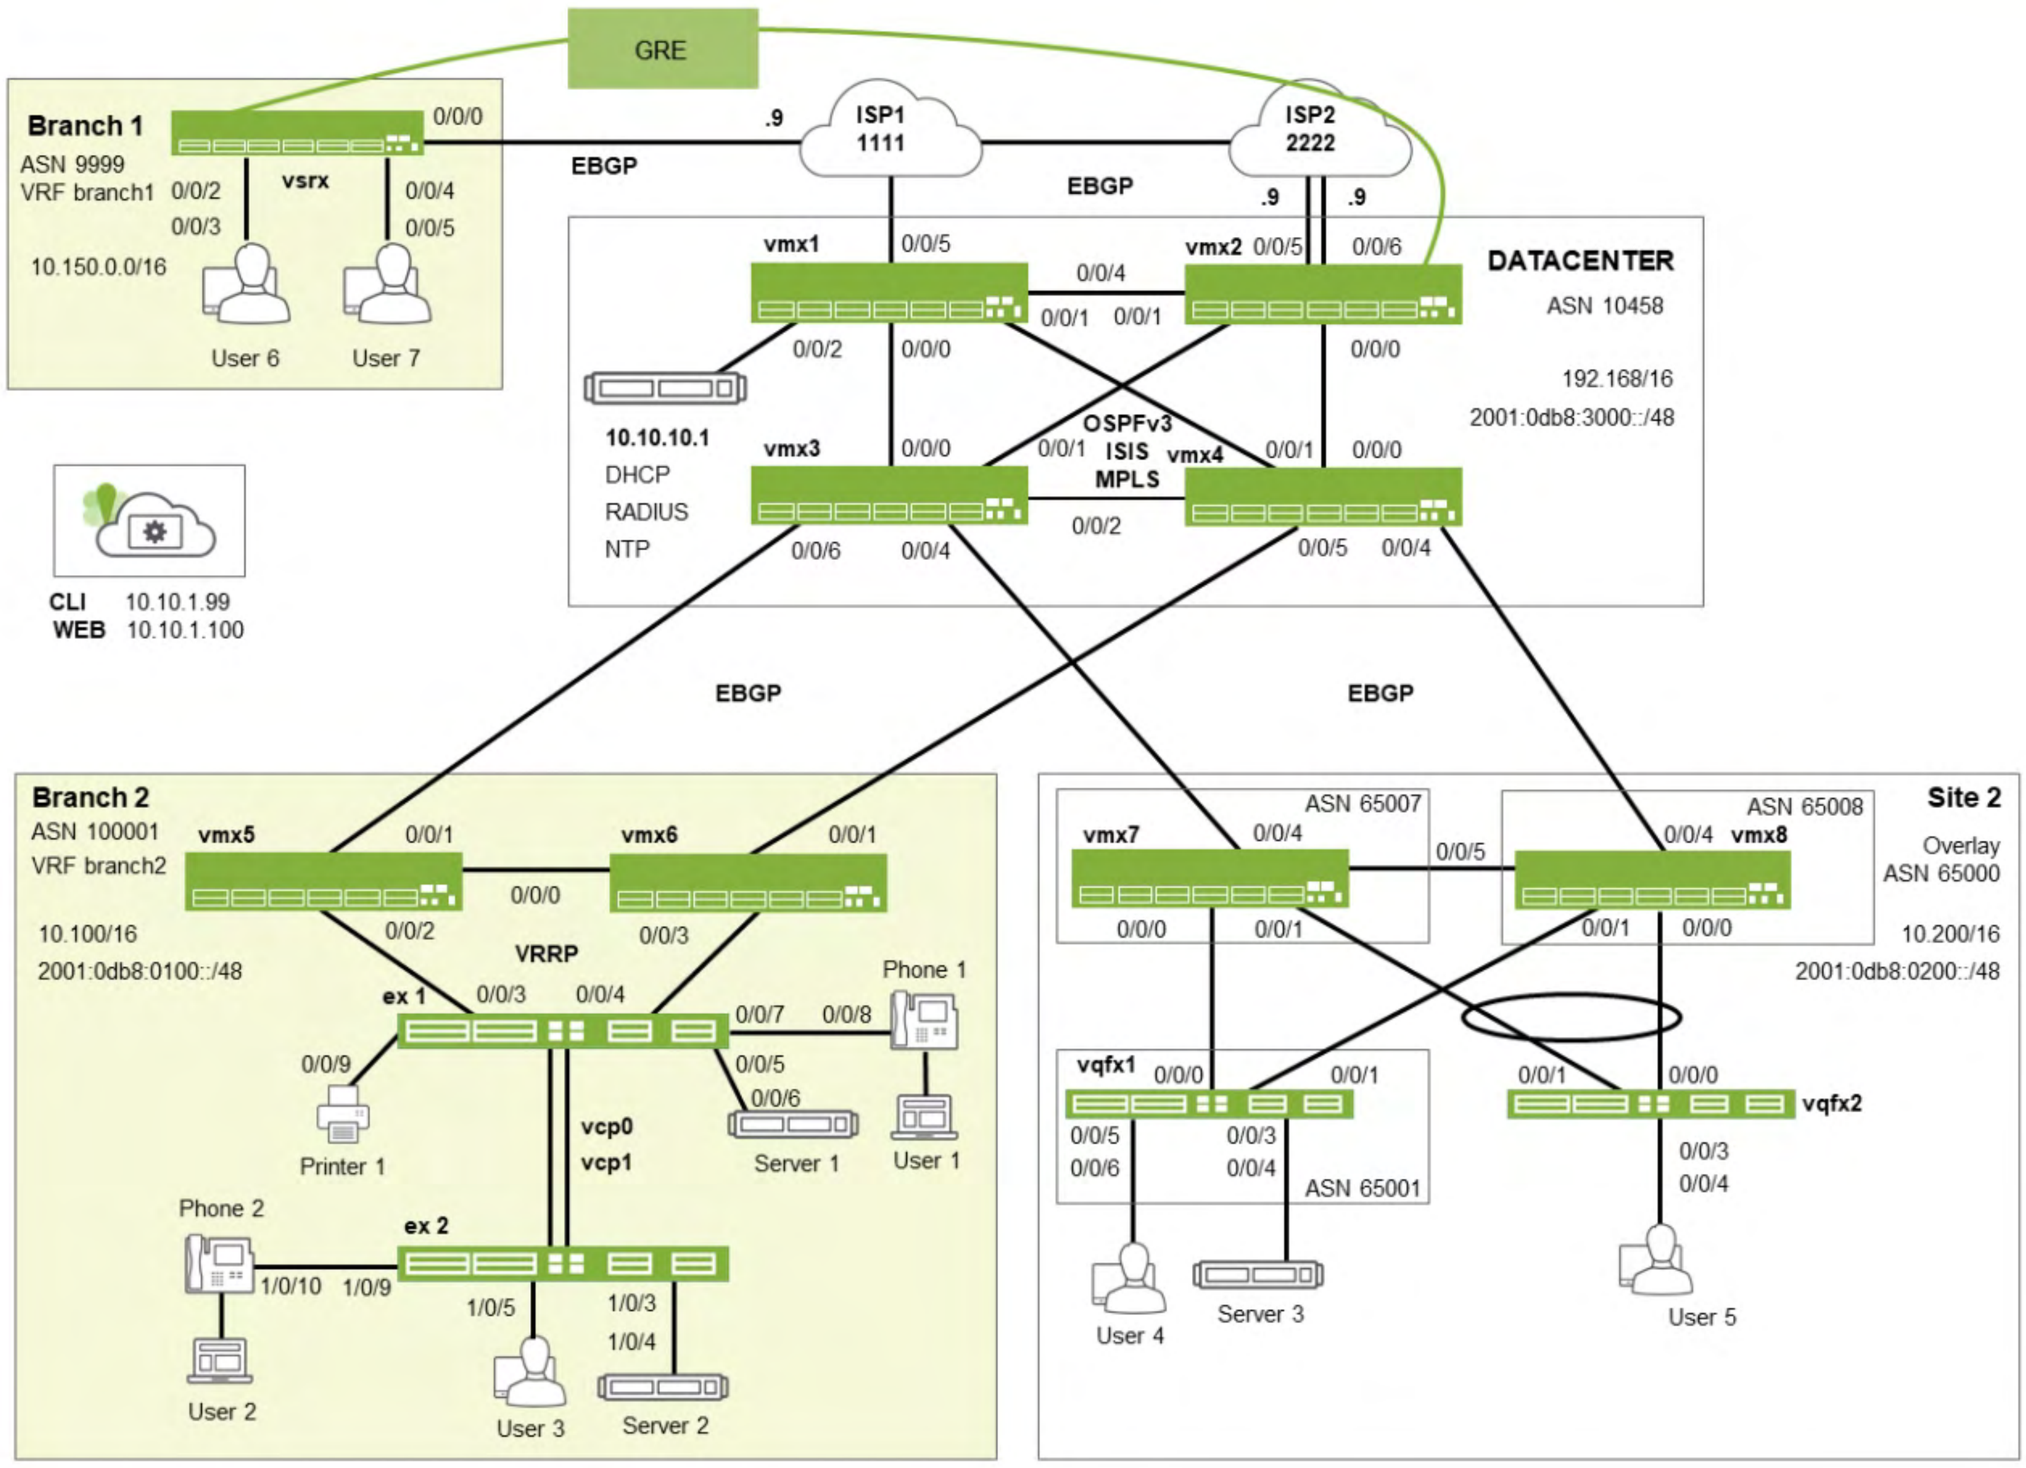

OSPFv2

In this section you will configure an OSPFv2 for IPv4 routing in the data center and over the GRE tunnel between the data center and the branch office.

Device	Interface	Routing Instance	Area
vMX2	GRE interface	branch1	0.0.0.0
vSRX	GRE interface	default	0.0.0.0

Table 46 OSPF information for GRE tunnel

5) Enable OSPFv2 over the GRE interface as shown in the table above.

© 2020 Juniper Networks, Inc. All rights reserved.

6) Ensure the data center OSPFv2 database contains an external type 1 LSA for the 10.150.99/24 network.
7) Ensure that vSRX OSPF database contains a default route and the IS-IS destinations received from vMX2. Configure OSPF on vSRX that the default route is the only external LSA installed in its routing table.

vMX2:

```
lab@vmx2# show routing-options
interface-routes {
    rib-group inet interfaces;
}
...
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 branch1.inet.0 ];
    }
    branch1-to-inet {
        import-rib [ branch1.inet.0 inet.0 ];
    }
    inet-to-branch1 {
        import-rib [ inet.0 branch1.inet.0 ];
    }
}

lab@vmx2# show routing-instances branch1
instance-type vrf;
interface gr-0/0/10.0;
route-distinguisher 192.168.0.2:1;
vrf-target target:10458:1;
routing-options {
    interface-routes {
        rib-group inet branch1-to-inet;
    }
    static {
        route 0.0.0.0/0 next-table inet.0;
    }
}
protocols {
    ospf {
        rib-group branch1-to-inet;
        export 2ospfbranch1;
        area 0.0.0.0 {
            interface gr-0/0/10.0 {
                hello-interval 1;
            }
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vmx2# show policy-options policy-statement 2ospfbranch1
term 1 {
    from protocol [ direct static isis ];
    then accept;
}
term 2 {
    then reject;
}
```

vSRX1:

```
lab@vsrx# show protocols ospf
export 2ospf;
import ospfimport;
area 0.0.0.0 {
    interface gr-0/0/0.0 {
        hello-interval 1;
    }
}

lab@vsrx# show policy-options policy-statement 2ospf
term 1 {
    from interface irb.99;
    then {
        external {
            type 1;
        }
        accept;
    }
}
term 2 {
    then reject;
}
```

```
lab@vsrx# show policy-options policy-statement ospfimport
term 1 {
    from {
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

Note: Verification must be done after completing the IS-IS and BGP tasks.

vMX2:

```
lab@vmx2# run show route table branch1.inet.0

branch1.inet.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


0.0.0.0/0          *[Static/5] 1d 17:26:44
                      to table inet.0
10.10.10.0/24      *[IS-IS/18] 1d 16:57:22, metric 20
                    > to 192.168.12.1 via ge-0/0/4.0

10.150.99.0/24     *[OSPF/150] 1d 17:25:57, metric 1, tag 0
                    > via gr-0/0/10.0
...
```

vSRX:

```
lab@vsrx# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *10.150.0.99      10.150.0.99     0x800001f1  2287  0x22 0xaeac  48
Router   173.17.99.2      173.17.99.2     0x80000155  2571  0x22 0x1127  48
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   0.0.0.0          173.17.99.2     0x800000ed   371  0x22 0xbfe1  36
Extern   10.10.10.0       173.17.99.2     0x80000033   171  0x22 0x9495  36
Extern  *10.150.99.0      10.150.0.99     0x80000175   184  0x22 0x575e  36
Extern   172.16.2.0       173.17.99.2     0x8000001d  1971  0x22 0xc5ed  36
Extern   172.16.3.0       173.17.99.2     0x8000001d  1771  0x22 0xbaf7  36
Extern   192.168.0.1      173.17.99.2     0x80000032  1571  0x22 0xdf09  36
Extern   192.168.0.2      173.17.99.2     0x8000001d  2971  0x22 0x9b6b  36
Extern   192.168.0.3      173.17.99.2     0x80000032  1371  0x22 0xcb1b  36
Extern   192.168.0.4      173.17.99.2     0x80000032  1171  0x22 0xc124  36
Extern   192.168.12.0     173.17.99.2     0x8000001d  2171  0x22 0x2bd1  36
Extern   192.168.13.0     173.17.99.2     0x80000032   971  0x22 0xbe14  36
Extern   192.168.14.0     173.17.99.2     0x80000032   771  0x22 0xb31e  36
Extern   192.168.23.0     173.17.99.2     0x8000001d  2371  0x22 0xb140  36
Extern   192.168.24.0     173.17.99.2     0x8000001d  2771  0x22 0xa64a  36
Extern   192.168.34.0     173.17.99.2     0x80000032   571  0x22 0xd6e6  36
```

The output above shows that vSRX receives the IS-IS routes from the data center and a default route.

© 2020 Juniper Networks, Inc. All rights reserved.

vSRX:

```
lab@vsrx# run show route protocol ospf

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[OSPF/150] 1d 17:30:17, metric 0, tag 0
                     >  via gr-0/0/0.0
173.17.99.0/24       [OSPF/10] 1w5d 21:28:22, metric 1
                     >  via gr-0/0/0.0
224.0.0.5/32        *[OSPF/10] 1w5d 21:29:00, metric 1
                        MultiRecv
```

The OSPF import policy ensures that the default route is the only external LSA that is installed in the routing table of vSRX.

8) Configure IS-IS in the data center using the information as shown in the tables below.

IS-IS Configuration

Device	Level	Interfaces
vMX1	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
		ge-0/0/2
vMX2	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/4
vMX3	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2
vMX4	2	lo0
		ge-0/0/0
		ge-0/0/1
		ge-0/0/2

Table 47 IS-IS information

© 2020 Juniper Networks, Inc. All rights reserved.

Device	Required IS-IS configuration
vMX1	Sysid: 0168.0000.0001 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled
vMX2	Sysid: 0168.0000.0002 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled
vMX3	Sysid: 0168.0000.0003 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2 IPv4 is enabled, IPv6 is disabled
vMX4	Sysid: 0168.0000.0004 Areaid: 49.0001.0192 Adjacency holddown: enabled Maximum Areas: 3 LSP life time: 1200 Attached bit evaluation: enabled SPF delay: 50 msec, SPF holddown: 2500 msec,

© 2020 Juniper Networks, Inc. All rights reserved.

SPF rapid runs: 2
IPv4 is enabled, IPv6 is disabled

9) Do not enable Level 1 routing.
10) Ensure no DIS is elected.
11) Ensure interface ge-0/0/2 on vMX1 participates in IS-IS, but no adjacency can be formed.
12) The maximum number of routes that can be redistributed in IS-IS should be 100.
13) Enable MD5 authentication for level 2. Use “juniper123” as the key.

vMX1:

```
lab@vmx1# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.13.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::1/64;
    }
}

lab@vmx1# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::1/64;
    }
}

lab@vmx1# show interfaces ge-0/0/4
unit 0 {
    family inet {
        address 192.168.12.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::1/64;
    }
}

lab@vmx1# show interfaces lo0
unit 0 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    family inet {
        address 192.168.0.1/32 {
            primary;
            preferred;
        }
    }
    family iso {
        address 49.0001.0192.0168.0000.0001.00;
    }
    family inet6 {
        address 2001:0db8:3000:0000::1/128;
    }
}
```

```
lab@vmx1# show protocols isis
no-ipv6-routing;
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"; ## SECRET-DATA
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    passive;
}
interface ge-0/0/4.0 {
    point-to-point;
}
interface lo0.0;
```

vMX2:

```
lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        address 192.168.23.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
}

lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
}

lab@vmx2# show interfaces lo0
unit 0 {
    family inet {
        address 192.168.0.2/32 {
            primary;
            preferred;
        }
    }
    family iso {
        address 49.0001.0192.0168.0000.0002.00;
    }
    family inet6 {
        address 2001:0db8:3000:0000::2/128;
    }
}

lab@vmx2# show protocols isis
export 2isis;
no-ipv6-routing;
rib-group inet inetzero-to-branch1;
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/4.0 {
    point-to-point;
}
interface lo0.0;

lab@vmx2# show policy-options policy-statement 2isis
term 1 {
    from protocol ospf;
    then accept;
}
```

vMX3:

```
lab@vmx3# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.13.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::3/64;
    }
}

lab@vmx3# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.23.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0023::3/64;
    }
}

lab@vmx3# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.3/24;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::3/64;
    }
}

lab@vmx3# show interfaces lo0
unit 0 {
    family inet {
        address 192.168.0.3/32 {
            primary;
            preferred;
        }
    }
    family iso {
        address 49.0001.0192.0168.0000.0003.00;
    }
    family inet6 {
        address 2001:0db8:3000:0000::3/128;
    }
}

lab@vmx3# show protocols isis
no-ipv6-routing;
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"; ## SECRET-DATA
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    point-to-point;
}
interface lo0.0;
```

vMX4:

```
lab@vmx4# show interfaces ge-0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
unit 0 {
    family inet {
        address 192.168.24.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0024::4/64;
    }
}

lab@vmx4# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::4/64;
    }
}

lab@vmx4# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::4/64;
    }
}

lab@vmx4# show interfaces lo0
unit 0 {
    family inet {
        address 192.168.0.4/32 {
            primary;
            preferred;
        }
    }
    family iso {
        address 49.0001.0192.0168.0000.0004.00;
    }
    family inet6 {
        address 2001:0db8:3000:0000::4/128;
    }
}

lab@vmx4# show protocols isis
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
no-ipv6-routing;
spf-options {
    delay 50;
    holddown 2500;
    rapid-runs 2;
}
level 2 {
    authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"; ## SECRET-DATA
    authentication-type md5;
    prefix-export-limit 100;
}
level 1 disable;
interface ge-0/0/0.0 {
    point-to-point;
}
interface ge-0/0/1.0 {
    point-to-point;
}
interface ge-0/0/2.0 {
    point-to-point;
}
interface lo0.0;
```

Verification:

vMX1:

```
lab@vmx1# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0001
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled

lab@vmx1# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx3           2  Up                   20
ge-0/0/1.0            vmx4           2  Up                   23
ge-0/0/4.0            vmx2           2  Up                   26

lab@vmx1# run show route protocol isis

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

192.168.0.2/32     *[IS-IS/18] 00:04:28, metric 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.0.3/32     *[IS-IS/18] 00:02:35, metric 10
                    > to 192.168.13.3 via ge-0/0/0.0
192.168.0.4/32     *[IS-IS/18] 00:00:25, metric 20
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.23.0/24    *[IS-IS/18] 00:02:35, metric 20
                      to 192.168.13.3 via ge-0/0/0.0
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.24.0/24    *[IS-IS/18] 00:04:28, metric 20
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.34.0/24    *[IS-IS/18] 00:02:35, metric 20
                    > to 192.168.13.3 via ge-0/0/0.0
...
```

vMX2:

