JNCIE-DC 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-DC-BOOK

Juniper Business Use Only

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-DC Certification Self-Study Bundle, Revision 1.b

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

Printed in USA.

Revision History:

Revision 1.a—December 2018
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.

Juniper Business Use Only

Contents

Contents 3
About This Publication 21
Education Services Offerings 21
Technical Publications 21
Juniper Networks Support 21
Target Audience 22
Exam strategy 22
Lab Preparation 22
Day Before the Exam 23
Exam Day 23
After the Exam 24
Workbook and Configuration Files 25
Chapter 1: Layer 2 Underlay 26
Part 1: MC-LAG 36
Task 1.1: ToR Configuration 37
Task 1.2: MC-LAG Peers: Server-Facing Configuration 37
Task 1.3: MC-LAG Peers: Core-Facing Configuration 37
Task 1.4: vMX LAG Configuration 37
Task 1.5: vMX MC-LAG Configuration 38
Task 1.6: vMX MC-LAG Gateway Configuration 38
Part 2: Multistage MC-LAG 39
Task 1.7: DC1 LAG 41

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.8: DC1 Multistage MC-LAG 41
Task 1.9: DC1 Gateway and OSPF Configuration 41
Task 1.10: DC1 Security 42
Task 1.11: DC2 LAG and MC-LAG Configuration 42
Task 1.12: DC2 Gateway, VRRP and OSPF Configuration 42
Task 1.13: MC-LAG Verification 42
Part 3: VCF 44
Task 1.14: VCF Configuration 45
Task 1.15: VCF Features 45
Task 1.16: Deconstructing the VCF 45
Task 1.17: Autoprovisioning a VCF 45
Task 1.18: VCF VLAN and Interface Configuration 45
Chapter 2: Clos IP Fabric 46
Part 1: EBGP-based Clos IP Fabric 49
Task 2.1: Core Network BGP Configuration 1 50
Task 2.2: Core Network BGP Configuration 2 51
Task 2.3: ISP Uplink Configuration 51
Task 2.4: IP-Fabric BGP Customer Configuration 51
Task 2.5: Server 100-110 Internet Feed. 52
Task 2.6: Server 113-115 Configuration 52
Task 2.7: IP-Fabric Customer Policy Configuration 52
Task 2.8: IP-Fabric Customer Policy Configuration 52
Task 2.9: ISP Policy Configuration 52

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: IBGP-based Clos IP Fabric 53
Task 2.10: IGP Configuration 54
Task 2.11: IGP Configuration 54
Task 2.12: IGP Configuration 54
Task 2.13: IPv4 IBGP Configuration 55
Task 2.14: IPv6 IBGP Configuration 55
Task 2.15: Server 101-110 55
Task 2.16: EBGP Configuration 55
Task 2.17: Local AS Advertisement 55
Task 2.18: BGP RIB Configuration 56
Task 2.19: AS65000 Customers 56
Task 2.20: IP-Transit Policy Configuration 56
Chapter 3: Controllerless Overlay 56
Part 1: A Basic EBGP-based Overlay Network 64
Task 3.1: Configure the EBGP Underlay 65
Task 3.2: EBGP Underlay Optimization 65
Task 3.3: IBGP Overlay Configuration 66
Task 3.4: Configure QFX Switch-level EVPN 66
Task 3.5: VXLAN Segment Configuration 66
Task 3.6: Active-Active Multihoming 66
Task 3.7: Configure vMX Virtual-Switch-Level EVPN 66
Task 3.8: Redundant Layer 3 VXLAN Gateway Configuration 67
Task 3.9: Inter-VXLAN Routing 67