```
lab@vmx2# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0002
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled

lab@vmx2# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx4           2  Up                   26
ge-0/0/1.0            vmx3           2  Up                   21
ge-0/0/4.0            vmx1           2  Up                   24

lab@vmx2# run show route protocol isis

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

10.10.10.0/24      *[IS-IS/18] 00:05:15, metric 20
                    > to 192.168.12.1 via ge-0/0/4.0
192.168.0.1/32     *[IS-IS/18] 00:05:15, metric 10
                    > to 192.168.12.1 via ge-0/0/4.0
192.168.0.3/32     *[IS-IS/18] 00:02:54, metric 10
                    > to 192.168.23.3 via ge-0/0/1.0
192.168.0.4/32     *[IS-IS/18] 00:01:11, metric 10
                    > to 192.168.24.4 via ge-0/0/0.0
192.168.13.0/24    *[IS-IS/18] 00:02:54, metric 20
                    > to 192.168.23.3 via ge-0/0/1.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.14.0/24    *[IS-IS/18] 00:01:11, metric 20
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 192.168.24.4 via ge-0/0/0.0
                      to 192.168.12.1 via ge-0/0/4.0
192.168.34.0/24     *[IS-IS/18] 00:01:11, metric 20
                      to 192.168.24.4 via ge-0/0/0.0
                    > to 192.168.23.3 via ge-0/0/1.0
...
```

vMX3:

```
lab@vmx3# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0003
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled

  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled

lab@vmx3# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx1           2  Up                    23
ge-0/0/1.0            vmx2           2  Up                    19
ge-0/0/2.0            vmx4           2  Up                    26

lab@vmx3# run show route protocol isis

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

10.10.10.0/24       *[IS-IS/18] 00:03:58, metric 20
                     > to 192.168.13.1 via ge-0/0/0.0
192.168.0.1/32      *[IS-IS/18] 00:03:58, metric 10
                     > to 192.168.13.1 via ge-0/0/0.0
192.168.0.2/32      *[IS-IS/18] 00:03:30, metric 10
                     > to 192.168.23.2 via ge-0/0/1.0
192.168.0.4/32      *[IS-IS/18] 00:01:22, metric 10
                     > to 192.168.34.4 via ge-0/0/2.0
192.168.12.0/24     *[IS-IS/18] 00:03:30, metric 20
                     > to 192.168.13.1 via ge-0/0/0.0
                       to 192.168.23.2 via ge-0/0/1.0
192.168.14.0/24     *[IS-IS/18] 00:01:22, metric 20
                     > to 192.168.13.1 via ge-0/0/0.0
                       to 192.168.34.4 via ge-0/0/2.0
192.168.24.0/24     *[IS-IS/18] 00:01:22, metric 20
                     > to 192.168.23.2 via ge-0/0/1.0
                       to 192.168.34.4 via ge-0/0/2.0
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX4:

```
lab@vmx4# run show isis overview | match "Sysid|Areaid|Adjac|Maxi|LSP|Attach|SPF|IPv4"
  Sysid: 0168.0000.0004
  Areaid: 49.0001.0192
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 50 msec, SPF holddown: 2500 msec, SPF rapid runs: 2
  IPv4 is enabled

lab@vmx4# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/0.0            vmx2           2  Up                   24
ge-0/0/1.0            vmx1           2  Up                   19
ge-0/0/2.0            vmx3           2  Up                   25

lab@vmx4# run show route protocol isis

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

10.10.10.0/24      *[IS-IS/18] 00:02:00, metric 20
                    > to 192.168.14.1 via ge-0/0/1.0
192.168.0.1/32     *[IS-IS/18] 00:02:00, metric 10
                    > to 192.168.14.1 via ge-0/0/1.0
192.168.0.2/32     *[IS-IS/18] 00:02:27, metric 10
                    > to 192.168.24.2 via ge-0/0/0.0
192.168.0.3/32     *[IS-IS/18] 00:02:01, metric 10
                    > to 192.168.34.3 via ge-0/0/2.0
192.168.12.0/24    *[IS-IS/18] 00:02:00, metric 20
                    > to 192.168.24.2 via ge-0/0/0.0
                      to 192.168.14.1 via ge-0/0/1.0
192.168.13.0/24    *[IS-IS/18] 00:02:00, metric 20
                    > to 192.168.14.1 via ge-0/0/1.0
                      to 192.168.34.3 via ge-0/0/2.0
192.168.23.0/24    *[IS-IS/18] 00:02:01, metric 20
                      to 192.168.24.2 via ge-0/0/0.0
                    > to 192.168.34.3 via ge-0/0/2.0
...
```

Configurations for Task Completion:

vMX1:

```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.1 interface lo0.0
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/4 unit 0 family iso
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0001.00
set protocols isis no-ipv6-routing
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"
set protocols isis level 2 authentication-type md5
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 passive
set protocols isis interface ge-0/0/4.0 point-to-point
set protocols isis interface lo0.0
```

vMX2:

```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set protocols ospf3 area 0.0.0.2 interface lo0.0
set routing-options rib-groups interfaces import-rib branch1.inet.0
set routing-options rib-groups branch1-to-inet import-rib [ branch1.inet.0 inet.0 ]
set routing-options rib-groups inet-to-branch1 import-rib [ inet.0 branch1.inet.0 ]
set routing-instances branch1 instance-type vrf
set routing-instances branch1 interface gr-0/0/10.0
set routing-instances branch1 route-distinguisher 192.168.0.2:1
set routing-instances branch1 vrf-target target:10458:1
set routing-instances branch1 routing-options interface-routes rib-group inet branch1-
to-inet
set routing-instances branch1 routing-options static route 0.0.0.0/0 next-table inet.0
set routing-instances branch1 protocols ospf rib-group branch1-to-inet
set routing-instances branch1 protocols ospf export 2ospfbranch1
set routing-instances branch1 protocols ospf area 0.0.0.0 interface gr-0/0/10.0 hello-
interval 1
set policy-options policy-statement 2ospfbranch1 term 1 from protocol [ direct static
isis ]
set policy-options policy-statement 2ospfbranch1 term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ospfbranch1 term 2 then reject
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/4 unit 0 family iso
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0002.00
set protocols isis export 2isis
set protocols isis no-ipv6-routing
set protocols isis rib-group inet inet-to-branch1
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"
set protocols isis level 2 authentication-type md5
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/4.0 point-to-point
set protocols isis interface lo0.0
set policy-options policy-statement 2isis term 1 from protocol ospf
set policy-options policy-statement 2isis term 1 then accept
```

vMX3:

```
set protocols ospf3 area 0.0.0.3 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family iso
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0003.00
set protocols isis no-ipv6-routing
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"
set protocols isis level 2 authentication-type md5
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 point-to-point
set protocols isis interface lo0.0
```

vMX4:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols ospf3 area 0.0.0.4 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/2.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 hello-interval 1
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 hello-interval 1
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family iso
set interfaces lo0 unit 0 family iso address 49.0001.0192.0168.0000.0004.00
set protocols isis no-ipv6-routing
set protocols isis spf-options delay 50
set protocols isis spf-options holddown 2500
set protocols isis spf-options rapid-runs 2
set protocols isis level 2 authentication-key "$9$7EdwgGDkTz6oJz69A1INdbsoJUjHm5Q"
set protocols isis level 2 authentication-type md5
set protocols isis level 2 prefix-export-limit 100
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0 point-to-point
set protocols isis interface ge-0/0/1.0 point-to-point
set protocols isis interface ge-0/0/2.0 point-to-point
set protocols isis interface lo0.0
```

vSRX:

```
set protocols ospf export 2ospf
set protocols ospf import ospfimport
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 hello-interval 1
set policy-options policy-statement 2ospf term 1 from interface irb.99
set policy-options policy-statement 2ospf term 1 then external type 1
set policy-options policy-statement 2ospf term 1 then accept
set policy-options policy-statement 2ospf term 2 then reject
set policy-options policy-statement ospfimport term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement ospfimport term 1 then accept
set policy-options policy-statement ospfimport term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 5: MPLS/LDP

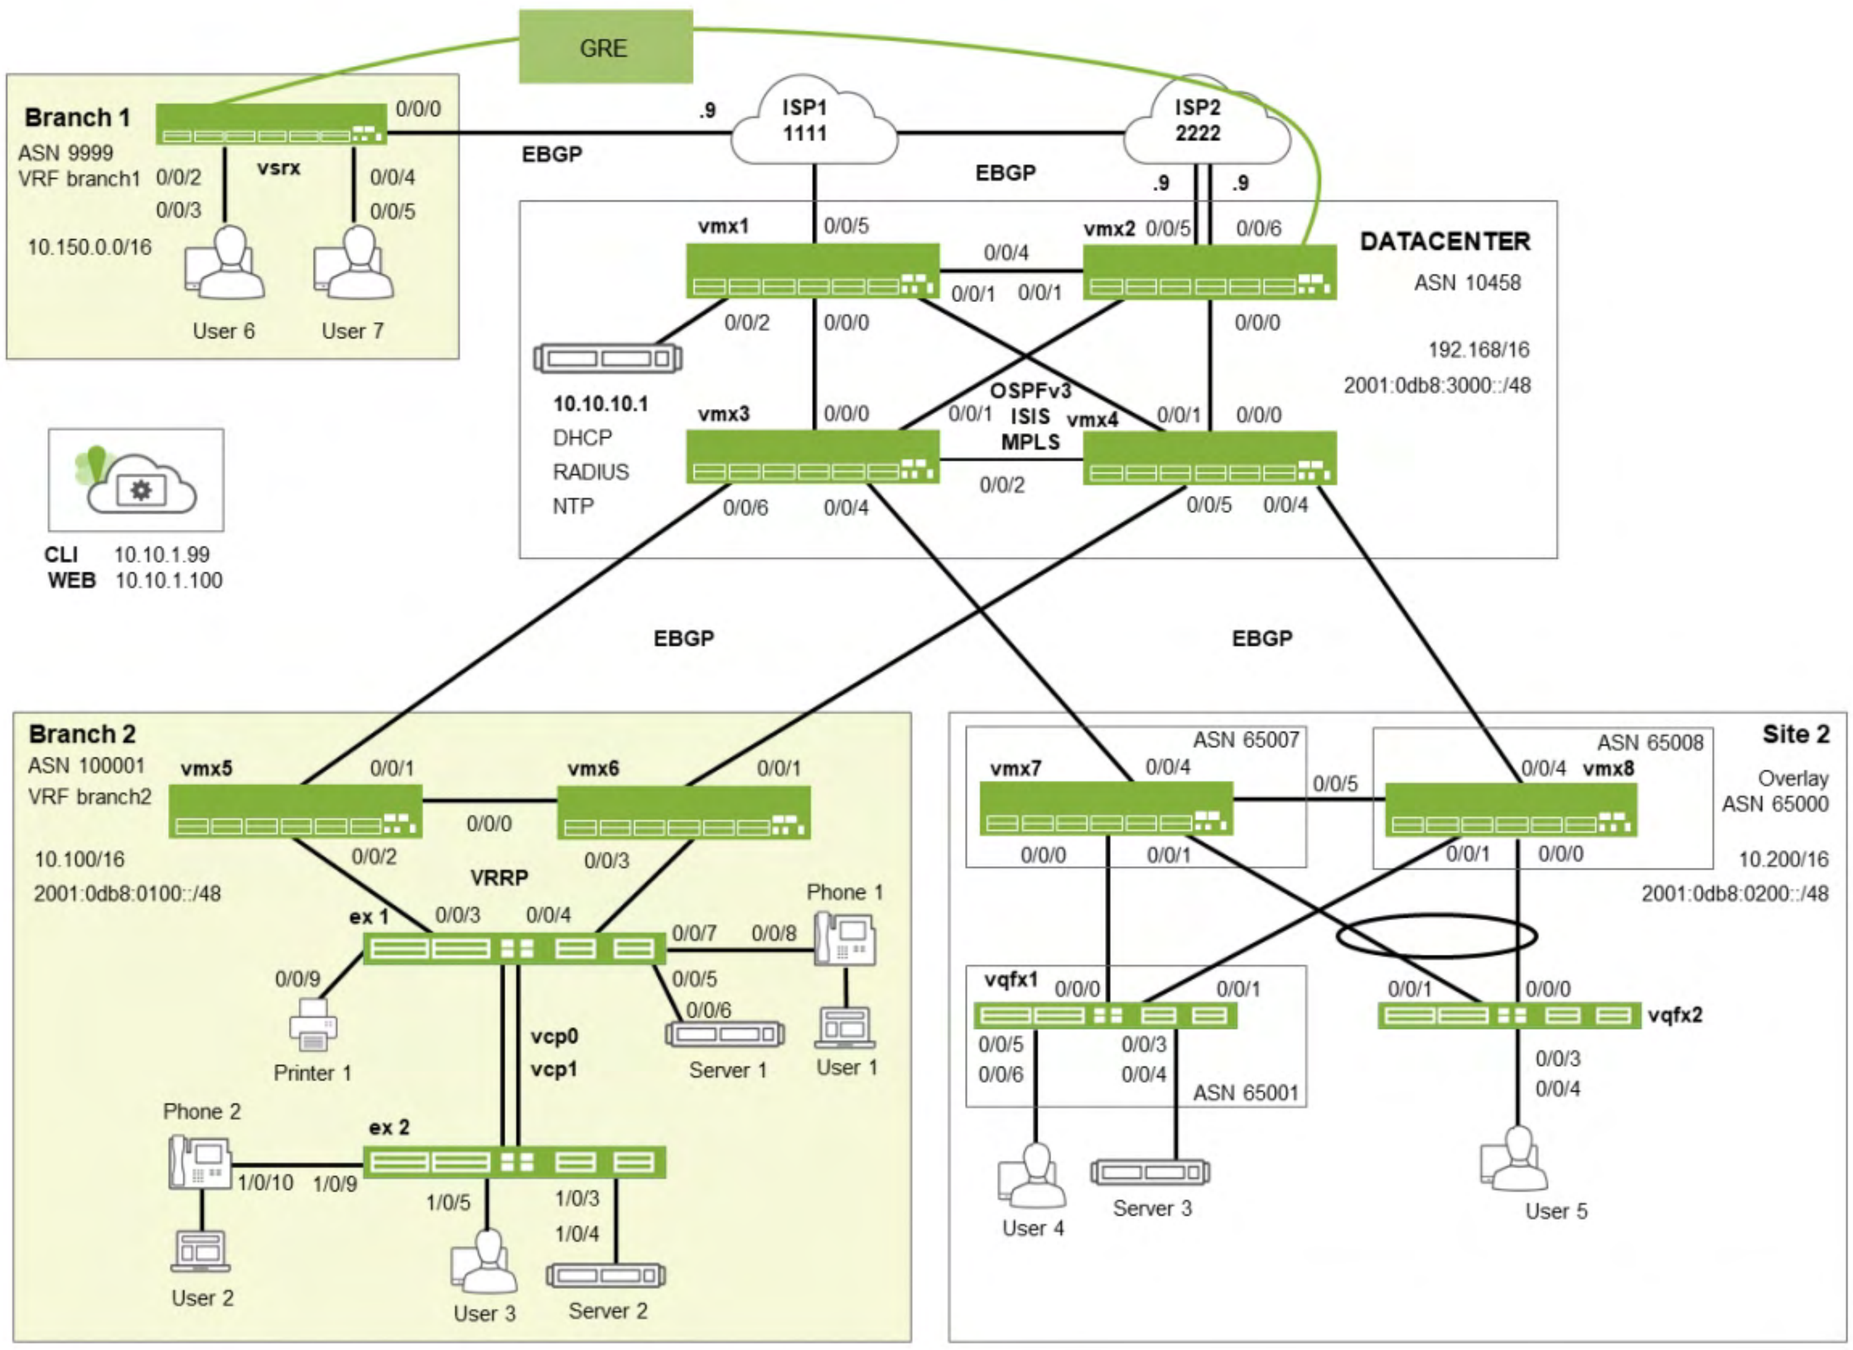

Data Center:

1) Enable MPLS forwarding between the data center vMX devices using the LDP protocol.

vMX1:

```
lab@vmx1# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 192.168.13.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::1/64;
    }
    family mpls;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vmx1# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::1/64;
    }
    family mpls;
}
```

```
lab@vmx1# show interfaces ge-0/0/4
unit 0 {
    family inet {
        address 192.168.12.1/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::1/64;
    }
    family mpls;
}

lab@vmx1# show protocols mpls
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/4.0;

lab@vmx1# show protocols ldp
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/4.0;
interface lo0.0;
```

vMX2:

```
lab@vmx2# show interfaces ge-0/0/0
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.24.2/24;
    }
    family iso;
    family inet6 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        address 2001:0db8:3000:0024::2/64;
    }
    family mpls;
}

lab@vmx2# show interfaces ge-0/0/1
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.23.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0023::2/64;
    }
    family mpls;
}

lab@vmx2# show interfaces ge-0/0/4
unit 0 {
    family inet {
        filter {
            input fbf-v4;
        }
        address 192.168.12.2/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0012::2/64;
    }
    family mpls;
}

lab@vmx2# show protocols mpls
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/4.0;

lab@vmx2# show protocols ldp
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/4.0;
interface lo0.0;
```

vMX3:

```
lab@vmx3# show interfaces ge-0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
unit 0 {
    family inet {
        address 192.168.13.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0013::3/64;
    }
    family mpls;
}

lab@vmx3# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.23.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0023::3/64;
    }
    family mpls;
}

lab@vmx3# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.3/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::3/64;
    }
    family mpls;
}

lab@vmx3# show protocols mpls
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/2.0;

lab@vmx3# show protocols ldp
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface lo0.0;
```

vMX4:

```
lab@vmx4# show interfaces ge-0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
unit 0 {
    family inet {
        address 192.168.24.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0024::4/64;
    }
    family mpls;
}

lab@vmx4# show interfaces ge-0/0/1
unit 0 {
    family inet {
        address 192.168.14.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0014::4/64;
    }
    family mpls;
}

lab@vmx4# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.34.4/24;
    }
    family iso;
    family inet6 {
        address 2001:0db8:3000:0034::4/64;
    }
    family mpls;
}

lab@vmx4# show protocols mpls
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/2.0;

lab@vmx4# show protocols ldp
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface lo0.0;
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

vMX1:

```
lab@vmx1# run show ldp neighbor
Address                                 Interface       Label space ID     Hold time
192.168.13.3                            ge-0/0/0.0      192.168.0.3:0        12
192.168.14.4                            ge-0/0/1.0      192.168.0.4:0        12
192.168.12.2                            ge-0/0/4.0      192.168.0.2:0        12

[edit]
lab@vmx1# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.2/32     *[LDP/9] 00:03:57, metric 1
                    > to 192.168.12.2 via ge-0/0/4.0
192.168.0.3/32     *[LDP/9] 00:03:56, metric 1
                    > to 192.168.13.3 via ge-0/0/0.0
192.168.0.4/32     *[LDP/9] 00:03:56, metric 1
                    > to 192.168.14.4 via ge-0/0/1.0
```

vMX2:

```
lab@vmx2# run show ldp neighbor
Address                                 Interface       Label space ID     Hold time
192.168.24.4                            ge-0/0/0.0      192.168.0.4:0        11
192.168.23.3                            ge-0/0/1.0      192.168.0.3:0        14
192.168.12.1                            ge-0/0/4.0      192.168.0.1:0        11

lab@vmx2# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:04:18, metric 1
                    > to 192.168.12.1 via ge-0/0/4.0
192.168.0.3/32     *[LDP/9] 00:04:18, metric 1
                    > to 192.168.23.3 via ge-0/0/1.0
192.168.0.4/32     *[LDP/9] 00:04:17, metric 1
                    > to 192.168.24.4 via ge-0/0/0.0
```

vMX3:

```
lab@vmx3# run show ldp neighbor
Address                                 Interface       Label space ID     Hold time
192.168.13.1                            ge-0/0/0.0      192.168.0.1:0        14
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.23.2                    ge-0/0/1.0      192.168.0.2:0        12
192.168.34.4                    ge-0/0/2.0      192.168.0.4:0        10

lab@vmx3# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:04:38, metric 1
                    > to 192.168.13.1 via ge-0/0/0.0
192.168.0.2/32     *[LDP/9] 00:04:38, metric 1
                    > to 192.168.23.2 via ge-0/0/1.0
192.168.0.4/32     *[LDP/9] 00:04:37, metric 1
                    > to 192.168.34.4 via ge-0/0/2.0
```

vMX4:

```
lab@vmx4# run show ldp neighbor
Address                         Interface       Label space ID     Hold time
192.168.24.2                    ge-0/0/0.0      192.168.0.2:0        11
192.168.14.1                    ge-0/0/1.0      192.168.0.1:0        10
192.168.34.3                    ge-0/0/2.0      192.168.0.3:0        12

lab@vmx4# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:04:58, metric 1
                    > to 192.168.14.1 via ge-0/0/1.0
192.168.0.2/32     *[LDP/9] 00:04:58, metric 1
                    > to 192.168.24.2 via ge-0/0/0.0
192.168.0.3/32     *[LDP/9] 00:04:58, metric 1
                    > to 192.168.34.3 via ge-0/0/2.0
```

2) Penultimate Hop Popping (PHP) should be disabled.

vMX1, vMX2, vMX3, and vMX4:

```
lab@vmx1# show protocols ldp
explicit-null;
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface ge-0/0/4.0;
interface lo0.0;
```

vMX1:

```
lab@vmx1# run show route table inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.2/32     *[LDP/9] 00:00:09, metric 1
                    > to 192.168.12.2 via ge-0/0/4.0, Push 0
192.168.0.3/32     *[LDP/9] 00:00:07, metric 1
                    > to 192.168.13.3 via ge-0/0/0.0, Push 0
192.168.0.4/32     *[LDP/9] 00:00:08, metric 1
                    > to 192.168.14.4 via ge-0/0/1.0, Push 0
```

vMX2:

```
lab@vmx2# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:01:15, metric 1
                    > to 192.168.12.1 via ge-0/0/4.0, Push 0
192.168.0.3/32     *[LDP/9] 00:01:15, metric 1
                    > to 192.168.23.3 via ge-0/0/1.0, Push 0
192.168.0.4/32     *[LDP/9] 00:01:14, metric 1
                    > to 192.168.24.4 via ge-0/0/0.0, Push 0
```

vMX3:

```
lab@vmx3# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:01:26, metric 1
                    > to 192.168.13.1 via ge-0/0/0.0, Push 0
192.168.0.2/32     *[LDP/9] 00:01:28, metric 1
                    > to 192.168.23.2 via ge-0/0/1.0, Push 0
192.168.0.4/32     *[LDP/9] 00:01:27, metric 1
                    > to 192.168.34.4 via ge-0/0/2.0, Push 0
```

vMX4:

```
lab@vmx4# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1/32     *[LDP/9] 00:01:44, metric 1
                    > to 192.168.14.1 via ge-0/0/1.0, Push 0
192.168.0.2/32     *[LDP/9] 00:01:44, metric 1
                    > to 192.168.24.2 via ge-0/0/0.0, Push 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.0.3/32     *[LDP/9] 00:01:44, metric 1
                    > to 192.168.34.3 via ge-0/0/2.0, Push 0
```

Configurations for Task Completion:

vMX1:

```
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/4.0
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/4.0
set protocols ldp interface lo0.0
set protocols ldp explicit-null
```

vMX2:

```
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/4.0
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/4.0
set protocols ldp interface lo0.0
set protocols ldp explicit-null
```

vMX3:

```
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0
set protocols ldp explicit-null
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX4:

```
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0
set protocols ldp explicit-null
```

Task 6: BGP

This task focuses on the implementation of the BGP. At various points in the topology you need to perform route redistribution between an IGP and BGP for end to end connectivity.

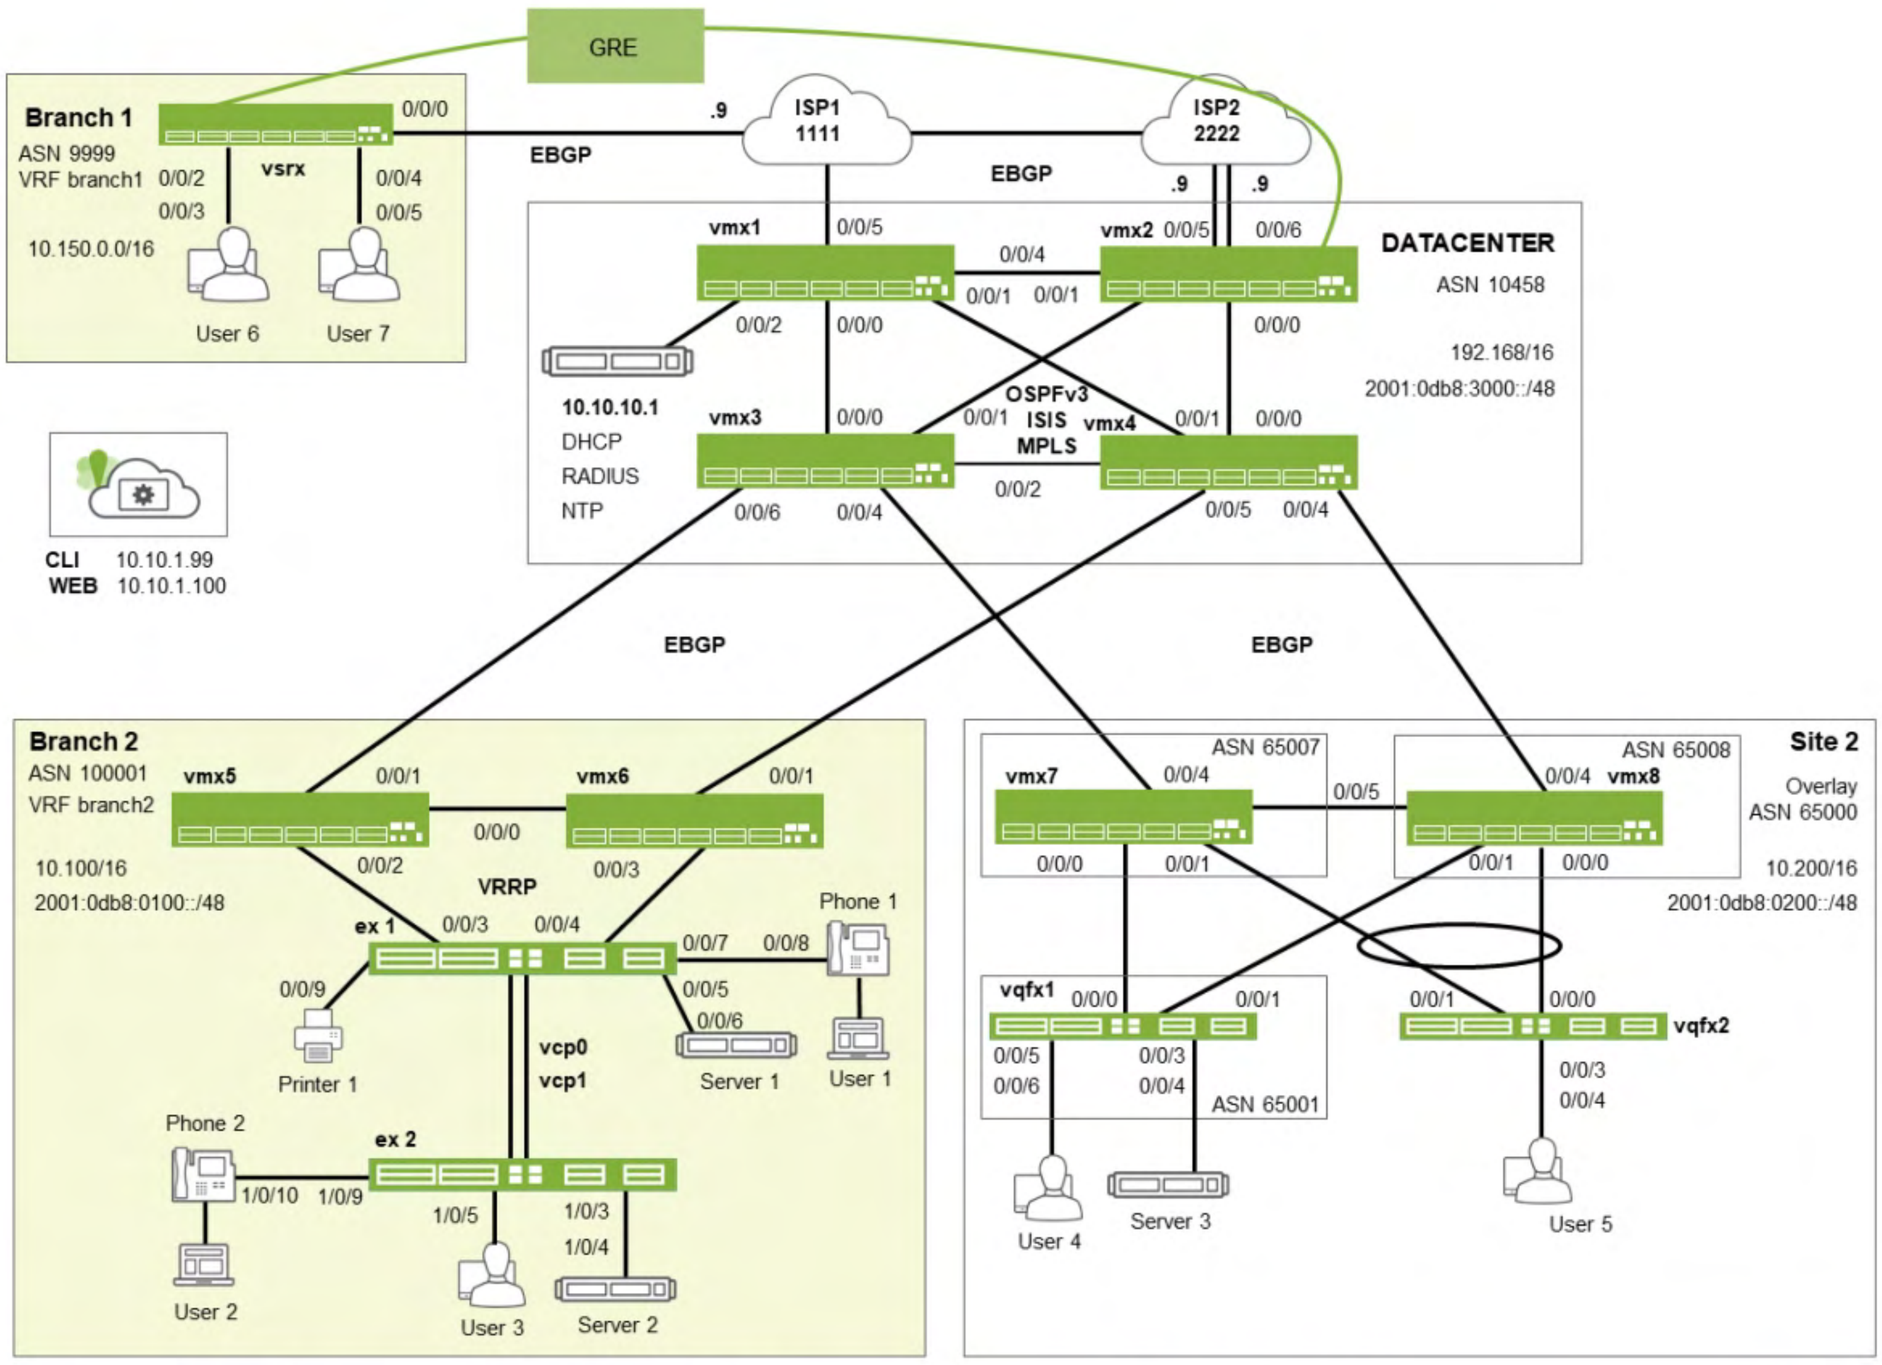

Device	Interface	Ipv4 Address	Ipv6 Address	Description

© 2020 Juniper Networks, Inc. All rights reserved.

vMX1	lo0	192.168.0.1/32	2001:0db8:3000:0000::1/128	IBGP source interface
	ge-0/0/5	172.16.1.1/24	2001:0db9:4000:0001::1/64	EBGP to AS 1111
vMX2	lo0	192.168.0.2/24	2001:0db8:3000:0000::2/128	IBGP source interface
				v4 EBGP to AS 2222
	ge-0/0/5	172.16.2.2/24	2001:0dba:4000:0002::2/64	v6 EBGP to AS 2222
	ge-0/0/6	172.16.3.2/24	Link local	v6 EBGP to AS 2222
vMX3	lo0	192.168.0.3/32	2001:0db8:3000:0000::3/128	IBGP source interface
	ge-0/0/4	192.168.37.3/24	2001:0db8:3000:0037::3/64	EBGP to AS 64513
	ge-0/0/6	192.168.35.3/24	2001:0db8:3000:0035::3/64	EBGP to AS 100001
vMX4	lo0	192.168.0.4/32	2001:0db8:3000:0000::4/128	IBGP source interface
	ge-0/0/4	192.168.48.4/24	2001:0db8:3000:0048::4/64	EBGP to AS 64513
	ge-0/0/5	192.168.46.4/24	2001:0db8:3000:0046::4/64	EBGP to AS 100001
vMX5	lo0	10.100.0.5/32	2001:0db8:0100:0000::5/128	IBGP source interface
	ge-0/0/1	192.168.35.5/24	2001:0db8:3000:0035::5/64	EBGP to AS 10458
vMX6	lo0	10.100.0.6/32	2001:0db8:0100:0000::6/128	IBGP source interface
	ge-0/0/1	192.168.46.6/24	2001:0db8:3000:0046::6/64	EBGP to AS 10458
vMX7	lo0	10.200.0.7/32	2001:0db8:3000:0000::7/128	IBGP source interface
	ge-0/0/4	192.168.37.7/24	2001:0db8:3000:0037::7/64	EBGP to AS 10458
vMX8	lo0	10.200.0.8/32	2001:0db8:3000:0000::8/128	IBGP source interface
	ge-0/0/4	192.168.48.8/24	2001:0db8:3000:0048::8/64	EBGP to AS 10458
vSRX	ge-0/0/0	10.1.1.2/30		EBGP to AS 1111

Table 48 BGP peering information

IBGP Data Center

1) Configure IBGP in the data center. Use the information as shown in the table above. vMX3 and vMX4 should not have an IBGP peering with each other. vMX1 and vMX2 should be configured as router reflectors and also accept routes received from each other.
2) All IBGP sessions on vMX1 and vMX2 should be configured without a neighbor statement.
3) Ensure IPv4, IPv6 and MPLS L3VPN NLRI is exchanged. Use separate peerings for IPv4 and IPv6 NLRI exchange.

vMX1:

```
lab@vmx1# show routing-options autonomous-system
10458;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.1;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    cluster 192.168.0.1;
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.1;
     family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.1;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::1;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::2;
}
```

vMX2:

```
lab@vmx2# show routing-options autonomous-system
10458;

lab@vmx2# show protocols bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
group ibgp {
    type internal;
    local-address 192.168.0.2;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    cluster 192.168.0.2;
    allow 192.168.0.0/29;
}
group rr {
    type internal;
    local-address 192.168.0.2;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.168.0.1;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    cluster 192.168.0.2;
    allow 2001:0db8:3000:0000::/64;
}
group rr-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::2;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::1;
}
```

vMX3:

```
lab@vmx3# show routing-options autonomous-system
10458;

lab@vmx3# show protocols bgp
group ibgp {
    type internal;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    local-address 192.168.0.3;
    family inet {
        unicast ;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    family inet6 {
        unicast;
    }
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
```

vMX4:

```
lab@vmx4# show routing-options autonomous-system
10458;

lab@vmx4# show protocols bgp
group ibgp {
    type internal;
    local-address 192.168.0.4;
    family inet {
        unicast ;
    }
    family inet-vpn {
        unicast;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    family inet6 {
        unicast {
            rib-group inetzero-to-branch2-v6;
        }
    }
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

vMX1:

```
lab@vmx1# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.0.2          10458         39         37       0       0       15:47 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
192.168.0.3          10458         21         19       0       0        8:10 Establ
  bgp.l3vpn.0: 0/0/0/0
192.168.0.4          10458          7          5       0       0        1:48 Establ
  bgp.l3vpn.0: 0/0/0/0
2001:db8:3000::2     10458         37         36       0       0       15:43 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::3     10458         20         19       0       0        8:01 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::4     10458          6          4       0       0        1:36 Establ
  inet6.0: 0/0/0/0
```

vMX2:

```
lab@vmx2# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.0.1          10458         39         39       0       0       16:22 Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
192.168.0.3          10458         22         21       0       0        8:41 Establ
  bgp.l3vpn.0: 0/0/0/0
192.168.0.4          10458          8          8       0       0        2:20 Establ
  bgp.l3vpn.0: 0/0/0/0
2001:db8:3000::1     10458         38         38       0       0       16:18 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::3     10458         21         20       0       0        8:33 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::4     10458          7          5       0       0        2:08 Establ
  inet6.0: 0/0/0/0
```

vMX3:

```
lab@vmx3# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.0.1          10458         21         21       0       0        9:07 Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  bgp.l3vpn.0: 0/0/0/0
192.168.0.2            10458         22         21       0       0        9:02 Establ
  bgp.l3vpn.0: 0/0/0/0
2001:db8:3000::1       10458          21         20       0       0        8:58 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::2       10458          21         20       0       0        8:55 Establ
  inet6.0: 0/0/0/0
```

vMX4:

```
lab@vmx4# run show bgp summary
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.0.1            10458         13         13       0       0        5:13 Establ
  bgp.l3vpn.0: 0/0/0/0
192.168.0.2            10458         14         13       0       0        5:09 Establ
  bgp.l3vpn.0: 0/0/0/0
2001:db8:3000::1       10458          13         12       0       0        5:01 Establ
  inet6.0: 0/0/0/0
2001:db8:3000::2       10458          12         12       0       0        4:57 Establ
  inet6.0: 0/0/0/0
```

4) Create an IPv4 default route on vMX1 and vMX2 and advertise these routes within the AS

5) vMX1 should not advertise the default route if the BGP peering with ISP1 goes down. vMX2 should not advertise the default route if the BGP peering with ISP2 goes down

vMX1:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}

lab@vmx1# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.1.9;
    }
    then accept;
}
term 2 {
    then reject;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}

lab@vmx1# show policy-options policy-statement 2ibgp
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        next-hop self;
        accept;
    }
}
term 4 {
    then reject;
}

lab@vmx1# show protocols bgp group ibgp export
export 2ibgp;
```

vMX2:

```
lab@vmx1# show routing-options
generate {
    route 0.0.0.0/0 policy if-isp-is-up;
}


lab@vmx2# show policy-options policy-statement if-isp-is-up
term 1 {
    from {
        protocol bgp;
        neighbor 172.16.200.176;
    }
    then accept;
}
term 2 {
    then reject;
}


lab@vmx2# show policy-options policy-statement 2ibgp
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        next-hop self;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        accept;
    }
}
term 4 {
    then reject;
}

lab@vmx2# show protocols bgp
group ibgp {
    export 2ibgp;
```

Verification:

Note: The EBGP sessions with ISP1 and ISP2 must be configured before you can test your solution

```
lab@vmx1# run show route

inet.0: 245 destinations, 470 routes (245 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 00:00:08
                    > to 172.16.1.9 via ge-0/0/5.0
...

lab@vmx1# run show bgp summary group isp1
Groups: 5 Peers: 7 Down peers: 0
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.9             1111        223        102       0       0        9:32 Establ
  inet.0: 134/222/222/0

lab@vmx1# set interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete


lab@vmx1# run show bgp summary | match 172.16.1.9
172.16.1.9             1111        419        218       0       1          30 Idle

lab@vmx1# run show route 0/0 exact
```

Since the EBGP peering with ISP1 is down, the generate route is not installed in the routing-table due to the applied policy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# delete interfaces ge-0/0/5 disable

[edit]
lab@vmx1# commit
commit complete
```

6) Create an IPv6 default route on vMX1 and vMX2 and advertise these routes within the AS

vMX1 & vMX2:

```
lab@vmx1# show routing-options
rib inet6.0 {
    static {
        route 0::/0 {
            discard;
            no-install;
        }
    }
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        accept;
    }
}
term 3 {
    then reject;
}

lab@vmx1# show protocols bgp group ibgp-v6 export
export 2ibgp-v6;
```

Verification:

```
lab@vmx1# run show route advertising-protocol bgp 2001:db8:3000::3

inet6.0: 102 destinations, 182 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                         100        I
```

IBGP Branch 2 and Site 2

© 2020 Juniper Networks, Inc. All rights reserved.

7) Configure IPv4 and IPv6 IBGP peerings in Branch 2 between vMX5 and vMX6. Peer using the loopback interfaces. The ASN is 100001.

vMX5:

```
lab@vmx5# show routing-options
autonomous-system 100001;

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    neighbor 2001:0db8:0100:0000::6;
}
```

vMX6:

```
lab@vmx6# show routing-options
autonomous-system 100001;

lab@vmx6# show protocols bgp
group ibgp {
    type internal;
    neighbor 10.100.0.5 {
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    neighbor 2001:0db8:0100:0000::5;
}
```

Verification:

```
lab@vmx5# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                       2          1          0          0          0          0
inet.0
                       8          4          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.100.0.6           100001       2064       2064       0       0    15:31:15 0/0/0/0
0/0/0/0
2001:db8:100::6      100001       2052       2051       0       0    15:31:15 Establ
  inet6.0: 0/0/0/0
```

EBGP Data Center

8) Configure an IPv4 and IPv6 EBGP peering between vMX1 and ISP1 and an IPv4 EBGP peering between vMX2 and ISP2.
9) Ensure that vMX2 can load balance towards ISP2 for IPv4 destinations. You are allowed to configure a static IPv4 route for this purpose. Use troubleshooting on vMX2 to find out how to configure the IPv4 peering with ISP2.

vMX1:

```
lab@vmx1# show protocols bgp
group isp1 {
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
group isp1-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0db9:4000:0001::9 {
        peer-as 1111;
   }
```

Verification:

```
lab@vmx1# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Unconfigured peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                      80         79          0          0          0          0
inet.0
                     226        224          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.9             1111        253         50       0       2       20:45
222/222/222/0        0/0/0/0
2001:db9:4000:1::9        1111        120         47       0       2       20:31 Establ
  inet6.0: 78/78/78/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

The task states to establish one IPv4 peering with ISP2 and we need to provide load balancing using both links towards ISP2. As ISP2 is outside of our control it must be preconfigured. This means that if we can debug the BGP packets received from ISP2 we might be able to find out valuable information to configure the BGP peering. Note that the BGP packets received from ISP2 can come in on interface ge-0/0/5 or interface ge-0/0/6. Keep this in mind when debugging the BGP packets.