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: IBGP-based Overlay Network 68
Task 3.10: IGP Configuration 69
Task 3.11: IGP Optimization 69
Task 3.12: IBGP Configuration 69
Task 3.13: Configure QFX Switch-Level EVPN 70
Task 3.14: VXLAN Segment Configuration and Active-Active Multihoming 70
Task 3.15: VXLAN Segment Configuration for VNI 101 70
Task 3.16: Additional VXLAN Segment Configuration 70
Task 3.17: MX Virtual-Switch with Selective VXAN Import 71
Task 3.18: Layer 3 VXLAN Gateway Configuration 71
Task 3.19: Layer 3 VXLAN Gateway Configuration 71
Task 3.20: Layer 3 VXLAN Gateway Configuration 71
Part 3: IP Fabric and Controllerless Overlay 72
Task 3.21: EBGP for the IP-Fabric 73
Task 3.22: Completing the IP-Fabric 74
Task 3.23: Controllerless Overlay 74
Task 3.24: Expanding the Controllerless Overlay 74
Task 3.25: IP-Fabric Routing Policies 75
Chapter 4: Data Center Interconnect 75
Part 1: MPLS and MPLS L3VPN 83
Task 4.1: IGP Configuration 85
Task 4.2: LDP Configuration 85
Task 4.3: MPLS LSP Options Configuration 85

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.4: BGP Configuration 85
Task 4.5: MPLS L3VPN JNCIE-BLUE 85
Task 4.6: MPLS L3VPN JNCIE-GREEN 86
Task 4.7: MPLS L3VPN JNCIE-RED 86
Part 2: MPLS, MPLS L3VPN and EVPN. 87
Task 4.8: RSVP Configuration 89
Task 4.9: RSVP LSP Configuration 89
Task 4.10: BGP Configuration 89
Task 4.11: EVPN JNCIE-PURPLE 89
Task 4.12: EVPN JNCIE-BLUE 89
Task 4.13: EVPN JNCIE-GREEN 89
Task 4.14: EVPN JNCIE-RED 90
Part 3: MPLS L3VPN and VLAN-aware EVPN 91
Task 4.15: EVPN North Configuration 93
Task 4.16: EVPN South Configuration 93
Task 4.17: EVPN Multihoming 93
Task 4.18: VLAN-based EVPN Configuration 94
Part 4: Controllerless overlay and EVPN DCI 95
Task 4.19: DC1 Controllerless Overlay Part 1: Routing and Signaling 95
Task 4.20: DC1 Controllerless Overlay Part 2: the Overlay 95
Task 4.21: DC2 Controllerless Overlay Part 1: The Underlay 96
Task 4.22: DC2 Controllerless Overlay part 2: The Overlay 96
Task 4.23: EVPN Stitching 96

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Chapter 5: Security 96
Part 1: Securing the datacenter 101
Task 5.1: vSRX Zone and Interface Configuration 101
Task 5.2: vSRX Policy Configuration 101
Task 5.3: vSRX NAT 102
Task 5.4: vSRX Screens 102
Task 5.5: vSRX RE protection 102
Task 5.6: Limiting bandwidth using a Stateless Firewall Filter 103
Chapter 6: Class of Service 103
Part 1: CoS Inside the Data Center. 107
Task 6.1: Forwarding Classes 108
Task 6.2: Scheduling 108
Task 6.3: WRED 108
Task 6.4: BA classifiers 108
Task 6.5: Rewrite Rules 109
Task 6.6: Traffic Classification on vQFX3 and vQFX4 109
Task 6.7: Traffic Classification on vQFX5 109
Chapter 7: Management 109
Part 1: Data Center Management 115
Task 7.1: Zero Touch Provisioning (ZTP) 115
Task 7.2: Python using NETCONF 115
Task 7.3: Junos Space 115
Task 7.4: On Box Scripting 1 115

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 7.5: On Box Scripting 2 116
Task 7.6: On Box Scripting 3 117
Super Lab 1 117
Topology 118
Task 1: Device management 119
Task 1.1: Zero Touch Provisioning (ZTP) 119
Task 1.2: Junos Space 119
Task 1.3: On Box Scripting 119
Task 1.4: NETCONF 119
Task 2: Underlay 119
Task 2.1: Multi Chassis Link Aggregation (MC-LAG) 119
Task 2.2: Layer 2 Underlay Configuration 120
Task 2.3: Layer 2 Underlay Features 120
Task 2.4: BUM Traffic Handling 120
Task 2.5: Layer 2 Redundancy 120
Task 2.6: Layer 3 Underlay MAIN DC 121
Task 2.7: Layer 3 Underlay Routing 121
Task 2.8: EBGP Underlay 121
Task 3: Controllerless Overlay 122
Task 3.1: EVPN Overlay Configuration 122
Task 3.2: EVPN/VLAN Service Configuration 122
Task 3.3: EVPN/VXLAN Service Tuning 123
Task 3.4: EVPN Signaling 123