In this example we are going to use the “monitor traffic” command to verify incoming BGP packets destinated to the control plane of vMX2.

```
lab@vmx2# run monitor traffic interface ge-0/0/5 extensive matching tcp
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on ge-0/0/5, capture size 1514 bytes

...
 PFE proto 2 (ipv4): (tos 0xc0, ttl  64, id 53971, offset 0, flags [none], proto: TCP
(6), length: 64) 172.16.200.176.55517 > 192.168.0.2.bgp: S 3936139201:3936139201(0) win
16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 3639125731 0,sackOK,eol>
23:50:28.618250 Out
        Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
```

The output above shows that we received an incoming tcp packet from 172.16.200.176 destined to the loopback address of vMX2 (192.168.0.2), port BGP.

We now know the source and destination address for the IPv4 BGP peering with ISP2. Due to the load balancing requirement we must make sure that IP 172.16.200.176 can be reached using both interfaces towards ISP2. For this we must configure a static route.

vMX2:

```
lab@vmx2# show routing-options
static {
    route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ];
}
```

Now configure the BGP peering with ISP2. Keep in mind that the “multihop” statement is needed as EBGP sessions by default are single-hop sessions (TTL checking).

vMX2:

```
lab@vmx2# show protocols bgp
group isp2 {
    neighbor 172.16.200.176 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        multihop;
        peer-as 2222;
    }
}
```

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called per-packet load balancing. The naming may be counter-intuitive. However, Junos per-packet load balancing is functionally equivalent to what other vendors may term per-flow load balancing. Configure load balancing to spread traffic across both links towards ISP2.

vMX2:

```
lab@vmx2# show policy-options policy-statement lbpp
term 1 {
    then {
        load-balance per-packet;
    }
}

lab@vmx2# show routing-options forwarding-table
export lbpp;
```

Verification:

```
lab@vmx2# run show bgp summary
…
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.200.176         2222        206        124       0       0        1:05 Establ
  inet.0: 89/222/222/0

lab@vmx2# run show route receive-protocol bgp 172.16.200.176

inet.0: 250 destinations, 255 routes (250 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 1.84.160.0/20           172.16.200.176                          2222 1111 33112 I
...

[edit]
lab@vmx2# run show route forwarding-table destination 1.84.160.1
...

Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
1.84.160.0/20      user     0                    indr  1048575   225
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                                   ulst   1048574      4
172.16.2.9                         ucst       660      3 ge-0/0/5.0
172.16.3.9                         ucst       659      4 ge-0/0/6.0
...
```

10) Configure two IPv6 peerings between vMX2 and ISP2. The IPv6 peering on the ge-0/0/6 interface should be established using the link local address. Use troubleshooting on vMX2 to find out how to configure the link local peering with ISP2.

First configure the IPv6 peering with ISP2 for the ge-0/0/5 link.

vMX2:

```
lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
}
```

The IPv6 peering over the ge-0/0/6 interface must be established using the IPv6 link local address. We need to find out the peering addresses using the “monitor traffic interface” command

```
lab@vmx2# run monitor traffic interface ge-0/0/6 no-resolve matching tcp
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on ge-0/0/6, capture size 96 bytes

01:19:33.147860 Out IP6 truncated-ip6 - 31 bytes missing!fe80::250:56ff:fea9:43ac.64092
> fe80::250:56ff:fea9:7fec.179: P 3276364454:3276364473(19) ack 2827393527 win 16384
<[|tcp]>
01:19:33.257758  In IP6 fe80::250:56ff:fea9:7fec.179 > fe80::250:56ff:fea9:43ac.64092: .
ack 19 win 16384 <nop,nop,timestamp 3644476929[|tcp]>
^C

lab@vmx2# run show interfaces ge-0/0/6 | match fe80
        Destination: fe80::/64, Local: fe80::250:56ff:fea9:43ac

lab@vmx2# show protocols bgp
group isp2-v6 {
    family inet6 {
        unicast;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    neighbor 2001:0dba:4000:0002::9 {
        peer-as 2222;
    }
    neighbor fe80::250:56ff:fea9:7fec {
        local-interface ge-0/0/6.0;
        peer-as 2222;
    }
}
```

As link-local addresses are only unique per interface, it is required to specify the “local-interface” command for the peering so that Junos known on which interface the peering must be established.

11) Advertise the IPv4 summary routes received from Branch2 and Site 2 to both ISPs. Also advertise the data center IPv4 summary route to the ISPs

vMX1:

```
lab@vmx1# show routing-options aggregate
route 192.168.0.0/16;

lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        accept;
    }
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp
group isp1 {
    export 2ebgp;
    neighbor 172.16.1.9 {
        peer-as 1111;
    }
}
```

Verification:

```
lab@vmx1# run show route advertising-protocol bgp 172.16.1.9
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.0.0/16          Self                                    I
```

vMX2:

```
lab@vmx2# show routing-options aggregate
route 192.168.0.0/16;

lab@vmx2# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
lab@vmx2# show protocols bgp group isp2
export 2ebgp;
neighbor 172.16.200.176 {
    multihop;
    peer-as 2222;
}
```

Verification:

```
lab@vmx2# run show route advertising-protocol bgp 172.16.200.176

inet.0: 251 destinations, 255 routes (251 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.0.0/16          Self                                    I
```

12) Advertise a single /32 route to both ISPs that encompasses the IPv6 networks of the data center, Branch 2 and Site 2.

vMX1:

```
lab@vmx1# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;

lab@vmx1# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx1# show protocols bgp group isp1-v6
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0db9:4000:0001::9 {
    peer-as 1111;
}
```

vMX2:

```
lab@vmx2# show routing-options rib inet6.0 aggregate
route 2001:0db8::/32;
```

```
lab@vmx2# show policy-options policy-statement 2ebgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx2# show protocols bgp group isp2-v6
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0dba:4000:0002::9 {
    peer-as 2222;
}
neighbor fe80::250:56ff:fea9:7fec {
    local-interface ge-0/0/6.0;
    peer-as 2222;
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX1:

```
lab@vmx1# run show route advertising-protocol bgp 2001:0db9:4000:0001::9

inet6.0: 102 destinations, 181 routes (102 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                     I
```

vMX2:

```
lab@vmx2# run show route advertising-protocol bgp 2001:0dba:4000:0002::9

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                     I

lab@vmx2# run show route advertising-protocol bgp fe80::250:56ff:fea9:7fec

inet6.0: 105 destinations, 184 routes (105 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:db8::/32           Self                                     I
```

13) Configure vMX1 so that ISP2 is the primary path inbound to your AS.

vMX1:

```
lab@vmx1# show policy-options policy-statement 2ebgp
term 1 {
    from {
        route-filter 192.168.0.0/16 exact;
        route-filter 10.100.0.0/16 exact;
        route-filter 10.200.0.0/16 exact;
    }
    then {
        as-path-prepend "10458 10458 10458 10458";
        accept;
    }
}
term 2 {
    then reject;
}

lab@vmx1# show protocols bgp group isp1
export 2ebgp;
neighbor 172.16.1.9 {
    peer-as 1111;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verification:

Note: You can review the summary route on the vr-device using the lab/lab123 credentials.

VR-Device:

```
lab@vr-device# run show route 192.168.0.0

ISP1.inet.0: 232 destinations, 235 routes (232 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 1d 22:41:14, localpref 100
                      AS path: 2222 10458 I, validation-state: unverified
                    >  to 10.20.20.2 via lt-0/0/0.1
                    [BGP/170] 2d 03:51:21, localpref 100
                      AS path: 10458 10458 10458 10458 10458 I, validation-state:
unverified
                    >  to 172.16.1.1 via ge-0/0/2.0
```

14) Ensure the ISP2 is the primary exit point for IPv4 and IPV6 destinations outside your AS.

vMX1:

```
lab@vmx1# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}
```

```
lab@vmx1# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
```

vMX2:

```
lab@vmx2# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    then {
        next-hop self;
    }
}

lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        accept;
    }
}
```

```
lab@vmx2# show protocols bgp | display set | match nhs
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs

lab@vmx2# show protocols bgp group isp2
import fromisp2;
export 2ebgp;
neighbor 172.16.200.176 {
    multihop;
    peer-as 2222;
}

lab@vmx2# show protocols bgp group isp2-v6
import fromisp2;
family inet6 {
    unicast;
}
export 2ebgp-v6;
neighbor 2001:0dba:4000:0002::9 {
    peer-as 2222;
}
neighbor fe80::250:56ff:fea9:7fec {
    local-interface ge-0/0/6.0;
    peer-as 2222;
}
```

Verification:

```
lab@vmx1# run show route protocol bgp

inet.0: 246 destinations, 472 routes (245 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
1.84.160.0/20      *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 1111 33112 I, validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 33112 I, validation-state: unverified
                    > to 172.16.1.9 via ge-0/0/5.0
1.176.0.0/12       *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 2222 33112 49129 16320 52954 I, validation-state:
unverified
                    > to 172.16.1.9 via ge-0/0/5.0
3.189.0.0/16       *[BGP/170] 01:11:57, localpref 150, from 192.168.0.2
                      AS path: 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                    > to 192.168.12.2 via ge-0/0/4.0
                    [BGP/170] 02:59:09, localpref 100
                      AS path: 1111 2222 33112 41164 7136 6370 35389 39823 16965 I,
validation-state: unverified
                    > to 172.16.1.9 via ge-0/0/5.0
...
```

15) Prefixes from Site 2 that are advertised to both ISPs should not contain the ASN of Site 2.

vMX1:

```
lab@vmx1# show protocols bgp | display set | match remove
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
```

vMX2:

```
lab@vmx2# show protocols bgp | display set | match remove-private
set protocols bgp group isp2 remove-private
set protocols bgp group isp2-v6 remove-private
```

16) vMX3 and vMX4 should learn the ISP prefixes from vMX1 and vMX2 through BGP. Use well known communities to ensure that vMX3 and vMX4 do not advertise the ISP prefixes to Branch 2 and Site 2.

vMX1:

```
lab@vmx1# show policy-options community no-export
members no-export;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx1# show policy-options policy-statement fromisp1
term 1 {
    then {
        community set no-export;
        accept;
    }
}

lab@vmx1# show protocols bgp group isp1
import fromisp1;
export 2ebgp;
remove-private;
neighbor 172.16.1.9 {
    peer-as 1111;
}
```

vMX2:

```
lab@vmx2# show policy-options community no-export
members no-export;


lab@vmx2# show policy-options policy-statement fromisp2
term 1 {
    then {
        local-preference 150;
        community set no-export;
        accept;
    }
}
```

17) Site 2 should learn the IPv4 and IPv6 default route from the data center. Site 2 should learn the management network 10.10.10/24 that originated in the data center and the IPv4 summarized networks from Branch 2 and the data center.

18) Ensure the devices in Branch1, Branch 2 and Site2 have full reachability to each other and the internet.

NOTE: The solution for Branch 2 is shown in VPN section of this superlab.

vMX1:

```
lab@vmx1# show policy-options policy-statement 2ibgp
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    then accept;
}
term 2 {
    from {
        protocol direct;
        route-filter 10.10.10.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        next-hop self;
        accept;
    }
}
term 4 {
    then reject;
}

lab@vmx1# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        next-hop self;
        accept;
    }
}
term 3 {
    then reject;
}
```

vMX2:

```
lab@vmx2# show policy-options policy-statement 2ibgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
term 1 {
    from {
        protocol aggregate;
        route-filter 192.168.0.0/16 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol ospf;
        route-filter 10.150.99.0/24 exact;
    }
    then accept;
}
term 3 {
    from {
        protocol aggregate;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
term 4 {
    then reject;
}

[edit]
lab@vmx2# show policy-options policy-statement 2ibgp-v6
term 1 {
    from {
        protocol aggregate;
        route-filter 2001:0db8::/32 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol static;
        route-filter ::/0 exact;
    }
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
term 3 {
    then reject;
}
```

vMX3:

```
lab@vmx3# show policy-options policy-statement nhs
term 1 {
    from {
        protocol bgp;
        external;
    }
    then {
        next-hop self;
    }
}

lab@vmx3# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.3;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    export nhs;
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    family inet6 {
        unicast;
    }
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site2 {
    neighbor 192.168.37.7 {
        peer-as 65007;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0037::7 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        family inet6 {
            unicast;
        }
        peer-as 65007;
    }
}
```

vMX4:

```
lab@vmx4# show policy-options
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx4# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.4;
    family inet {
        unicast;
    }
    family inet-vpn {
        unicast;
    }
    export nhs;
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    family inet6 {
        unicast;
    }
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
group site2 {
    neighbor 192.168.48.8 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        peer-as 65008;
    }
}
group site2-v6 {
    neighbor 2001:0db8:3000:0048::8 {
        family inet6 {
            unicast;
        }
        peer-as 65008;
    }
}
```

vMX5:

```
lab@vmx5# show policy-options
...
policy-statement 2bgp {
    term 1 {
        from {
            protocol static;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol static;
            route-filter 2001:0db8:0100::/48 exact;
        }
        then accept;
    }
}
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx5# show protocols bgp
group ibgp {
    type internal;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    export nhs;
    neighbor 10.100.0.6 {
        local-address 10.100.0.5;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::5;
    export nhs;
    neighbor 2001:0db8:0100:0000::6;
}
group dc {
    export 2bgp;
    neighbor 192.168.35.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0035::3 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vMX6:

```
lab@vmx6# show policy-options
...
policy-statement 2bgp {
    term 1 {
        from {
            protocol static;
            route-filter 10.100.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol static;
            route-filter 2001:0db8:0100::/48 exact;
        }
        then accept;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
policy-statement nhs {
    term 1 {
        from {
            protocol bgp;
            external;
        }
        then {
            next-hop self;
        }
    }
}

lab@vmx6# show protocols bgp
group ibgp {
    type internal;
    export nhs;
    neighbor 10.100.0.5 {
        local-address 10.100.0.6;
    }
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:0100:0000::6;
    export nhs;
    neighbor 2001:0db8:0100:0000::5;
}
group dc {
    export 2bgp;
    neighbor 192.168.46.4 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    neighbor 2001:0db8:3000:0046::4 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

vSRX1:

```
lab@vsrx# show routing-options
static {
    route 192.168.0.2/32 next-hop 10.1.1.1;
}
aggregate {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    route 10.150.0.0/16;
}

lab@vsrx# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            route-filter 10.150.0.0/16 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement frombgp {
    term 1 {
        from {
            route-filter 192.168.0.0/16 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
lab@vsrx# show protocols bgp
group ebgp {
    import frombgp;
    export 2bgp;
    neighbor 10.1.1.1 {
        peer-as 1111;
        local-as 9999;
    }
}
```

Verification Branch 1:

```
lab@vsrx# run show bgp summary
Threading mode: BGP I/O
Groups: 1 Peers: 1 Down peers: 0
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.1                1111        201          1       0       0           0 1/224/1/0
0/0/0/0

[edit]
lab@vsrx# run show route protocol bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[BGP/170] 00:49:52, localpref 100
                      AS path: 1111 2222 10458 I, validation-state: unverified
                    >  to 10.1.1.1 via ge-0/0/0.0
...
```

```
lab@vsrx# run show route advertising-protocol bgp 10.1.1.1

inet.0: 238 destinations, 239 routes (14 active, 0 holddown, 224 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.150.0.0/16           Self                                    I
```

As Branch2 and Site2 announce more specifics / contributing routes of the aggregate route announced in the data center you need to ensure that vMX1 and vMX2 create the aggregate route 2001:db8::/32 based on local contributing routes. We can do this with route policy as shown below. If you forget this step, the aggregate route might not be accepted by Branch1 or Site2 due to the BGP AS loop prevention mechanism. The reason is that the aggregate route inherits the AS path of the contributing routes.

vMX1:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
}
```

vMX2:

```
lab@vmx1# show policy-options policy-statement aggregate-contributor
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        route-filter 2001:db8:3000::/48 orlonger;
    }
    then accept;
}
term 2 {
    then reject;
}
```

```
lab@vmx1# show routing-options
    aggregate {
        route 2001:0db8::/32 policy aggregate-contributor;
    }
}
```

vMX7:

```
lab@vmx7# show protocols bgp
group dc {
    export 2bgp;
    local-as 65007;
    neighbor 192.168.37.3 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    local-as 65007;
    neighbor 2001:0db8:3000:0037::3 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
```

```
lab@vmx7# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}
...

lab@vmx7# show policy-options
policy-statement 2bgp {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
```

vMX8:

```
lab@vmx8# show protocols bgp
group dc {
    export 2bgp;
    local-as 65008;
    neighbor 192.168.48.4 {
        peer-as 10458;
    }
}
group dc-v6 {
    export 2bgp-v6;
    local-as 65008;
    neighbor 2001:0db8:3000:0048::4 {
        family inet6 {
            unicast;
        }
        peer-as 10458;
    }
}
lab@vmx8# show routing-options
rib inet6.0 {
    aggregate {
        route 2001:0db8:0200::/48;
    }
}
aggregate {
    route 10.200.0.0/16;
}
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx8# show policy-options
policy-statement 2bgp {
    term 1 {
        from {
            protocol aggregate;
            route-filter 10.200.0.0/16 exact;
        }
        then accept;
    }
}
policy-statement 2bgp-v6 {
    term 1 {
        from {
            protocol aggregate;
            route-filter 2001:0db8:0200::/48 exact;
        }
        then accept;
    }
}
```

Verification Site 2:

vMX7:

```
lab@vmx7# run show route protocol bgp

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:02:59, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
10.10.10.0/24      *[BGP/170] 00:02:59, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
10.150.99.0/24     *[BGP/170] 00:02:59, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0
192.168.0.0/16     *[BGP/170] 00:02:59, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 192.168.37.3 via ge-0/0/4.0

inet6.0: 9 destinations, 9 routes (8 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