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4: Data Center Interconnect 123
Task 4.1: DCI Configuration 123
Task 4.2: DC1 Redundancy 125
Task 4.3: (Inter) Data Center Connectivity 126
Task 4.4: Data Center Connectivity to External Server 126
Task 5: Security 126
Task 5.1: User Authentication and Authorization 126
Task 5.2: SRX Configuration 126
Task 5.3: RE Protection 127
Task 6: Class of Service 127
Task 6.1: Forwarding Classes 127
Task 6.2: Schedulers and Traffic Profiles 127
Task 6.3: Policing 128
Super Lab 2 128
Topology 128
Task 1: Device Management 129
Task 1.1: Zero Touch Provisioning (ZTP) 129
Task 1.2: Junos Space 129
Task 1.3: On Box Scripting 129
Task 1.4: NETCONF 129
Task 2: Underlay 129
Task 2.1: Layer 2 Underlay Configuration 129
Task 2.2: Multi Chassis Link Aggregation (MC-LAG) 130

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.3: Layer 2 Underlay Features 130
Task 2.4: BUM Traffic Handling 130
Task 2.5: Layer 2 Redundancy 130
Task 2.6: Layer 3 Underlay MAIN DC 131
Task 2.7: Layer 3 Underlay Routing DC1 131
Task 2.8: Layer 2 Underlay Switching vQFX1 131
Task 2.9: EBGP Underlay 131
Task 3: Controllerless Overlay 132
Task 3.1: EVPN Overlay Configuration 132
Task 3.2: EVPN/VLAN Service Configuration 132
Task 3.3: EVPN/VXLAN Service Tuning 133
Task 3.4: EVPN Signaling 133
Task 4: Data Center Interconnect 133
Task 4.1: DCI Configuration 133
Task 4.2: DC1 Configuration 135
Task 4.3: (Inter) Data Center Connectivity 136
Task 4.4: Data Center Connectivity to External Server 136
Task 4.5: EVPN VXLAN / EVPN MPLS Interworking 136
Task 5: Security 136
Task 5.1: User Authentication and Authorization 136
Task 5.2: SRX Configuration 136
Task 5.3: RE Protection 137
Task 6: Class of Service 137

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 6.1: Forwarding Classes 137
Task 6.2: Schedulers and Traffic Profiles 137
Task 6.3: Policing 138
Appendices 139
Appendix – Chapter 1: Layer 2 Underlay 139
Part 1: MC-LAG 139
Task 1.1: ToR Configuration 140
Task 1.2: MC-LAG Peers: Server-Facing Configuration 143
Task 1.3: MC-LAG Peers: Core-Facing Configuration 152
Task 1.4: vMX LAG Configuration 153
Task 1.5: vMX MC-LAG Configuration 158
Task 1.6: vMX MC-LAG Gateway Configuration 165
Part 2: Multistage MC-LAG 169
Task 1.7: DC1 LAG 170
Task 1.8: DC1 Multistage MC-LAG 173
Task 1.9: DC1 Gateway and OSPF Configuration 180
Task 1.10: DC1 Security 185
Task 1.11: DC2 LAG and MC-LAG Configuration 187
Task 1.12: DC2 Gateway, VRRP, and OSPF Configuration 198
Task 1.13: MC-LAG Verification 205
Part 3: VCF 217
Task 1.14: VCF Configuration 218
Task 1.15: VCF Features 222

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.16: Deconstructing the VCF 224
Task 1.17: Autoprovisioning a VCF 227
Task 1.18: VCF VLAN and Interface Configuration 231
Appendix – Chapter 2: Clos IP Fabric 246
Part 1: EBGP-based Clos IP Fabric 246
Task 2.1: Core Network BGP Configuration 1 248
Task 2.2: Core Network BGP Configuration 2 253
Task 2.3: ISP Uplink Configuration 254
Task 2.4: IP-Fabric BGP Customer Configuration 256
Task 2.5: Server 100-110 Internet Feed. 260
Task 2.6: Server 113-115 Configuration 262
Task 2.7: IP-Fabric Customer Policy Configuration 269
Task 2.8: IP-Fabric Customer Policy Configuration 270
Task 2.9: ISP Policy Configuration 277
Part 2: IBGP-based Clos IP Fabric 282
Task 2.10: IGP Configuration 283
Task 2.11: IGP Configuration 288
Task 2.12: IGP Configuration 291
Task 2.13: IPv4 IBGP Configuration 293
Task 2.14: IPv6 IBGP Configuration 296
Task 2.15: Server 101-110 299
Task 2.16: EBGP Configuration 302
Task 2.17: Local AS Advertisement 313