::/0               *[BGP/170] 00:02:55, localpref 100
                      AS path: 10458 I, validation-state: unverified
                    > to 2001:db8:3000:37::3 via ge-0/0/4.0
2001:db8::/32      *[BGP/170] 00:02:55, localpref 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      AS path: 10458 I, validation-state: unverified
    > to 2001:db8:3000:37::3 via ge-0/0/4.0
```

Configurations for Task Completion:

vMX1:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.1
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp cluster 192.168.0.1
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.1
set protocols bgp group rr family inet unicast
set protocols bgp group rr family inet-vpn unicast
set protocols bgp group rr neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.1
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::1
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::2
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.1.9
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set policy-options policy-statement 2ibgp term 3 from protocol aggregate
set policy-options policy-statement 2ibgp term 3 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2ibgp term 3 then next-hop self
set policy-options policy-statement 2ibgp term 3 then accept
set policy-options policy-statement 2ibgp term 4 then reject
set protocols bgp group ibgp export 2ibgp
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
set protocols bgp group isp1 neighbor 172.16.1.9 peer-as 1111
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp1 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp1-v6 export 2ebgp-v6
set protocols bgp group isp1-v6 family inet6 unicast
set protocols bgp group isp1-v6 neighbor 2001:0db9:4000:0001::9 peer-as 1111
set policy-options policy-statement 2ebgp term 1 then as-path-prepend "10458 10458 10458
10458"
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set protocols bgp group isp1 remove-private
set protocols bgp group isp1-v6 remove-private
set policy-options community no-export members no-export
set policy-options policy-statement fromisp1 term 1 then community set no-export
set policy-options policy-statement fromisp1 term 1 then accept
set protocols bgp group isp1 import fromisp1
set policy-options policy-statement 2ibgp term 1 from protocol aggregate
set policy-options policy-statement 2ibgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ibgp term 1 then accept
set policy-options policy-statement 2ibgp term 2 from protocol direct
set policy-options policy-statement 2ibgp term 2 from route-filter 10.10.10.0/24 exact
set policy-options policy-statement 2ibgp term 2 then accept
insert policy-options policy-statement 2ibgp term 1 before term 3
insert policy-options policy-statement 2ibgp term 2 before term 3
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
set policy-options policy-statement 2ibgp-v6 term 2 then next-hop self

set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::/48 orlonger
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX2:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.2
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp cluster 192.168.0.2
set protocols bgp group ibgp allow 192.168.0.0/29
set protocols bgp group rr type internal
set protocols bgp group rr local-address 192.168.0.2
set protocols bgp group rr family inet unicast
set protocols bgp group rr family inet-vpn unicast
set protocols bgp group rr neighbor 192.168.0.1
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 cluster 192.168.0.2
set protocols bgp group ibgp-v6 allow 2001:0db8:3000:0000::/64
set protocols bgp group rr-v6 type internal
set protocols bgp group rr-v6 local-address 2001:0db8:3000:0000::2
set protocols bgp group rr-v6 family inet6 unicast
set protocols bgp group rr-v6 neighbor 2001:0db8:3000:0000::1
set routing-options generate route 0.0.0.0/0 policy if-isp-is-up
set policy-options policy-statement if-isp-is-up term 1 from protocol bgp
set policy-options policy-statement if-isp-is-up term 1 from neighbor 172.16.200.176
set policy-options policy-statement if-isp-is-up term 1 then accept
set policy-options policy-statement if-isp-is-up term 2 then reject
set policy-options policy-statement 2ibgp term 3 from protocol aggregate
set policy-options policy-statement 2ibgp term 3 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2ibgp term 3 then next-hop self
set policy-options policy-statement 2ibgp term 3 then accept
set policy-options policy-statement 2ibgp term 4 then reject
set protocols bgp group ibgp export 2ibgp
set routing-options rib inet6.0 static route 0::/0 discard
set routing-options rib inet6.0 static route 0::/0 no-install
set policy-options policy-statement 2ibgp-v6 term 2 from protocol static
set policy-options policy-statement 2ibgp-v6 term 2 from route-filter ::/0 exact
set policy-options policy-statement 2ibgp-v6 term 2 then accept
set policy-options policy-statement 2ibgp-v6 term 3 then reject
set protocols bgp group ibgp-v6 export 2ibgp-v6
set routing-options static route 172.16.200.176/32 next-hop [ 172.16.2.9 172.16.3.9 ]
set protocols bgp group isp2 neighbor 172.16.200.176 multihop
set protocols bgp group isp2 neighbor 172.16.200.176 peer-as 2222
set policy-options policy-statement lbpp term 1 then load-balance per-packet
set routing-options forwarding-table export lbpp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group isp2-v6 family inet6 unicast
set protocols bgp group isp2-v6 neighbor 2001:0dba:4000:0002::9 peer-as 2222
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec local-interface ge-
0/0/6.0
set protocols bgp group isp2-v6 neighbor fe80::250:56ff:fea9:7fec peer-as 2222
set routing-options aggregate route 192.168.0.0/16
set policy-options policy-statement 2ebgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2ebgp term 1 then accept
set policy-options policy-statement 2ebgp term 2 then reject
set protocols bgp group isp2 export 2ebgp
set routing-options rib inet6.0 aggregate route 2001:0db8::/32
set policy-options policy-statement 2ebgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ebgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ebgp-v6 term 1 then accept
set policy-options policy-statement 2ebgp-v6 term 2 then reject
set protocols bgp group isp2-v6 export 2ebgp-v6
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group rr export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group rr-v6 export nhs
set policy-options policy-statement fromisp2 term 1 then local-preference 150
set policy-options policy-statement fromisp2 term 1 then accept
set protocols bgp group isp2 import fromisp2
set protocols bgp group isp2-v6 import fromisp2
set policy-options community no-export members no-export
set policy-options policy-statement fromisp2 term 1 then community set no-export
set policy-options policy-statement 2ibgp term 1 from protocol aggregate
set policy-options policy-statement 2ibgp term 1 from route-filter 192.168.0.0/16 exact
set policy-options policy-statement 2ibgp term 1 then accept
set policy-options policy-statement 2ibgp term 2 from protocol ospf
set policy-options policy-statement 2ibgp term 2 from route-filter 10.150.99.0/24 exact
set policy-options policy-statement 2ibgp term 2 then accept
insert policy-options policy-statement 2ibgp term 1 before term 3
insert policy-options policy-statement 2ibgp term 2 before term 3
set policy-options policy-statement 2ibgp term 3 then local-preference 150
set policy-options policy-statement 2ibgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2ibgp-v6 term 1 from route-filter 2001:0db8::/32
exact
set policy-options policy-statement 2ibgp-v6 term 1 then accept
insert policy-options policy-statement 2ibgp-v6 term 1 before term 2
set policy-options policy-statement 2ibgp-v6 term 2 then next-hop self
set policy-options policy-statement 2ibgp-v6 term 2 then local-preference 150
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement aggregate-contributor term 1 from route-filter
2001:db8:3000::/48 orlonger
set policy-options policy-statement aggregate-contributor term 1 then accept
set policy-options policy-statement aggregate-contributor term 2 then reject
set routing-options rib inet6.0 aggregate route 2001:0db8::/32 policy aggregate-
contributor
```

vMX3:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.3
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::3
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site2 neighbor 192.168.37.7 peer-as 65007
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0037::7 peer-as 65007
```

vMX4:

```
set routing-options autonomous-system 10458
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 192.168.0.4
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp neighbor 192.168.0.1
set protocols bgp group ibgp neighbor 192.168.0.2
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 family inet6 unicast
set protocols bgp group ibgp-v6 local-address 2001:0db8:3000:0000::4
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::1
set protocols bgp group ibgp-v6 neighbor 2001:0db8:3000:0000::2
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp advertise-inactive
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group site2 neighbor 192.168.48.8 peer-as 65008
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 family inet6 unicast
set protocols bgp group site2-v6 neighbor 2001:0db8:3000:0048::8 peer-as 65008
```

vMX5:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 10.100.0.6 local-address 10.100.0.5
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::5
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::6
set policy-options policy-statement 2bgp term 1 from protocol static
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol static
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.35.3 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0035::3 peer-as 10458
```

vMX6:

```
set routing-options autonomous-system 100001
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 10.100.0.5 local-address 10.100.0.6
set protocols bgp group ibgp-v6 type internal
set protocols bgp group ibgp-v6 local-address 2001:0db8:0100:0000::6
set protocols bgp group ibgp-v6 neighbor 2001:0db8:0100:0000::5
set policy-options policy-statement 2bgp term 1 from protocol static
set policy-options policy-statement 2bgp term 1 from route-filter 10.100.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol static
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0100::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
set policy-options policy-statement nhs term 1 from protocol bgp
set policy-options policy-statement nhs term 1 from external
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement nhs term 1 then next-hop self
set protocols bgp group ibgp export nhs
set protocols bgp group ibgp-v6 export nhs
set protocols bgp group dc export 2bgp
set protocols bgp group dc neighbor 192.168.46.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0046::4 peer-as 10458
```

vMX7:

```
set protocols bgp group dc export 2bgp
set protocols bgp group dc local-as 65007
set protocols bgp group dc neighbor 192.168.37.3 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 local-as 65007
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0037::3 peer-as 10458
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
```

vMX8:

```
set protocols bgp group dc export 2bgp
set protocols bgp group dc local-as 65008
set protocols bgp group dc neighbor 192.168.48.4 peer-as 10458
set protocols bgp group dc-v6 export 2bgp-v6
set protocols bgp group dc-v6 local-as 65008
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 family inet6 unicast
set protocols bgp group dc-v6 neighbor 2001:0db8:3000:0048::4 peer-as 10458
set routing-options rib inet6.0 aggregate route 2001:0db8:0200::/48
set routing-options aggregate route 10.200.0.0/16
set policy-options policy-statement 2bgp term 1 from protocol aggregate
set policy-options policy-statement 2bgp term 1 from route-filter 10.200.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp-v6 term 1 from protocol aggregate
set policy-options policy-statement 2bgp-v6 term 1 from route-filter 2001:0db8:0200::/48
exact
set policy-options policy-statement 2bgp-v6 term 1 then accept
```

vSRX:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-options aggregate route 10.150.0.0/16
set policy-options policy-statement 2bgp term 1 from route-filter 10.150.0.0/16 exact
set policy-options policy-statement 2bgp term 1 then accept
set policy-options policy-statement 2bgp term 2 then reject
set policy-options policy-statement frombgp term 1 from route-filter 192.168.0.0/16
exact
set policy-options policy-statement frombgp term 1 then accept
set policy-options policy-statement frombgp term 2 then reject
set protocols bgp group ebgp import frombgp
set protocols bgp group ebgp export 2bgp
set protocols bgp group ebgp neighbor 10.1.1.1 peer-as 1111
set protocols bgp group ebgp neighbor 10.1.1.1 local-as 9999
```

Task 7: VPN

MPLS Layer 3 VPN

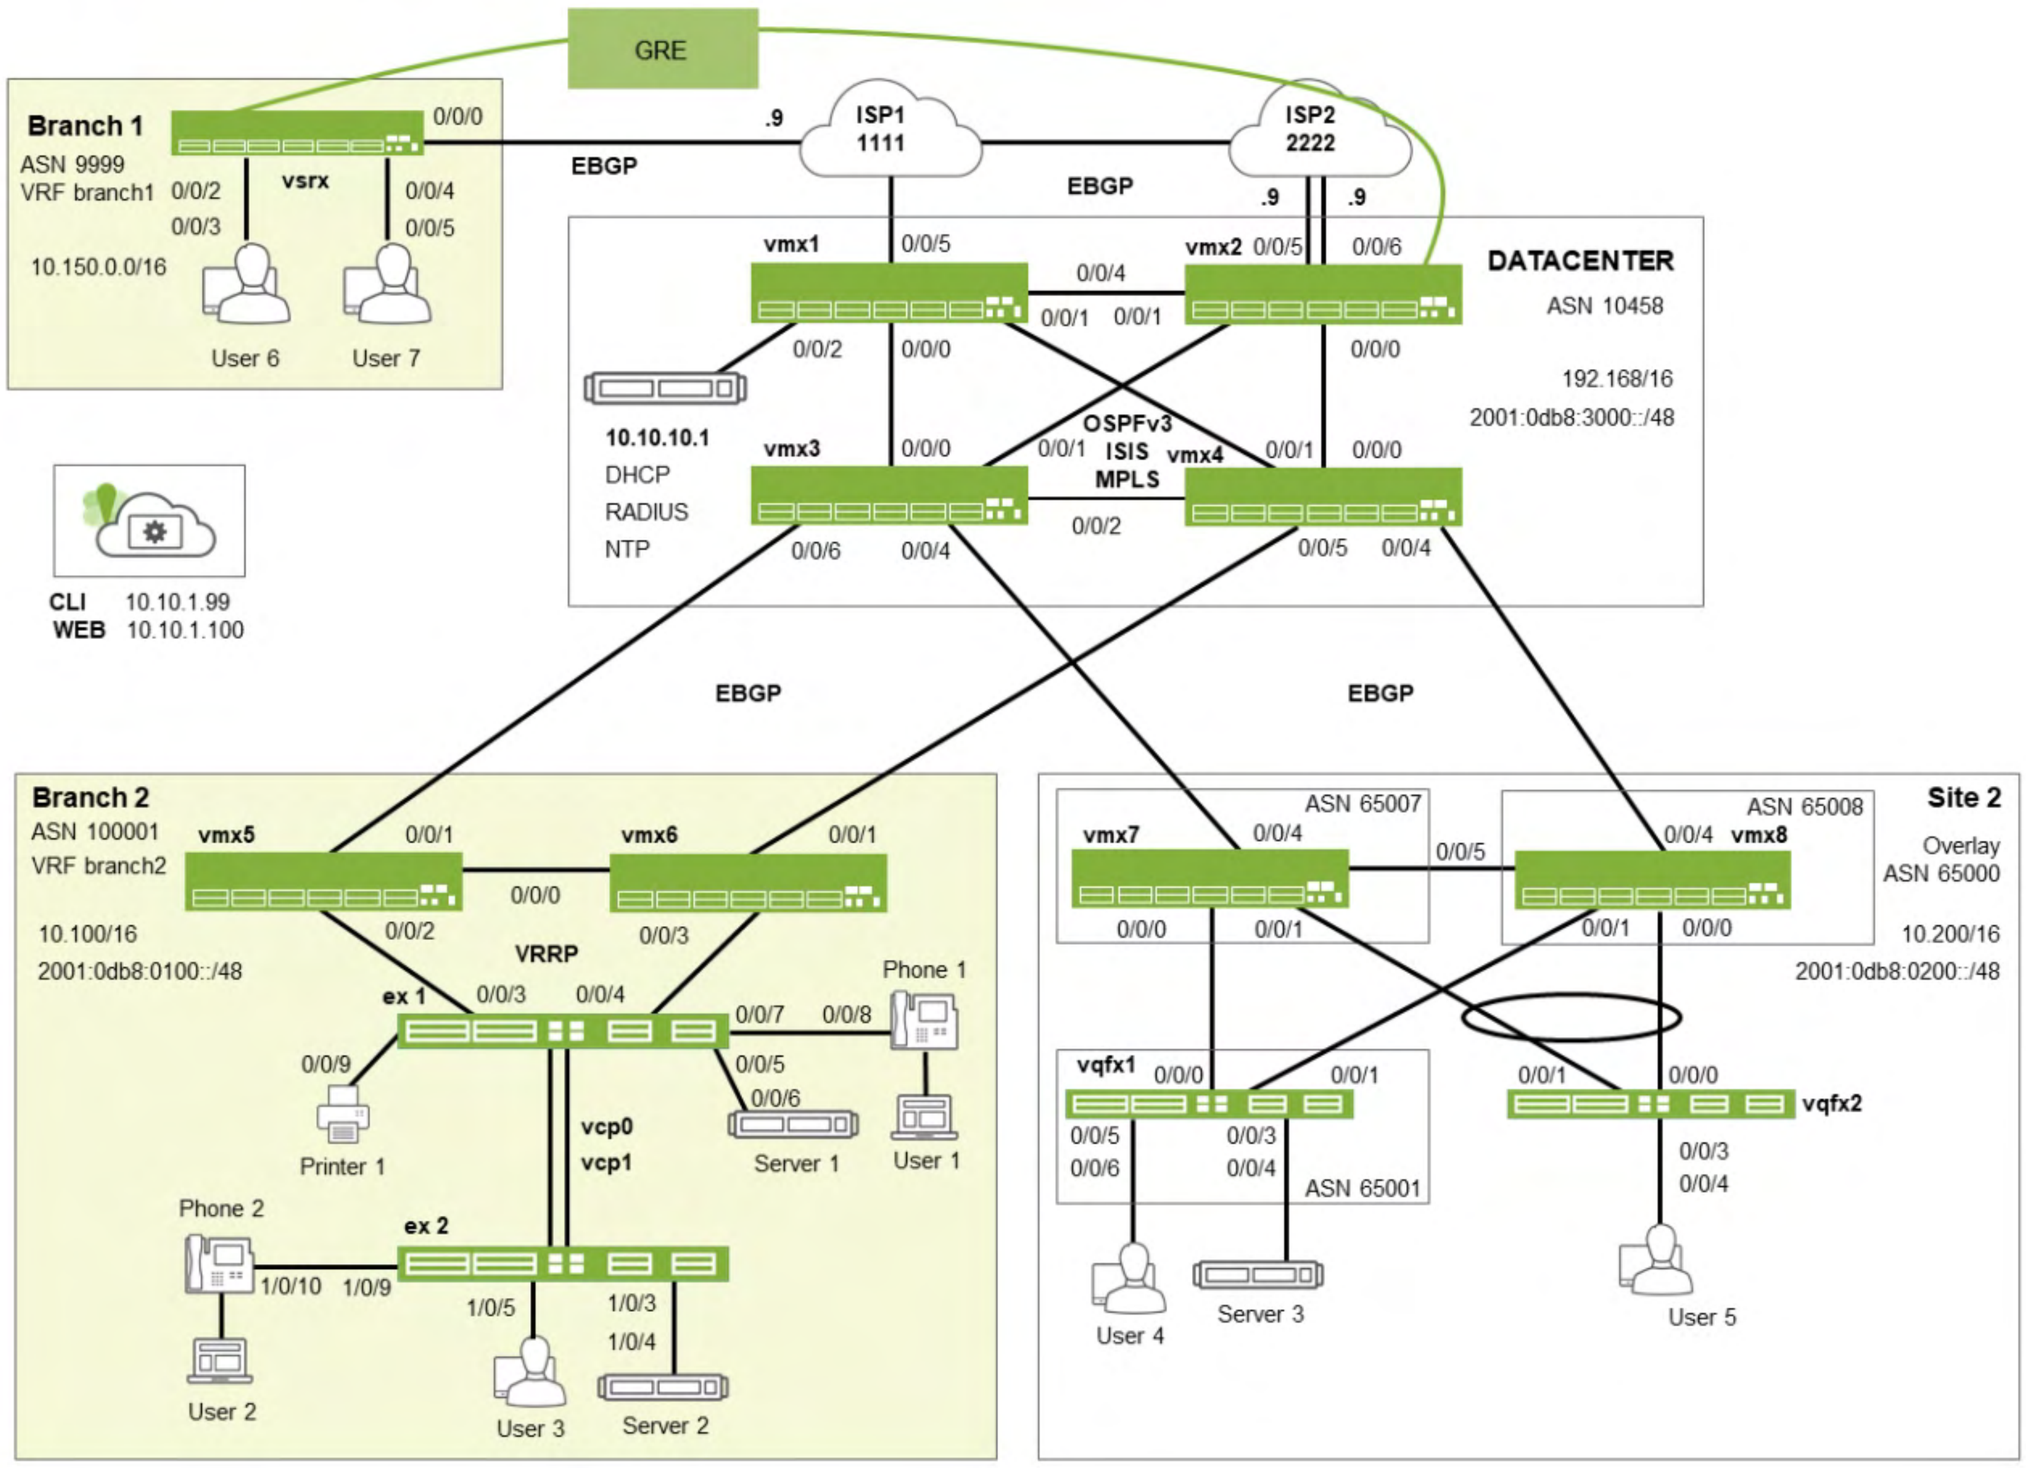

© 2020 Juniper Networks, Inc. All rights reserved.

L3VPN

In this task you will configure the Layer 3 MPLS VPN

1) Configure a MPLS L3VPN on vMX2, vMX3 and vMX4 connecting Branch1 and Branch 2 for IPv4 routing. Use the information as shown in the table below

Router	Interface	Routing instance	PE-CE routing
vMX2	GRE tunnel	branch1	OSPFv2
vMX3	ge-0/0/6.0	branch2	EBGP
vMX4	ge-0/0/5.0	branch2	EBGP

2) Configure vMX2 to ensure Branch1 can access Site 2 and the Internet for IPv4 traffic using inet.0 routing table.
3) Configure vMX3 and vMX4 to ensure Branch 2 can access Site 2 and the Internet for IPv4 and IPv6 traffic using the inet6.0 routing table.
4) Branch 1 should only learn the IPv4 and IPv6 default route from the data center

vMX2:

Note: You should have already configured vMX2 in a previous section. Just verify your configuration is correct.