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.18: BGP RIB Configuration 318
Task 2.19: AS65000 Customers 322
Task 2.20: IP-Transit Policy Configuration 325
Appendix – Chapter 3: Controllerless Overlay 330
Part 1: A basic EBGP-based Overlay Network 330
Task 3.1: Configure the EBGP underlay 331
Task 3.2: EBGP Underlay Optimization 336
Task 3.3: IBGP Overlay Configuration 340
Task 3.4: Configure QFX Switch-Level EVPN 346
Task 3.5: VXLAN Segment Configuration 350
Task 3.6: Active-Active Multihoming 359
Task 3.7: Configure vMX Virtual-Switch-Level EVPN 370
Task 3.8: Redundant Layer 3 VXLAN Gateway Configuration 377
Task 3.9: Inter-VXLAN Routing 380
Part 2: IBGP-based Overlay Network 383
Task 3.10: IGP Configuration 384
Task 3.11: IGP Optimization 387
Task 3.12: IBGP Configuration 391
Task 3.13: Configure QFX Switch-Level EVPN 394
Task 3.14: VXLAN Segment Configuration and Active-Active Multihoming ... 397
Task 3.15: VXLAN Segment Configuration for VNI 101 403
Task 3.16: Additional VXLAN Segment Configuration 406
Task 3.17: MX Virtual-Switch with Selective VXAN Import 415

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.18: Layer 3 VXLAN Gateway Configuration 422
Task 3.19: Layer 3 VXLAN Gateway Configuration 428
Task 3.20: Layer 3 VXLAN Gateway Configuration 432
Part 3: IP Fabric and Controllerless Overlay 437
Task 3.21: EBGP for the IP-Fabric 438
Task 3.22: Completing the IP-Fabric 444
Task 3.23: Controllerless Overlay 447
Task 3.24: Expanding the Controllerless Overlay 452
Task 3.25: IP-Fabric Routing Policies 457
Appendix – Chapter 4: Data Center Interconnect 462
Part 1: MPLS and MPLS L3VPN 462
Task 4.1: IGP Configuration 464
Task 4.2: LDP Configuration 467
Task 4.3: MPLS LSP Options Configuration 475
Task 4.4: BGP Configuration 479
Task 4.5: MPLS L3VPN JNCIE-BLUE 482
Task 4.6: MPLS L3VPN JNCIE-GREEN 489
Task 4.7: MPLS L3VPN JNCIE-RED 494
Part 2: MPLS, MPLS L3VPN and EVPN. 502
Task 4.8: RSVP Configuration 504
Task 4.9: RSVP LSP Configuration 506
Task 4.10: BGP Configuration 511
Task 4.11: EVPN JNCIE-PURPLE 513

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.12: EVPN JNCIE-BLUE 518
Task 4.13: EVPN JNCIE-GREEN 523
Task 4.14: EVPN JNCIE-RED 529
Part 3: MPLS L3VPN and VLAN-Aware EVPN. 536
Task 4.16: EVPN North Configuration 538
Task 4.17: EVPN South Configuration 545
Task 4.18: EVPN Multihoming 556
Task 4.19: VLAN-Based EVPN Configuration 576
Part 4: Controllerless Overlay and EVPN DCI 581
Task 4.20: DC1 Controllerless Overlay Part 1: routing and signaling 581
Task 4.21: DC1 Controllerless Overlay Part 2: the overlay 584
Task 4.22: DC2 Controllerless Overlay Part 1: the underlay 589
Task 4.23: DC2 Controllerless Overlay Part 2: the overlay 593
Task 4.24: EVPN Stitching 603
Appendix – Chapter 5: Security 615
Part 1: Securing the Data Center 615
Task 5.1: vSRX Zone and Interface Configuration 615
Task 5.2: vSRX Policy Configuration 619
Task 5.3: vSRX NAT 626
Task 5.4: vSRX Screens 634
Task 5.5: vSRX RE Protection 637
Task 5.6: Limiting Bandwidth using a Stateless Firewall Filter 640
Appendix – Chapter 6: Class of Service 643

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: CoS Inside the Data Center 643
Task 6.1: Forwarding Classes 644
Task 6.2: Scheduling 645
Task 6.3: WRED 651
Task 6.4: BA Classifiers 654
Task 6.5: Rewrite Rules 656
Task 6.6: Traffic Classification on vQFX3 and vQFX4 658
Appendix – Chapter 7: Management 665
Part 1: Data Center Management 665
Task 7.1: Zero Touch Provisioning (ZTP) 666
Task 7.2: Python using NETCONF 668
Task 7.3: Junos Space 669
Task 7.4: On Box Scripting 1 676
Task 7.5: On Box Scripting 2 678
Task 7.6: On Box Scripting 3 682
Appendix – Super Lab 1 684
Task 1: Device Management 684
Task 1.1: Zero Touch Provisioning (ZTP) 684
Task 1.2 : Junos Space 684
Task 1.3: On Box Scripting 684
Task 1.4: NETCONF 685
Task 2: Underlay 685
Task 2.1: Multi Chassis Link Aggregation (MC-LAG) 685

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.2: Layer 2 Underlay Configuration 688
Task 2.3: Layer 2 Underlay Features 689
Task 2.4: BUM Traffic Handling 690
Task 2.5: Layer 2 Redundancy 690
Task 2.6: Layer 3 Underlay MAIN DC 690
Task 2.7: Layer 3 Underlay Routing 691
Task 2.8: EBGP Underlay 692
Task 3: Controllerless Overlay 694
Task 3.1: EVPN Overlay Configuration 694
Task 3.2: EVPN/VLAN Service Configuration 697
Task 3.3: EVPN/VXLAN Service Tuning 700
Task 3.4: EVPN Signaling 701
Task 4: Data Center Interconnect 702
Task 4.1: DCI Configuration 702
Task 4.2: DC1 Redundancy 708
Task 4.3: (Inter) Data Center Connectivity 708
Task 4.4: Data Center Connectivity to External Server 708
Task 5: Security 709
Task 5.1: User Authentication and Authorization 709
Task 5.2: SRX Configuration 709
Task 5.3: RE Protection 710
Task 6: Class of Service 710
Task 6.1: Forwarding Classes 710

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 6.2: Schedulers and Traffic Profiles 711
Task 6.3: Policing 712
Appendix – Super Lab 2 713
Task 1: Device Management 713
Task 1.1: Zero Touch Provisioning (ZTP) 713
Task 1.2 : Junos Space 713
Task 1.3: On Box Scripting 713
Task 1.4: NETCONF 714
Task 2: Underlay 714
Task 2.1: Layer 2 Underlay Configuration 714
Task 2.2: Multi Chassis Link Aggregation (MC-LAG) 716
Task 2.3: Layer 2 Underlay Features 719
Task 2.4: BUM Traffic Handling 719
Task 2.5: Layer 2 Redundancy 720
Task 2.6: Layer 3 Underlay MAIN DC 720
Task 2.7: Layer 3 Underlay Routing DC1 720
Task 2.8: Layer 2 Underlay Switching vQFX1 721
Task 2.9: EBGP Underlay DC2 722
Task 3: Controllerless Overlay 724
Task 3.1: EVPN Overlay Configuration 724
Task 3.2: EVPN/VLAN Service Configuration 726
Task 3.3: EVPN/VXLAN Service Tuning 730
Task 3.4: EVPN signaling 730

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4: Data Center Interconnect 730
Task 4.1: DCI Configuration 730
Task 4.2: DC1 Redundancy 736
Task 4.3: (Inter) Data Center Connectivity 736
Task 4.4: Data Center Connectivity to External Server 737
Task 4.5: EVPN VXLAN / EVPN MPLS Interworking 738
Task 5: Security 738
Task 5.1: User Authentication and Authorization 738
Task 5.2: SRX Configuration 739
Task 5.3: RE Protection 739
Task 6: Class of Service 740
Task 6.1: Forwarding Classes 740
Task 6.2: Schedulers and Traffic Profiles 741
Task 6.3: Policing 742