```
lab@vmx2# show routing-instances
branch1 {
    instance-type vrf;
    interface gr-0/0/10.0;
    route-distinguisher 192.168.0.2:1;
    vrf-target target:10458:1;
    routing-options {
        interface-routes {
            rib-group inet branch1-to-inet;
        }
        static {
            route 0.0.0.0/0 next-table inet.0;
        }
    }
    protocols {
        ospf {
            rib-group branch1-to-inet;
            export 2ospfbranch1;
            area 0.0.0.0 {
                interface gr-0/0/10.0 {
                    hello-interval 1;
                }
            }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
}

lab@vmx2# show protocols isis
...
rib-group inet inet-to-branch1;

lab@vmx2# show routing-options
interface-routes {
    rib-group inet interfaces;
}
...
rib-groups {
    interfaces {
        import-rib [ inet.0 fbf.inet.0 branch1.inet.0 ];
    }
    branch1-to-inet {
        import-rib [ branch1.inet.0 inet.0 ];
    }
    inet-to-branch1 {
        import-rib [ inet.0 branch1.inet.0 ];
    }
}
...
```

vMX3:

```
lab@vmx3# show routing-options
rib-groups {
    branch2-to-inet-v4 {
        import-rib [ branch2.inet.0 inet.0 ];
    }
    branch2-to-inet-v6 {
        import-rib [ branch2.inet6.0 inet6.0 ];
    }
    inet-to-branch2-v4 {
        import-rib [ inet.0 branch2.inet.0 ];
        import-policy to-branch2-v4;
    }
    inet-to-branch2-v6 {
        import-rib [ inet6.0 branch2.inet6.0 ];
        import-policy to-branch2-v6;
    }
}

lab@vmx3# show policy-options
...
policy-statement to-branch2-v4 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    term 1 {
        from {
            route-filter 0.0.0.0/0 exact;
        }
        to rib branch2.inet.0;
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement to-branch2-v6 {
    term 1 {
        from {
            route-filter ::/0 exact;
        }
        to rib branch2.inet6.0;
        then accept;
    }
    term 2 {
        then reject;
    }
}

lab@vmx3# show routing-instances
branch2 {
    instance-type vrf;
    interface ge-0/0/6.0;
    route-distinguisher 192.168.0.3:1;
    vrf-target target:10458:1;
    protocols {
        bgp {
            group branch2 {
                family inet {
                    unicast {
                        rib-group branch2-to-inet-v4;
                    }
                }
                neighbor 192.168.35.5 {
                    peer-as 100001;
                }
            }
            group branch2-v6 {
                family inet6 {
                    unicast {
                        rib-group branch2-to-inet-v6;
                    }
                }
                neighbor 2001:0db8:3000:0035::5 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    family inet6 {
                        unicast;
                    }
                    peer-as 100001;
                }
            }
        }
    }
}

lab@vmx3# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.3;
    family inet {
        unicast {
            rib-group inet-to-branch2-v4;
        }
    }
    family inet-vpn {
        unicast;
    }
    export nhs;
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::3;
    family inet6 {
        unicast {
            rib-group inet-to-branch2-v6;
        }
    }
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
...
```

vMX4:

```
lab@vmx4# show routing-options
rib-groups {
    branch2-to-inet-v4 {
        import-rib [ branch2.inet.0 inet.0 ];
    }
    branch2-to-inet-v6 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        import-rib [ branch2.inet6.0 inet6.0 ];
    }
    inet-to-branch2-v4 {
        import-rib [ inet.0 branch2.inet.0 ];
        import-policy to-branch2-v4;
    }
    inet-to-branch2-v6 {
        import-rib [ inet6.0 branch2.inet6.0 ];
        import-policy to-branch2-v6;
    }
}

lab@vmx4# show policy-options
...
policy-statement to-branch2-v4 {
    term 1 {
        from {
            route-filter 0.0.0.0/0 exact;
        }
        to rib branch2.inet.0;
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement to-branch2-v6 {
    term 1 {
        from {
            route-filter ::/0 exact;
        }
        to rib branch2.inet6.0;
        then accept;
    }
    term 2 {
        then reject;
    }
}

lab@vmx4# show routing-instances
branch2 {
    instance-type vrf;
    interface ge-0/0/5.0;
    route-distinguisher 192.168.0.4:1;
    vrf-target target:10458:1;
    protocols {
        bgp {
            group branch2 {
                family inet {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            unicast {
                rib-group branch2-to-inet-v4;
            }
        }
        neighbor 192.168.46.6 {
            peer-as 100001;
        }
    }
    group branch2-v6 {
        family inet6 {
            unicast {
                rib-group branch2-to-inet-v6;
            }
            }
            neighbor 2001:0db8:3000:0046::6 {
                family inet6 {
                    unicast;
                }
                peer-as 100001;
            }
        }
    }
}

lab@vmx4# show protocols bgp
advertise-inactive;
group ibgp {
    type internal;
    local-address 192.168.0.4;
    family inet {
        unicast {
            rib-group inet-to-branch2-v4;
        }
    }
    family inet-vpn {
        unicast;
    }
    export nhs;
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
}
group ibgp-v6 {
    type internal;
    local-address 2001:0db8:3000:0000::4;
    family inet6 {
        unicast {
            rib-group inet-to-branch2-v6;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    export nhs;
    neighbor 2001:0db8:3000:0000::1;
    neighbor 2001:0db8:3000:0000::2;
}
...
```

Verification:

vMX2:

```
lab@vmx2# run show route receive-protocol bgp 192.168.0.3 table bgp.l3vpn.0

bgp.l3vpn.0: 8 destinations, 12 routes (8 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.0.3:1:10.100.0.0/16
*                         192.168.0.3                  100        100001 I
  192.168.0.3:1:192.168.35.0/24
*                         192.168.0.3                  100        I
```

vMX3:

```
lab@vmx3# run show route receive-protocol bgp 192.168.0.2 table bgp.l3vpn.0

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.0.2:1:10.150.99.0/24
*                         192.168.0.2          1       100        I
  192.168.0.2:1:173.17.99.0/24
*                         192.168.0.2                  100        I
```

vMX4:

```
lab@vmx4# run show route receive-protocol bgp 192.168.0.2 table bgp.l3vpn.0

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.0.2:1:10.150.99.0/24
*                         192.168.0.2          1       100        I
  192.168.0.2:1:173.17.99.0/24
*                         192.168.0.2                  100        I
```

EX1:

Verify from EX1 reachability to a destination inside the L3VPN. For example, 10.150.99.6 (user6) in Branch1.

```
lab@ex1# run ping 10.150.99.6 source 10.100.100.1 routing-instance user1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
PING 10.150.99.4 (10.150.99.6): 56 data bytes
64 bytes from 10.150.99.6: icmp_seq=0 ttl=59 time=3.708 ms
64 bytes from 10.150.99.6: icmp_seq=1 ttl=59 time=4.156 ms

lab@ex1# run traceroute 10.150.99.6 source 10.100.100.1 routing-instance user1
traceroute to 10.150.99.4 (10.150.99.4) from 10.100.100.1, 30 hops max, 40 byte packets
 1  10.100.100.254 (10.100.100.254)  2.511 ms  1.465 ms  4.064 ms
 2  10.100.156.5 (10.100.156.5)  1.711 ms  2.079 ms  2.267 ms
 3  192.168.35.3 (192.168.35.3)  3.103 ms  4.661 ms  2.218 ms
 4  192.168.23.2 (192.168.23.2)  4.721 ms  4.256 ms  3.854 ms
     MPLS Label=299888 CoS=0 TTL=1 S=1
 5  173.17.99.99 (173.17.99.99)  3.813 ms  3.662 ms  3.295 ms
 6  10.150.99.6 (10.150.99.6)  4.319 ms  4.642 ms  3.613 ms
```

Verify from EX1 reachability to a destination outside of the L3VPN. For example, 192.168.0.2 (vMX2 loopback) in the Datacenter.

```
lab@ex1# run ping 192.168.0.2 source 10.100.100.1 routing-instance user1
PING 192.168.0.2 (192.168.0.2): 56 data bytes
64 bytes from 192.168.0.2: icmp_seq=0 ttl=61 time=4.495 ms
64 bytes from 192.168.0.2: icmp_seq=1 ttl=61 time=4.122 ms
```

Note: You will not be able to verify reachability to site2 until after you have configured the overlay and underlay networks. The next command was gathered after that was completed.

Verify from EX1 reachability to a destination outside of the L3VPN. For example, 10.200.101.3 (server 3) in Site 2.

```
lab@ex1# run ping 10.200.101.3 source 10.100.100.1 routing-instance user1
PING 10.200.101.3 (10.200.101.3): 56 data bytes
64 bytes from 10.200.101.3: icmp_seq=0 ttl=60 time=337.247 ms
64 bytes from 10.200.101.3: icmp_seq=1 ttl=60 time=304.056 ms
```

Configurations for Task Completion:

vMX3:

```
set routing-options rib-groups branch2-to-inet-v4 import-rib [ branch2.inet.0 inet.0 ]
set routing-options rib-groups branch2-to-inet-v6 import-rib [ branch2.inet6.0 inet6.0 ]
set routing-options rib-groups inet-to-branch2-v4 import-rib [ inet.0 branch2.inet.0 ]
set routing-options rib-groups inet-to-branch2-v4 import-policy to-branch2-v4
set routing-options rib-groups inet-to-branch2-v6 import-rib [ inet6.0 branch2.inet6.0 ]
set routing-options rib-groups inet-to-branch2-v6 import-policy to-branch2-v6
set policy-options policy-statement to-branch2-v4 term 1 from route-filter 0.0.0.0/0
exact
set policy-options policy-statement to-branch2-v4 term 1 to rib branch2.inet.0
set policy-options policy-statement to-branch2-v4 term 1 then accept
set policy-options policy-statement to-branch2-v4 term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set policy-options policy-statement to-branch2-v6 term 1 from route-filter ::/0 exact
set policy-options policy-statement to-branch2-v6 term 1 to rib branch2.inet6.0
set policy-options policy-statement to-branch2-v6 term 1 then accept
set policy-options policy-statement to-branch2-v6 term 2 then reject
set routing-instances branch2 instance-type vrf
set routing-instances branch2 interface ge-0/0/6.0
set routing-instances branch2 route-distinguisher 192.168.0.3:1
set routing-instances branch2 vrf-target target:10458:1
set routing-instances branch2 protocols bgp group branch2 family inet unicast rib-group
branch2-to-inet-v4
set routing-instances branch2 protocols bgp group branch2 neighbor 192.168.35.5 peer-as
100001
set routing-instances branch2 protocols bgp group branch2-v6 family inet6 unicast rib-
group branch2-to-inet-v6
set routing-instances branch2 protocols bgp group branch2-v6 neighbor
2001:0db8:3000:0035::5 family inet6 unicast
set routing-instances branch2 protocols bgp group branch2-v6 neighbor
2001:0db8:3000:0035::5 peer-as 100001
set protocols bgp group ibgp family inet unicast rib-group inet-to-branch2-v4
set protocols bgp group ibgp-v6 family inet6 unicast rib-group inet-to-branch2-v6
```

vMX4:

```
set routing-options rib-groups branch2-to-inet-v4 import-rib [ branch2.inet.0 inet.0 ]
set routing-options rib-groups branch2-to-inet-v6 import-rib [ branch2.inet6.0 inet6.0 ]
set routing-options rib-groups inet-to-branch2-v4 import-rib [ inet.0 branch2.inet.0 ]
set routing-options rib-groups inet-to-branch2-v4 import-policy to-branch2-v4
set routing-options rib-groups inet-to-branch2-v6 import-rib [ inet6.0 branch2.inet6.0 ]
set routing-options rib-groups inet-to-branch2-v6 import-policy to-branch2-v6
set policy-options policy-statement to-branch2-v4 term 1 from route-filter 0.0.0.0/0
exact
set policy-options policy-statement to-branch2-v4 term 1 to rib branch2.inet.0
set policy-options policy-statement to-branch2-v4 term 1 then accept
set policy-options policy-statement to-branch2-v4 term 2 then reject
set policy-options policy-statement to-branch2-v6 term 1 from route-filter ::/0 exact
set policy-options policy-statement to-branch2-v6 term 1 to rib branch2.inet6.0
set policy-options policy-statement to-branch2-v6 term 1 then accept
set policy-options policy-statement to-branch2-v6 term 2 then reject
set routing-instances branch2 instance-type vrf
set routing-instances branch2 interface ge-0/0/5.0
set routing-instances branch2 route-distinguisher 192.168.0.4:1
set routing-instances branch2 vrf-target target:10458:1
set routing-instances branch2 protocols bgp group branch2 family inet unicast rib-group
branch2-to-inet-v4
set routing-instances branch2 protocols bgp group branch2 neighbor 192.168.46.6 peer-as
100001
set routing-instances branch2 protocols bgp group branch2-v6 family inet6 unicast rib-
group branch2-to-inet-v6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-instances branch2 protocols bgp group branch2-v6 neighbor
2001:0db8:3000:0046::6 family inet6 unicast
set routing-instances branch2 protocols bgp group branch2-v6 neighbor
2001:0db8:3000:0046::6 peer-as 100001
set protocols bgp group ibgp family inet unicast rib-group inet-to-branch2-v4
set protocols bgp group ibgp-v6 family inet6 unicast rib-group inet-to-branch2-v6
```

Task 8: EVPN/VXLAN

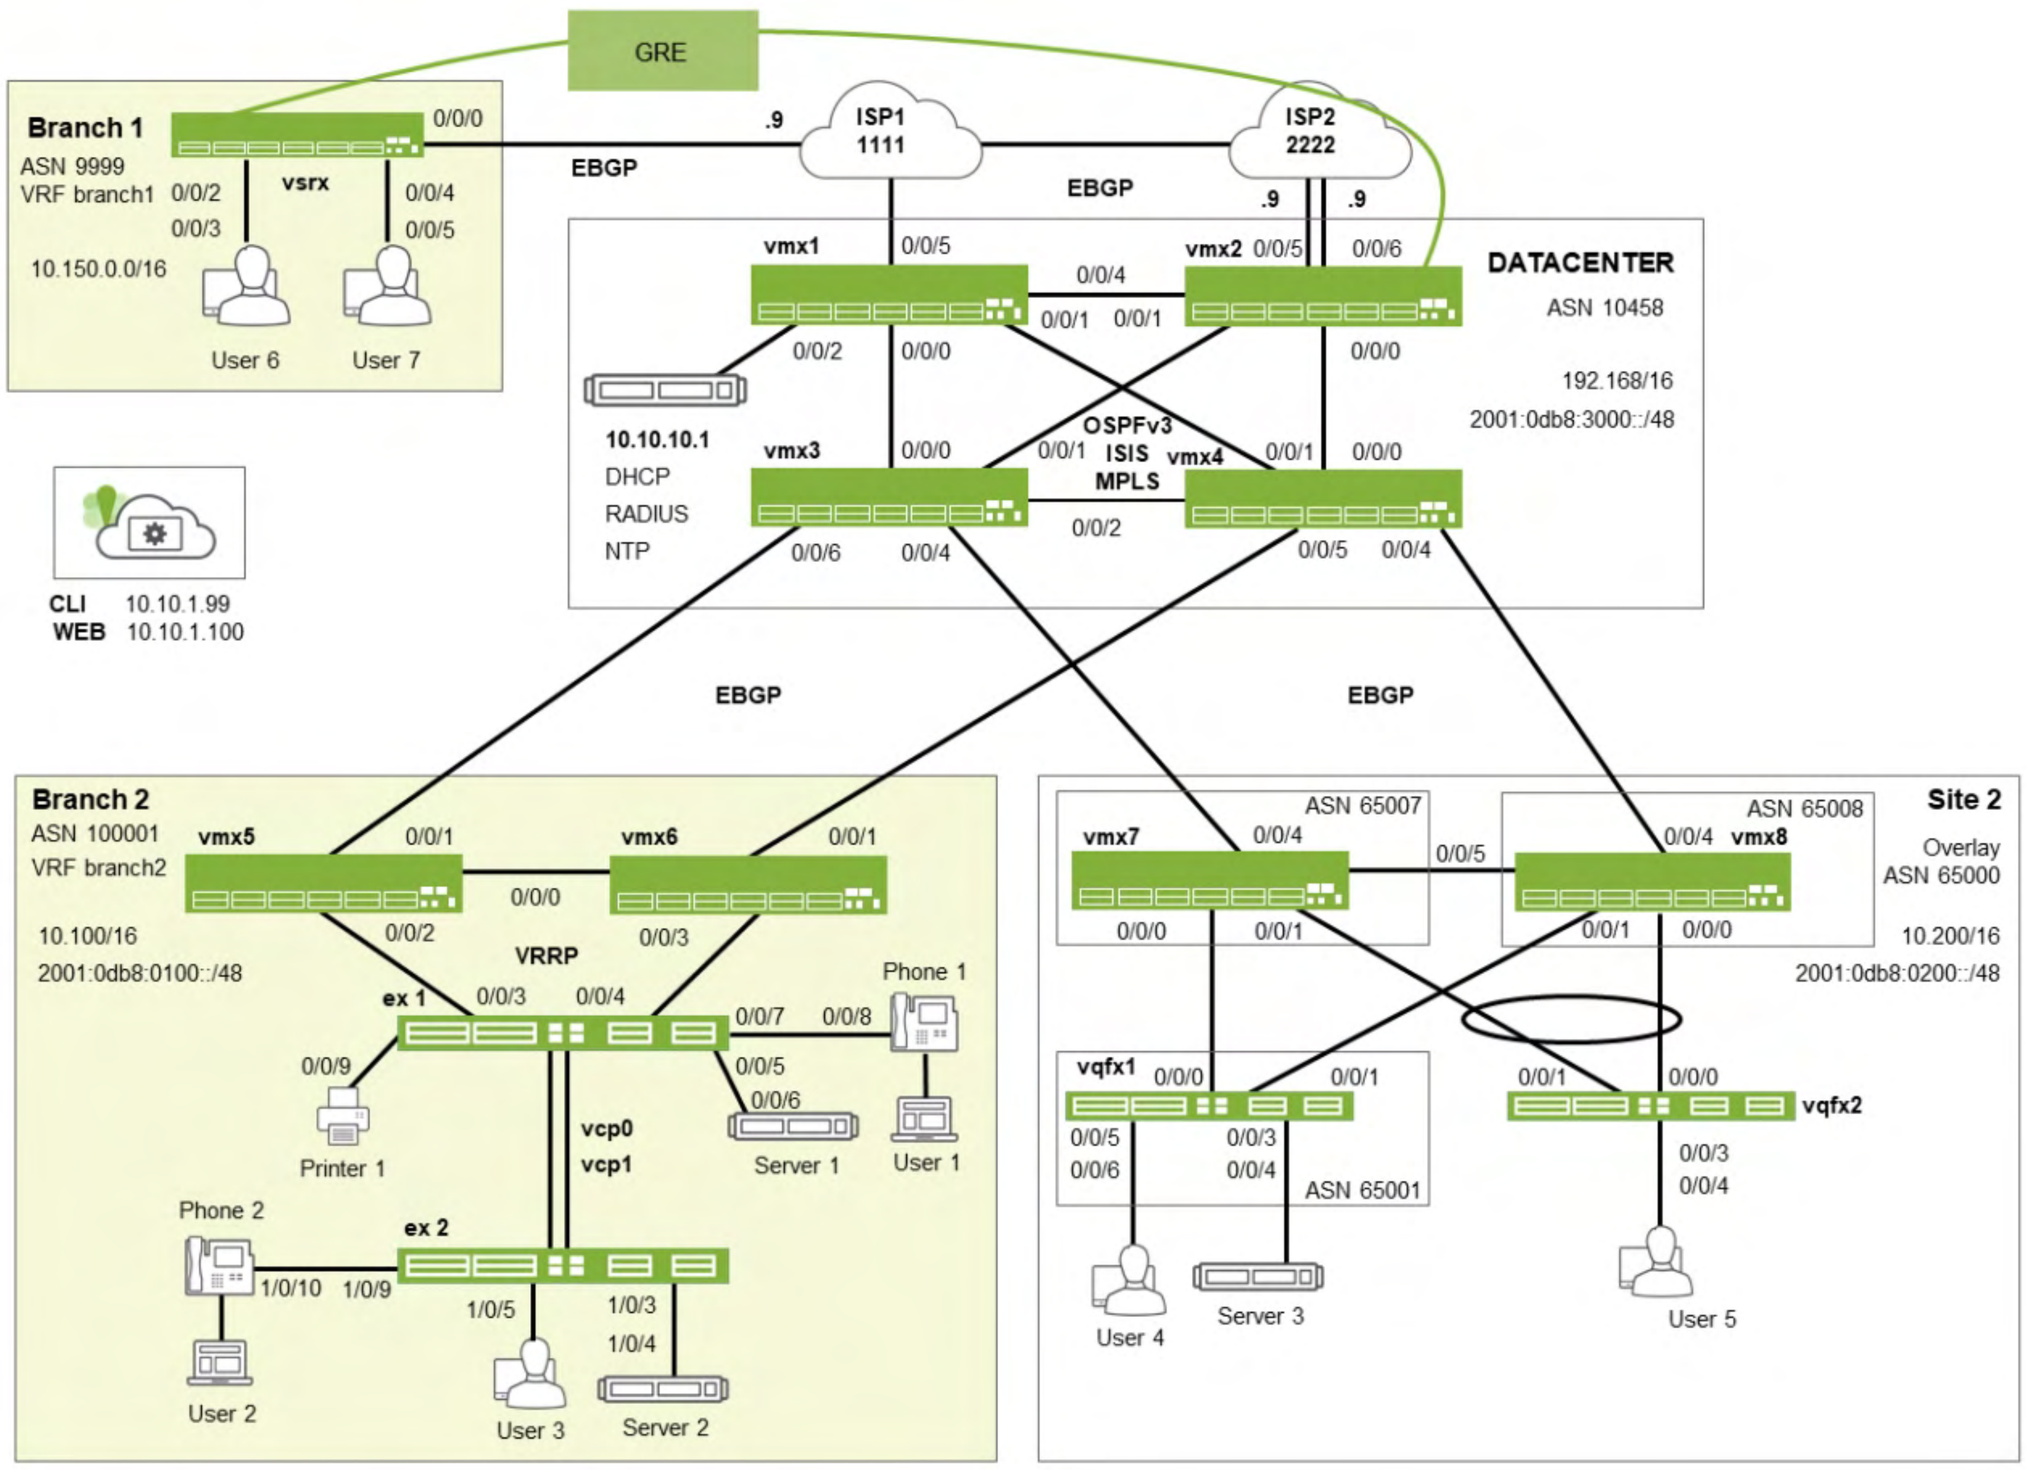

EVPN VXLAN

In this section you will configure the layer 2 connectivity in Site 2 using the EVPN/VXLAN technology.

1) Configure the EBGP underlay in Site 2 to exchange device loopback reachability. Use the AS numbers as shown in the topology. Use the interface addresses to establish the EBGP peerings.

© 2020 Juniper Networks, Inc. All rights reserved.

2) Configure the IBGP EVPN overlay for Site 2. vMX7 and vMX8 are router reflectors for the EVPN overlay.
3) Ensure that virtual gateway addresses respond to ICMP packets.
4) Configure one EVPN VLAN aware service on vMX7 and vMX8 for VLAN 100 and 101
5) Ensure the devices in Site 2 can reach each other, the devices in Branch2 and the internet

vMX7:

```
lab@vmx7# show protocols bgp group underlay
export export-loopback;
local-as 65007;
multipath {
    multiple-as;
}
neighbor 10.200.17.1 {
    peer-as 65001;
}

lab@vmx7# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

vMX8:

```
lab@vmx8# show protocols bgp group underlay
export export-loopback;
local-as 65008;
multipath {
    multiple-as;
}
neighbor 10.200.18.1 {
    peer-as 65001;
}

lab@vmx8# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

vQFX1:

```
lab@vQFX1# show protocols bgp group underlay
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
export export-loopback;
local-as 65001;
multipath {
    multiple-as;
}
neighbor 10.200.17.7 {
    peer-as 65007;
}
neighbor 10.200.18.8 {
    peer-as 65008;
}

lab@vQFX1# show policy-options policy-statement export-loopback
term 1 {
    from interface lo0.0;
    then accept;
}
```

Verification:

vMX7:

```
lab@vmx7# run show bgp summary group underlay
Groups: 3 Peers: 3 Down peers: 0
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.17.1           65001          9         12       0       1        2:33 2/2/2/0
0/0/0/0
```

vMX8:

```
lab@vmx8# run show bgp summary group underlay
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.18.1           65001         16         15       0       1        3:43 2/7/7/0
0/0/0/0
```

vQFX1:

```
lab@vQFX1# run show bgp summary group underlay
Threading mode: BGP I/O
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.17.7           65007         16         11       0       1        3:50 Establ
  inet.0: 6/6/6/0
10.200.18.8           65008         16         15       0       1        3:46 Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 6/6/6/0
```

6) Configure the IBGP EVPN overlay for Site 2. vMX7 and vMX8 are router reflectors for the EVPN overlay. Leaf devices should only peer with Spine devices.

vMX7:

```
lab@vmx7# show protocols bgp group overlay
type internal;
local-address 10.200.0.7;
family evpn {
    signaling;
}
cluster 10.200.0.7;
local-as 65000;
multipath;
neighbor 10.200.0.1;
neighbor 10.200.0.8;
```

vMX8:

```
lab@vmx8# show protocols bgp group overlay
type internal;
local-address 10.200.0.8;
family evpn {
    signaling;
}
cluster 10.200.0.8;
local-as 65000;
multipath;
neighbor 10.200.0.1;
neighbor 10.200.0.7;
```

vQFX1:

```
lab@vQFX1# show protocols bgp group overlay
type internal;
local-address 10.200.0.1;
family evpn {
    signaling;
}
local-as 65000;
multipath;
neighbor 10.200.0.7;
neighbor 10.200.0.8;
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX7:

```
lab@vmx7# run show bgp summary group overlay
Groups: 4 Peers: 7 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        17          9          0          0          0          0
bgp.evpn.0
                        50         35          0          0          0          0
inet6.0
                         3          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1            65000      13100      13050       0       3  4d 3:07:05 Establ
  bgp.evpn.0: 10/10/10/0
  evpn1.evpn.0: 10/10/10/0
  __default_evpn__.evpn.0: 0/0/0/0
  __default_evpn__.evpn.0: 0/0/0/0
10.200.0.8            65000      13055      13052       0       0  4d 3:24:18 Establ
  bgp.evpn.0: 20/35/35/0
  evpn1.evpn.0: 20/35/35/0
  __default_evpn__.evpn.0: 0/0/0/0

lab@vmx7# run show bgp summary group overlay
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1            65000         20         15       0       0        6:11 Establ
  bgp.evpn.0: 8/8/8/0
10.200.0.8            65000         17         17       0       0        5:31 Establ
  bgp.evpn.0: 0/8/8/0
```

vMX8:

```
lab@vmx8# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:64:00  Apr 08 01:42:34
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:64:00  Apr 08 01:42:34
2001:db8:200:100::254
     100        00:50:56:a9:57:23  10.200.0.1                     Apr 14 10:16:00
     100        02:05:86:71:77:13  00:78:78:78:78:78:78:78:78:02  Apr 18 12:33:57
10.200.100.5

2001:db8:200:100::5

fe80::205:86ff:fe71:7713
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
     100        02:05:86:71:d9:1b  10.200.0.1                    Apr 08 02:09:02
10.200.100.4

2001:db8:200:100::4

fe80::205:86ff:fe71:d91b
     100        2c:6b:f5:2c:a7:f0  irb.100                       Apr 08 01:42:34
10.200.100.251

2001:db8:200:100::251

fe80::2e6b:f500:642c:a7f0
     100        2c:6b:f5:9c:a0:f0  10.200.0.7                    Apr 08 01:42:33
10.200.100.250

2001:db8:200:100::250

fe80::2e6b:f500:649c:a0f0
     101        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:65:00  Apr 08 01:42:34
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:65:00  Apr 08 01:42:34
2001:db8:200:101::254
     101        00:50:56:a9:09:a3  10.200.0.1                    Apr 14 10:16:00
     101        02:05:86:71:d9:13  10.200.0.1                    Apr 08 02:09:13
10.200.101.3

2001:db8:200:101::3

fe80::205:86ff:fe71:d913
     101        2c:6b:f5:2c:a7:f0  irb.101                       Apr 08 01:42:34
10.200.101.251

2001:db8:200:101::251

fe80::2e6b:f500:652c:a7f0
     101        2c:6b:f5:9c:a0:f0  10.200.0.7                    Apr 08 01:42:33
10.200.101.250

2001:db8:200:101::250

fe80::2e6b:f500:659c:a0f0

lab@vmx8# run show bgp summary group overlay
Groups: 4 Peers: 7 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                       3          3          0          0          0          0
inet.0
                      13          9          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
bgp.evpn.0
                     50         35          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1           65000      13100      13056       0       3  4d 3:07:22 Establ
  bgp.evpn.0: 10/10/10/0
  evpn1.evpn.0: 10/10/10/0
  __default_evpn__.evpn.0: 0/0/0/0
10.200.0.7           65000      13053      13054       0       0  4d 3:24:39 Establ
  bgp.evpn.0: 20/35/35/0
  evpn1.evpn.0: 20/35/35/0
  __default_evpn__.evpn.0: 0/0/0/0
```

```
lab@vmx8# run show bgp summary group overlay
...
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.1           65000         17         12       0       0        5:03 Establ
  bgp.evpn.0: 8/8/8/0
10.200.0.7           65000         18         16       0       0        5:07 Establ
  bgp.evpn.0: 0/8/8/0
```

vQFX1:

```
lab@vQFX1# run show bgp summary group overlay
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.evpn.0
                     90         45          0          0          0          0
inet.0
                     14         14          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.200.0.7           65000      13051      13099       0       0  4d 3:07:41 Establ
  __default_evpn__.evpn.0: 0/0/0/0
  bgp.evpn.0: 25/45/45/0
  default-switch.evpn.0: 25/45/45/0
10.200.0.8           65000      13057      13099       0       0  4d 3:07:37 Establ
  __default_evpn__.evpn.0: 0/0/0/0
  bgp.evpn.0: 20/45/45/0
  default-switch.evpn.0: 20/45/45/0
```

7) Ensure that virtual gateway addresses respond to ICMP packets.
8) Configure one EVPN VLAN aware service on vMX7 and vMX8 for VLAN 100 and 101

vMX7:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vmx7# show interfaces ae15
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:02;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:78:02;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}

lab@vmx7# show policy-options
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}
community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;

lab@vmx7# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
    interface ae15.0;
    route-distinguisher 10.200.0.7:1;
    vrf-import fabric;
    vrf-target target:65000:1000;
    protocols {
        evpn {
            encapsulation vxlan;
            extended-vni-list [ 100 101 ];
        }
    }
    bridge-domains {
        vlan100 {
            domain-type bridge;
            vlan-id 100;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        routing-interface irb.100;
        vxlan {
            vni 100;
            ingress-node-replication;
        }
    }
    vlan101 {
        domain-type bridge;
        vlan-id 101;
        routing-interface irb.101;
        vxlan {
            vni 101;
            ingress-node-replication;
        }
    }
}

lab@vmx7# show interfaces irb
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.250/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::250/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.250/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::250/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

vMX8:

```
lab@vmx8# show interfaces ae15
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
aggregated-ether-options {
    lacp {
        active;
        system-id 00:78:78:78:78:02;
    }
}
unit 0 {
    esi {
        00:78:78:78:78:78:78:78:78:02;
        all-active;
    }
    family bridge {
        interface-mode trunk;
        vlan-id-list [ 100 101 ];
    }
}
```

```
lab@vmx8# show policy-options
...
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}

community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;
```

```
lab@vmx8# show routing-instances
evpn1 {
    vtep-source-interface lo0.0;
    instance-type virtual-switch;
    interface ae15.0;
    route-distinguisher 10.200.0.8:1;
    vrf-import fabric;
    vrf-target target:65000:1000;
    protocols {
        evpn {
            encapsulation vxlan;
            extended-vni-list [ 100 101 ];
        }
    }
    bridge-domains {
        vlan100 {
            domain-type bridge;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        vlan-id 100;
        routing-interface irb.100;
        vxlan {
            vni 100;
            ingress-node-replication;
        }
    }
    vlan101 {
        domain-type bridge;
        vlan-id 101;
        routing-interface irb.101;
        vxlan {
            vni 101;
            ingress-node-replication;
        }
    }
}

lab@vmx8# show interfaces irb
unit 100 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.100.251/24 {
            virtual-gateway-address 10.200.100.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:100::251/64 {
            virtual-gateway-address 2001:0db8:0200:100::254;
        }
    }
}
unit 101 {
    virtual-gateway-accept-data;
    family inet {
        address 10.200.101.251/24 {
            virtual-gateway-address 10.200.101.254;
        }
    }
    family inet6 {
        address 2001:0db8:0200:101::251/64 {
            virtual-gateway-address 2001:0db8:0200:101::254;
        }
    }
}
```

vQFX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@vQFX1# show switch-options
vtep-source-interface lo0.0;
route-distinguisher 10.200.0.1:1;
vrf-import fabric;
vrf-target target:65000:1;

lab@vQFX1# show protocols evpn
vni-options {
    vni 100 {
        vrf-target target:65000:100;
    }
    vni 101 {
        vrf-target target:65000:101;
    }
}
encapsulation vxlan;
multicast-mode ingress-replication;
extended-vni-list [ 100 101 ];

lab@vQFX1# show policy-options
policy-statement fabric {
    term 1 {
        from community [ esi evpn1 vni100 vni101 ];
        then accept;
    }
}
community esi members target:65000:1;
community evpn1 members target:65000:1000;
community vni100 members target:65000:100;
community vni101 members target:65000:101;

lab@vQFX1# show vlans
default {
    vlan-id 1;
}
server {
    vlan-id 101;
    vxlan {
        vni 101;
    }
}
user {
    vlan-id 100;
    vxlan {
        vni 100;
    }
}
```

Verification:

© 2020 Juniper Networks, Inc. All rights reserved.

vMX7:

```
lab@vmx7# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:64:00  May 06 07:24:33
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:64:00  May 06 07:24:33
2001:db8:200:100::254
     100        02:05:86:71:76:1b  10.200.0.1                     May 06 07:36:34
10.200.100.4
fe80::205:86ff:fe71:761b
     100        02:05:86:71:9b:13  00:78:78:78:78:78:78:78:78:02  May 06 07:36:07
10.200.100.5

fe80::205:86ff:fe71:9b13
     100        2c:6b:f5:2c:a7:f0  10.200.0.8                     May 06 07:24:33
10.200.100.251

2001:db8:200:100::251

fe80::2e6b:f500:642c:a7f0
     100        2c:6b:f5:9c:a0:f0  irb.100                        May 06 07:24:33
10.200.100.250

2001:db8:200:100::250

fe80::2e6b:f500:649c:a0f0
     101        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:65:00  May 06 07:24:33
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:65:00  May 06 07:24:33
2001:db8:200:101::254
     101        02:05:86:71:76:13  10.200.0.1                     May 06 07:36:32
10.200.101.3

fe80::205:86ff:fe71:7613
     101        2c:6b:f5:2c:a7:f0  10.200.0.8                     May 06 07:24:33
10.200.101.251

2001:db8:200:101::251

fe80::2e6b:f500:652c:a7f0
     101        2c:6b:f5:9c:a0:f0  irb.101                        May 06 07:24:33
10.200.101.250

2001:db8:200:101::250
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fe80::2e6b:f500:659c:a0f0
```

vMX8:

```
lab@vmx8# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:64:00  May 06 07:24:33
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:64:00  May 06 07:24:33
2001:db8:200:100::254
     100        02:05:86:71:76:1b  10.200.0.1                     May 06 07:36:34
10.200.100.4
     100        02:05:86:71:9b:13  00:78:78:78:78:78:78:78:78:02  May 06 07:35:30
10.200.100.5

fe80::205:86ff:fe71:9b13
     100        2c:6b:f5:2c:a7:f0  irb.100                        May 06 07:24:33
10.200.100.251

2001:db8:200:100::251

fe80::2e6b:f500:642c:a7f0
     100        2c:6b:f5:9c:a0:f0  10.200.0.7                     May 06 07:24:33
10.200.100.250

2001:db8:200:100::250

fe80::2e6b:f500:649c:a0f0
     101        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:65:00  May 06 07:24:33
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:65:00  May 06 07:24:33
2001:db8:200:101::254
     101        02:05:86:71:76:13  10.200.0.1                     May 06 07:36:32
     101        2c:6b:f5:2c:a7:f0  irb.101                        May 06 07:24:33
10.200.101.251

2001:db8:200:101::251

fe80::2e6b:f500:652c:a7f0
     101        2c:6b:f5:9c:a0:f0  10.200.0.7                     May 06 07:24:33
10.200.101.250

2001:db8:200:101::250

fe80::2e6b:f500:659c:a0f0
```


© 2020 Juniper Networks, Inc. All rights reserved.

vQFX1:

```
lab@vQFX1# run show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
     100        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:64:00  May 06 07:33:41
10.200.100.254
     100        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:64:00  May 06 07:33:41
2001:db8:200:100::254
     100        02:05:86:71:76:1b  xe-0/0/5.0                     May 06 07:36:34
10.200.100.4

fe80::205:86ff:fe71:761b
     100        02:05:86:71:9b:13  00:78:78:78:78:78:78:78:78:02  May 06 07:36:07
10.200.100.5

fe80::205:86ff:fe71:9b13
     100        2c:6b:f5:2c:a7:f0  10.200.0.8                     May 06 07:33:41
10.200.100.251

2001:db8:200:100::251

fe80::2e6b:f500:642c:a7f0
     100        2c:6b:f5:9c:a0:f0  10.200.0.7                     May 06 07:33:41
10.200.100.250

2001:db8:200:100::250

fe80::2e6b:f500:649c:a0f0
     101        00:00:5e:00:01:01  05:00:00:00:00:00:00:00:65:00  May 06 07:33:41
10.200.101.254
     101        00:00:5e:00:02:01  05:00:00:00:00:00:00:00:65:00  May 06 07:33:41
2001:db8:200:101::254
     101        02:05:86:71:76:13  xe-0/0/3.0                     May 06 07:36:31
10.200.101.3

fe80::205:86ff:fe71:7613
     101        2c:6b:f5:2c:a7:f0  10.200.0.8                     May 06 07:33:41
10.200.101.251

2001:db8:200:101::251

fe80::2e6b:f500:652c:a7f0
     101        2c:6b:f5:9c:a0:f0  10.200.0.7                     May 06 07:33:41
10.200.101.250

2001:db8:200:101::250
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fe80::2e6b:f500:659c:a0f0
```

9) Ensure the devices in Site 2 can reach each other, the device in Site1 and the Internet

vQFX1:

```
lab@vQFX1# run ping 10.200.100.5 routing-instance user4
PING 10.200.100.5 (10.200.100.5): 56 data bytes
64 bytes from 10.200.100.5: icmp_seq=0 ttl=64 time=578.779 ms
64 bytes from 10.200.100.5: icmp_seq=1 ttl=64 time=502.334 ms
64 bytes from 10.200.100.5: icmp_seq=2 ttl=64 time=504.193 ms
^C
--- 10.200.100.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 502.334/528.435/578.779/35.606 ms

lab@vQFX1# run ping 10.200.101.3 routing-instance user4
PING 10.200.101.3 (10.200.101.3): 56 data bytes
64 bytes from 10.200.101.3: icmp_seq=0 ttl=63 time=700.634 ms
64 bytes from 10.200.101.3: icmp_seq=1 ttl=63 time=700.138 ms
^C
--- 10.200.101.3 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 700.138/700.386/700.634/0.248 ms
```

vQFX2:

```
lab@vQFX2# run ping 10.100.100.2 routing-instance user5
PING 10.100.100.2 (10.100.100.2): 56 data bytes
64 bytes from 10.100.100.2: icmp_seq=0 ttl=60 time=300.049 ms
64 bytes from 10.100.100.2: icmp_seq=1 ttl=60 time=300.238 ms
^C
--- 10.100.100.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 300.049/300.144/300.238/0.094 ms

lab@vQFX2# run ping 10.100.100.1 routing-instance user5
PING 10.100.100.1 (10.100.100.1): 56 data bytes
64 bytes from 10.100.100.1: icmp_seq=0 ttl=60 time=300.708 ms
^C
--- 10.100.100.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 300.708/300.708/300.708/0.000 ms
```

EX1:

© 2020 Juniper Networks, Inc. All rights reserved.

Note: If you have not previously verified DHCP is working you may need to bounce the ge-1/0/4 and ge-0/0/6 interfaces by disabling and enabling them.