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

About This Publication

The JNCIE-DC lab exam workbook is designed for candidates to practice technologies covered in the JNCIE-DC 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 (also 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-DC lab exam preparation workbook consists of 6 virtual MX Series devices running Junos version 16.1; 1 virtual SRX Series device running Junos version 12.1; 6 virtual QFX Series devices running Junos version 15.1; 1 Junos Space with Network Director running release 15.2; 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-DC lab exam. Although not required, it is highly recommended that you have passed the JNCIS-ENT and JNCIP-DC written exams before you start using this workbook. The *JNCIE-DC Certification Self-Study Bundle* is a hands-on guide to validate the skills needed to pass the official JNCIE-DC lab exam. This guide is based on official JNCIE-DC exam objectives and is targeted at JNCIP-DC 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-DC exam.

Exam strategy

Lab Preparation

To earn the JNCIE-DC credential you must first pass the following Juniper Networks written exams: JNCIA-Junos, JNCIS-ENT, and JNCIP-DC. Finally, you must also pass the very challenging 8 hour JNCIE-DC lab exam. You must have extensive hands-on experience with data center technologies on vSRX, vMX, and vQFX platforms. When preparing for the JNCIE-DC lab exam ensure you fully understand the exam topics and objectives. The latest update for the JNCIE-DC 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-DC lab exam as it covers very challenging tasks specifically written for JNCIE-DC lab preparation. The workbook tasks are written at the same difficulty level as the JNCIE-DC 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-DC lab exam.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Day Before the Exam

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

Exam Day

When you wake up, dress comfortable and have a good breakfast. Plan to arrive 30 minutes before the exam start time. Remember to bring a valid ID with you, exam registration confirmation, and if you prefer, your own keyboard and mouse. If you forget your ID, the JNCIE lab exam proctor will not allow you to take the exam. When it's time to enter the exam room, the JNCIE lab exam proctor will direct you to your workstation. Using the SecureCRT software, you can access the JNCIE console server, which provides console access to the exam devices. There are Juniper Networks Data Center PDF files installed on your computer, which you can use to lookup certain commands and features. You do not have Internet access.

When the proctor tells you that the exams begins, you are allowed to open the exam documents. Before immediately starting with the first task, we advise you to read the entire exam and study the topology first to fully understand all requirements and how tasks might correlate to each other.

Be familiar with the Windows Notepad application as it is a real time saver. For example, if you have to apply BGP configuration to multiple devices, with minimal changes, Notepad is your friend. For the JNCIE-DC exam also some basic Linux knowledge is required. Make sure you know to navigate through a Linux file system, use the VI application to edit files and be able to start/restart Linux services.

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

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

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

If you are struggling with a task for more than 10-15 minutes, check the task and determine whether or not it will have impact on the rest of your exam. If it does not impact any other task on the exam, move to the next task and revisit the question again later on during the exam.

Do not assume any other requirement then the specific ones mentioned in the tasks. The JNCIE exam is about fulfilling the tasks. It is not about displaying your understanding of best practices, scalable network design, or anything other than fulfilling the task. So, it might be that your route-reflector is not redundant, the cluster-id is completely random, the active hold-time is 90 seconds, and the sessions are not authenticated. As long as the tasks do not require any of these things, no one cares!

Sometimes there are multiple ways to complete a task. In that case, you can choose your preferred option. If you are unsure about a task requirement then ask the proctor. Do not be shy. While the JNCIE lab exam proctor cannot assist you in solving the tasks, he or she can certainly help you in clarifying the task requirements.

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

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

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

After the Exam

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

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Workbook and Configuration Files

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

Important: To access the devices in this environment please use the jncie/jncie123 user credentials. For access to the Junos Space Web UI use the super/jncie123 credentials.

Note: This workbook has unique configurations for most parts, chapters and Super Labs. You must load the start configuration on all devices in the topology to ensure that the environment is correct for each chapter and part. For example, if you would like to start with Chapter 4, you must load the initial configuration files for Chapter 4. Then, as you complete Part 1, you will need to load the start configuration for Part 2. The chapter's labs and exercises do not generally build upon the configuration you executed in the previous part. In addition to start configurations, final configurations have been provided for each chapter and can be used to verify that all tasks are completed and working properly.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Chapter 1: Layer 2 Underlay

Virtual Chassis Fabric (VCF) and Multi-Chassis Link Aggregation Group (MC-LAG) offer two different approaches to constructing layer 2 underlay networks.