```
lab@EX1# run show dhcp client binding routing-instance server1
IP address        Hardware address   Expires     State      Interface
10.100.101.50     cc:e1:94:d2:6e:49  26005       BOUND      ge-0/0/6.0

{master:0}[edit]
lab@EX1# run show dhcp client binding routing-instance server2
IP address        Hardware address   Expires     State      Interface
10.100.101.51     b8:c2:53:97:97:a7  26000       BOUND      ge-1/0/4.0
```

vQFX2:

```
lab@vQFX2# run ping 10.100.101.50 routing-instance user5
PING 10.100.101.50 (10.100.101.50): 56 data bytes
64 bytes from 10.100.101.50: icmp_seq=0 ttl=60 time=380.841 ms
^C
--- 10.100.101.50 ping statistics ---
2 packets transmitted, 1 packets received, 50% packet loss
round-trip min/avg/max/stddev = 380.841/380.841/380.841/0.000 ms

lab@vQFX2# run ping 10.100.101.51 routing-instance user5
PING 10.100.101.51 (10.100.101.51): 56 data bytes
64 bytes from 10.100.101.51: icmp_seq=0 ttl=60 time=301.208 ms
^C
--- 10.100.101.51 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 301.208/301.208/301.208/0.000 ms
```

Configurations for Task Completion:

vMX7:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65007
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.17.1 peer-as 65001
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.7
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 10.200.0.7
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.1
set protocols bgp group overlay neighbor 10.200.0.8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ae15 aggregated-ether-options lacp system-id 00:78:78:78:78:02
set interfaces ae15 unit 0 esi 00:78:78:78:78:78:78:78:02
set interfaces ae15 unit 0 esi all-active
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface ae15.0
set routing-instances evpn1 route-distinguisher 10.200.0.7:1
set routing-instances evpn1 vrf-import fabric
set routing-instances evpn1 vrf-target target:65000:1000
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list 100
set routing-instances evpn1 protocols evpn extended-vni-list 101
set routing-instances evpn1 bridge-domains vlan100 domain-type bridge
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan100 vxlan ingress-node-replication
set routing-instances evpn1 bridge-domains vlan101 domain-type bridge
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set routing-instances evpn1 bridge-domains vlan101 vxlan ingress-node-replication
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 100 family inet address 10.200.100.250/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::250/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 virtual-gateway-accept-data
set interfaces irb unit 101 family inet address 10.200.101.250/24 virtual-gateway-
address 10.200.101.254
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::250/64 virtual-
gateway-address 2001:0db8:0200:101::254
```

vMX8:

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65008
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.18.1 peer-as 65001
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group overlay local-address 10.200.0.8
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 10.200.0.8
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.1
set protocols bgp group overlay neighbor 10.200.0.7
set interfaces ae15 aggregated-ether-options lacp system-id 00:78:78:78:78:02
set interfaces ae15 unit 0 esi 00:78:78:78:78:78:78:78:78:02
set interfaces ae15 unit 0 esi all-active
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set routing-instances evpn1 vtep-source-interface lo0.0
set routing-instances evpn1 instance-type virtual-switch
set routing-instances evpn1 interface ae15.0
set routing-instances evpn1 route-distinguisher 10.200.0.8:1
set routing-instances evpn1 vrf-import fabric
set routing-instances evpn1 vrf-target target:65000:1000
set routing-instances evpn1 protocols evpn encapsulation vxlan
set routing-instances evpn1 protocols evpn extended-vni-list 100
set routing-instances evpn1 protocols evpn extended-vni-list 101
set routing-instances evpn1 bridge-domains vlan100 domain-type bridge
set routing-instances evpn1 bridge-domains vlan100 vlan-id 100
set routing-instances evpn1 bridge-domains vlan100 routing-interface irb.100
set routing-instances evpn1 bridge-domains vlan100 vxlan vni 100
set routing-instances evpn1 bridge-domains vlan100 vxlan ingress-node-replication
set routing-instances evpn1 bridge-domains vlan101 domain-type bridge
set routing-instances evpn1 bridge-domains vlan101 vlan-id 101
set routing-instances evpn1 bridge-domains vlan101 routing-interface irb.101
set routing-instances evpn1 bridge-domains vlan101 vxlan vni 101
set routing-instances evpn1 bridge-domains vlan101 vxlan ingress-node-replication
set interfaces irb unit 100 virtual-gateway-accept-data
set interfaces irb unit 100 family inet address 10.200.100.251/24 virtual-gateway-
address 10.200.100.254
set interfaces irb unit 100 family inet6 address 2001:0db8:0200:100::251/64 virtual-
gateway-address 2001:0db8:0200:100::254
set interfaces irb unit 101 virtual-gateway-accept-data
set interfaces irb unit 101 family inet address 10.200.101.251/24 virtual-gateway-
address 10.200.101.254
set interfaces irb unit 101 family inet6 address 2001:0db8:0200:101::251/64 virtual-
gateway-address 2001:0db8:0200:101::254
```

vQFX1:

© 2020 Juniper Networks, Inc. All rights reserved.

```
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay local-as 65001
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.200.17.7 peer-as 65007
set protocols bgp group underlay neighbor 10.200.18.8 peer-as 65008
set policy-options policy-statement export-loopback term 1 from interface lo0.0
set policy-options policy-statement export-loopback term 1 then accept
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 10.200.0.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65000
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 10.200.0.7
set protocols bgp group overlay neighbor 10.200.0.8
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 10.200.0.1:1
set switch-options vrf-import fabric
set switch-options vrf-target target:65000:1
set protocols evpn vni-options vni 100 vrf-target target:65000:100
set protocols evpn vni-options vni 101 vrf-target target:65000:101
set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list [ 100 101 ]
set policy-options community esi members target:65000:1
set policy-options community evpn1 members target:65000:1000
set policy-options community vni100 members target:65000:100
set policy-options community vni101 members target:65000:101
set policy-options policy-statement fabric term 1 from community [ esi evpn1 vni100
vni101 ]
set policy-options policy-statement fabric term 1 then accept
set vlans server vxlan vni 101
set vlans user vxlan vni 100
```

Task 9: Class of Service

This task is focused on Class of Services.

EX Class of Service

In this section you will configure virtual chassis EX1 in Branch 2 to support Class of Service.

1) Configure the following Forwarding Classes and map them to the output Queues as indicated in Table 11.

Forwarding Class	Queue	Schedulers L2 and L3 core Interfaces	Schedulers L2 access ports

© 2020 Juniper Networks, Inc. All rights reserved.

control	7	control-sched-core	control-sched-user
voice	5	voice-sched-core	voice-sched-user
server-gold	3	server-gold-sched	server-gold-sched
server-silver	2	server-silver-sched	server-silver-sched
server-bronze	1	server-bronze-sched	server-bronze-sched
best-effort	0	best-effort-sched	best-effort-sched

Table 49

EX1:

```
lab@ex1# show class-of-service forwarding-classes
queue 0 best-effort;
queue 1 server-bronze;
queue 2 server-silver;
queue 3 server-gold;
queue 5 voice;
queue 7 control;
```

2) Configure schedulers with the parameters as described in the Table below.

Scheduler	Parameter	Value
control-sched-core	priority	strict-high
	shaping rate	5%
	buffer-size	10%
control-sched-user	priority	strict-high
	shaping rate	2%
	buffer-size	10%
voice-sched-core	priority	strict-high
	shaping rate	10%
	buffer-size	5% exact
voice-sched-user	priority	strict-high
	shaping rate	5%
	buffer-size	5% exact
server-gold-sched	priority	strict-high
	shaping rate	25%
	buffer-size	10% exact
server-silver-sched	priority	low
	transmit rate	35%
	buffer-size	50%
server-bronze-sched	priority	low
	transmit rate	remainder

© 2020 Juniper Networks, Inc. All rights reserved.

best-effort-sched	buffer-size	remainder
	priority	low
	transmit rate	remainder
	buffer-size	remainder

Table 50

EX1:

```
lab@ex1# show class-of-service schedulers
control-sched-core {
    shaping-rate percent 5;
    buffer-size percent 10;
    priority strict-high;
}
control-sched-user {
    shaping-rate percent 2;
    buffer-size percent 10;
    priority strict-high;
}
voice-sched-core {
    shaping-rate percent 10;
    buffer-size {
        percent 5;
        exact;
    }
    priority strict-high;
}
voice-sched-user {
    shaping-rate percent 5;
    buffer-size {
        percent 5;
        exact;
    }
    priority strict-high;
}
server-gold-sched {
    shaping-rate percent 25;
    buffer-size {
        percent 10;
        exact;
    }
    priority strict-high;
}
server-silver-sched {
    transmit-rate percent 35;
    buffer-size percent 50;
    priority low;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
server-bronze-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority low;
}
best-effort-sched {
    transmit-rate {
        remainder;
    }
    buffer-size {
        remainder;
    }
    priority low;
}
```

3) Map the Schedulers to the forwarding classes as outlined in Table 13.

Forwarding Class	Schedulers L2 and L3 core Interfaces	Schedulers L2 access ports
control	control-sched-core	control-sched-user
voice	voice-sched-core	voice-sched-user
server-gold	server-gold-sched	server-gold-sched
server-silver	server-silver-sched	server-silver-sched
server-bronze	server-bronze-sched	server-bronze-sched
best-effort	best-effort-sched	best-effort-sched

Table 13

EX1:

```
lab@ex1# show class-of-service scheduler-maps
cos-sched-map-core {
    forwarding-class control scheduler control-sched-core;
    forwarding-class voice scheduler voice-sched-core;
    forwarding-class server-gold scheduler server-gold-sched;
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
cos-sched-map-user {
    forwarding-class control scheduler control-sched-user;
    forwarding-class voice scheduler voice-sched-user;
    forwarding-class server-gold scheduler server-gold-sched;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    forwarding-class server-silver scheduler server-silver-sched;
    forwarding-class server-bronze scheduler server-bronze-sched;
    forwarding-class best-effort scheduler best-effort-sched;
}
```

4) Apply the schedulers to the physical interfaces on virtual chassis EX1. Make sure that OoB interfaces are not affected.

EX1:

```
lab@ex1# show class-of-service interfaces
ge-*/0/* {
    scheduler-map cos-sched-map-user;
}
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
}
ge-0/0/4 {
    scheduler-map cos-sched-map-core;
}
```

5) Configure IPv4, IPv6 and dot1p Behavior Aggregate classifiers as shown in the table below. Make sure that the DSCP bits are configured as default code-point aliases. Ensure that default DSCP classifiers are used.

Forwarding Class	Dot1p	Dot1p Loss-Priority	DSCP Value	DSCP Loss-Priority
control	111	low	110000	low
			111000	low
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high
server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 14

EX1:

```
lab@ex1# show class-of-service classifiers
dscp ba-classifier {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
forwarding-class voice {
    loss-priority low code-points ef;
}
forwarding-class server-gold {
    loss-priority low code-points af41;
}
forwarding-class server-silver {
    loss-priority low code-points af21;
    loss-priority high code-points af23;
}
forwarding-class server-bronze {
    loss-priority low code-points af11;
    loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
dscp-ipv6 ba-classifier-v6 {
    import default;
    forwarding-class control {
        loss-priority low code-points [ nc1 nc2 ];
    }
    forwarding-class voice {
        loss-priority low code-points ef;
    }
    forwarding-class server-gold {
        loss-priority low code-points af41;
    }
    forwarding-class server-silver {
        loss-priority low code-points af21;
        loss-priority high code-points af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-points af11;
        loss-priority high code-points af13;
    }
    forwarding-class best-effort {
        loss-priority low code-points be;
    }
}
ieee-802.1 cos {
    forwarding-class best-effort {
        loss-priority low code-points 000;
    }
    forwarding-class server-bronze {
        loss-priority low code-points 001;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    forwarding-class server-silver {
        loss-priority low code-points 010;
    }
    forwarding-class server-gold {
        loss-priority low code-points 011;
    }
    forwarding-class voice {
        loss-priority low code-points 101;
    }
    forwarding-class control {
        loss-priority low code-points 111;
    }
}
```

6) Apply the classifiers to all interfaces as described in the Table below. Make sure that the OoB interfaces are not affected.

Virtual chassis	Classifiers	Interfaces
EX1	BA IPv4, IPv6	All L3 interfaces
	Dot1p	All L2 interfaces

Table 15

EX1:

```
lab@ex1# show class-of-service interfaces
ge-*/0/* {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
    }
}
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-0/0/4 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
}
ge-0/0/6 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-0/0/8 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-1/0/4 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
ge-1/0/10 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
    }
}
```

7) Configure the rewrite rules for marking the outgoing traffic with correct DSCP values.

Forwarding class	Dot1p	Dot1p Loss-Priority	DSCP Value	DSCP Loss-Priority
control	111	low	110000	low
			111000	low
voice	101		101110	low
server-gold	011		100010	low
server-silver	010		010010	low
			010110	high

© 2020 Juniper Networks, Inc. All rights reserved.

server-bronze	001		001010	low
			001110	high
best-effort	000		000000	high

Table 16

EX1:

```
lab@ex1# show class-of-service rewrite-rules
dscp rewrite-v4 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
dscp-ipv6 rewrite-v6 {
    forwarding-class control {
        loss-priority low code-point nc1;
        loss-priority high code-point nc2;
    }
    forwarding-class voice {
        loss-priority low code-point ef;
    }
    forwarding-class server-gold {
        loss-priority low code-point af41;
    }
    forwarding-class server-silver {
        loss-priority low code-point af21;
        loss-priority high code-point af23;
    }
    forwarding-class server-bronze {
        loss-priority low code-point af11;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        loss-priority high code-point af13;
    }
    forwarding-class best-effort {
        loss-priority high code-point be;
    }
}
ieee-802.1 cos {
    forwarding-class best-effort {
        loss-priority low code-point 000;
    }
    forwarding-class server-bronze {
        loss-priority low code-point 001;
    }
    forwarding-class server-silver {
        loss-priority low code-point 010;
    }
    forwarding-class server-gold {
        loss-priority low code-point 011;
    }
    forwarding-class voice {
        loss-priority low code-point 101;
    }
    forwarding-class control {
        loss-priority low code-point 111;
    }
}
```

8) Apply the rewrite-rules to all interfaces as described in the table below. Make sure that the OoB interfaces are not affected.

Virtual chassis	Rewrite-rules	Interfaces
EX1	IPv4, IPv6	All L3 interfaces
	Dot1p rewrite rules	All L2 interfaces

Table 17

EX1:

```
ge-*/0/* {
    scheduler-map cos-sched-map-user;
    unit * {
        classifiers {
            ieee-802.1 cos;
        }
        rewrite-rules {
            ieee-802.1 cos;
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/3 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-0/0/4 {
    scheduler-map cos-sched-map-core;
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-0/0/6 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-0/0/8 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
ge-1/0/4 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
ge-1/0/10 {
    unit * {
        classifiers {
            dscp ba-classifier;
            dscp-ipv6 ba-classifier-v6;
        }
        rewrite-rules {
            dscp rewrite-v4;
            dscp-ipv6 rewrite-v6;
        }
    }
}
```

Configurations for Task Completion:

EX1:

```
set class-of-service forwarding-classes queue 0 best-effort
set class-of-service forwarding-classes queue 1 server-bronze
set class-of-service forwarding-classes queue 2 server-silver
set class-of-service forwarding-classes queue 3 server-gold
set class-of-service forwarding-classes queue 5 voice
set class-of-service forwarding-classes queue 7 control
set class-of-service schedulers control-sched-core shaping-rate percent 5
set class-of-service schedulers control-sched-core buffer-size percent 10
set class-of-service schedulers control-sched-core priority strict-high
set class-of-service schedulers control-sched-user shaping-rate percent 2
set class-of-service schedulers control-sched-user buffer-size percent 10
set class-of-service schedulers control-sched-user priority strict-high
set class-of-service schedulers voice-sched-core shaping-rate percent 10
set class-of-service schedulers voice-sched-core buffer-size percent 5
set class-of-service schedulers voice-sched-core buffer-size exact
set class-of-service schedulers voice-sched-core priority strict-high
set class-of-service schedulers voice-sched-user shaping-rate percent 5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service schedulers voice-sched-user buffer-size percent 5
set class-of-service schedulers voice-sched-user buffer-size exact
set class-of-service schedulers voice-sched-user priority strict-high
set class-of-service schedulers server-gold-sched shaping-rate percent 25
set class-of-service schedulers server-gold-sched buffer-size percent 10
set class-of-service schedulers server-gold-sched buffer-size exact
set class-of-service schedulers server-gold-sched priority strict-high
set class-of-service schedulers server-silver-sched transmit-rate percent 35
set class-of-service schedulers server-silver-sched buffer-size percent 50
set class-of-service schedulers server-silver-sched priority low
set class-of-service schedulers server-bronze-sched transmit-rate remainder
set class-of-service schedulers server-bronze-sched buffer-size remainder
set class-of-service schedulers server-bronze-sched priority low
set class-of-service schedulers best-effort-sched transmit-rate remainder
set class-of-service schedulers best-effort-sched buffer-size remainder
set class-of-service schedulers best-effort-sched priority low
set class-of-service scheduler-maps cos-sched-map-core forwarding-class control
scheduler control-sched-core
set class-of-service scheduler-maps cos-sched-map-core forwarding-class voice scheduler
voice-sched-core
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map-core forwarding-class best-effort
scheduler best-effort-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class control
scheduler control-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class voice scheduler
voice-sched-user
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-gold
scheduler server-gold-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-silver
scheduler server-silver-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class server-bronze
scheduler server-bronze-sched
set class-of-service scheduler-maps cos-sched-map-user forwarding-class best-effort
scheduler best-effort-sched
set class-of-service interfaces ge-*/0/* scheduler-map cos-sched-map-user
set class-of-service interfaces ge-0/0/3 scheduler-map cos-sched-map-core
set class-of-service interfaces ge-0/0/4 scheduler-map cos-sched-map-core
set class-of-service classifiers dscp ba-classifier import default
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc1
set class-of-service classifiers dscp ba-classifier forwarding-class control loss-
priority low code-points nc2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service classifiers dscp ba-classifier forwarding-class voice loss-priority
low code-points ef
set class-of-service classifiers dscp ba-classifier forwarding-class server-gold loss-
priority low code-points af41
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority low code-points af21
set class-of-service classifiers dscp ba-classifier forwarding-class server-silver loss-
priority high code-points af23
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority low code-points af11
set class-of-service classifiers dscp ba-classifier forwarding-class server-bronze loss-
priority high code-points af13
set class-of-service classifiers dscp ba-classifier forwarding-class best-effort loss-
priority low code-points be
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 import default
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc1
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class control
loss-priority low code-points nc2
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class voice loss-
priority low code-points ef
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-gold
loss-priority low code-points af41
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority low code-points af21
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
silver loss-priority high code-points af23
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority low code-points af11
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class server-
bronze loss-priority high code-points af13
set class-of-service classifiers dscp-ipv6 ba-classifier-v6 forwarding-class best-effort
loss-priority low code-points be
set class-of-service classifiers ieee-802.1 cos forwarding-class best-effort loss-
priority low code-points 000
set class-of-service classifiers ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-points 001
set class-of-service classifiers ieee-802.1 cos forwarding-class server-silver loss-
priority low code-points 010
set class-of-service classifiers ieee-802.1 cos forwarding-class server-gold loss-
priority low code-points 011
set class-of-service classifiers ieee-802.1 cos forwarding-class voice loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 cos forwarding-class control loss-priority
low code-points 111
set class-of-service interfaces ge-*/0/* unit * classifiers ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-0/0/3 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-0/0/4 unit * classifiers dscp ba-classifier
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service interfaces ge-0/0/4 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-0/0/6 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-0/0/6 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-0/0/8 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-0/0/8 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-1/0/4 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-1/0/4 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service interfaces ge-1/0/10 unit * classifiers dscp ba-classifier
set class-of-service interfaces ge-1/0/10 unit * classifiers dscp-ipv6 ba-classifier-v6
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class voice loss-priority
low code-point ef
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-gold loss-
priority low code-point af41
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority low code-point af21
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-silver loss-
priority high code-point af23
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority low code-point af11
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class server-bronze loss-
priority high code-point af13
set class-of-service rewrite-rules dscp rewrite-v4 forwarding-class best-effort loss-
priority high code-point be
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority low code-point nc1
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class control loss-
priority high code-point nc2
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class voice loss-
priority low code-point ef
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-gold
loss-priority low code-point af41
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority low code-point af21
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-silver
loss-priority high code-point af23
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority low code-point af11
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class server-bronze
loss-priority high code-point af13
set class-of-service rewrite-rules dscp-ipv6 rewrite-v6 forwarding-class best-effort
loss-priority high code-point be
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class best-effort loss-
priority low code-point 000
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-bronze loss-
priority low code-point 001
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-silver loss-
priority low code-point 010
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class server-gold loss-
priority low code-point 011
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class voice loss-priority
low code-point 101
set class-of-service rewrite-rules ieee-802.1 cos forwarding-class control loss-priority
low code-point 111
set class-of-service interfaces ge-*/0/* unit * rewrite-rules ieee-802.1 cos
set class-of-service interfaces ge-0/0/3 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/3 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-0/0/6 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/6 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-0/0/8 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-0/0/8 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-1/0/4 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-1/0/4 unit * rewrite-rules dscp-ipv6 rewrite-v6
set class-of-service interfaces ge-1/0/10 unit * rewrite-rules dscp rewrite-v4
set class-of-service interfaces ge-1/0/10 unit * rewrite-rules dscp-ipv6 rewrite-v6
```

This completes Super Lab 2.

© 2020 Juniper Networks, Inc. All rights reserved.