The VCF is a Juniper proprietary solution that allows you to combine up to 20 devices that can be managed as a single device. A VCF is constructed along the lines of a 3-stage Clos topology and runs Juniper proprietary protocols. You'll mostly see topologies wherein the VCF is presented as a 3 stage Clos topology that is folded onto itself:

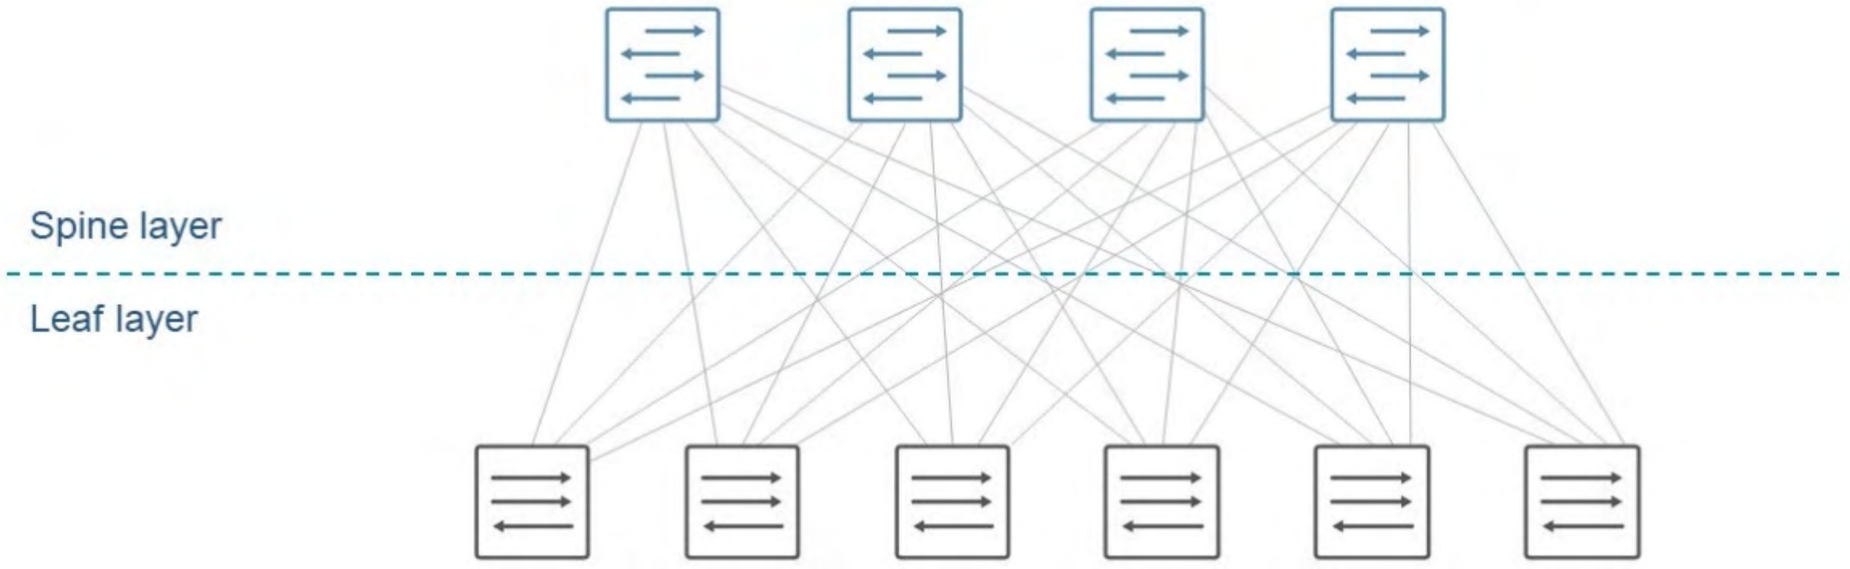

The MC-LAG approach is different in that it does not require any proprietary protocols. In an MC-LAG setup, every device runs its own control plane operations and is managed separately. MC-LAG peers are configured to appear as a single switch using 802.3ad:

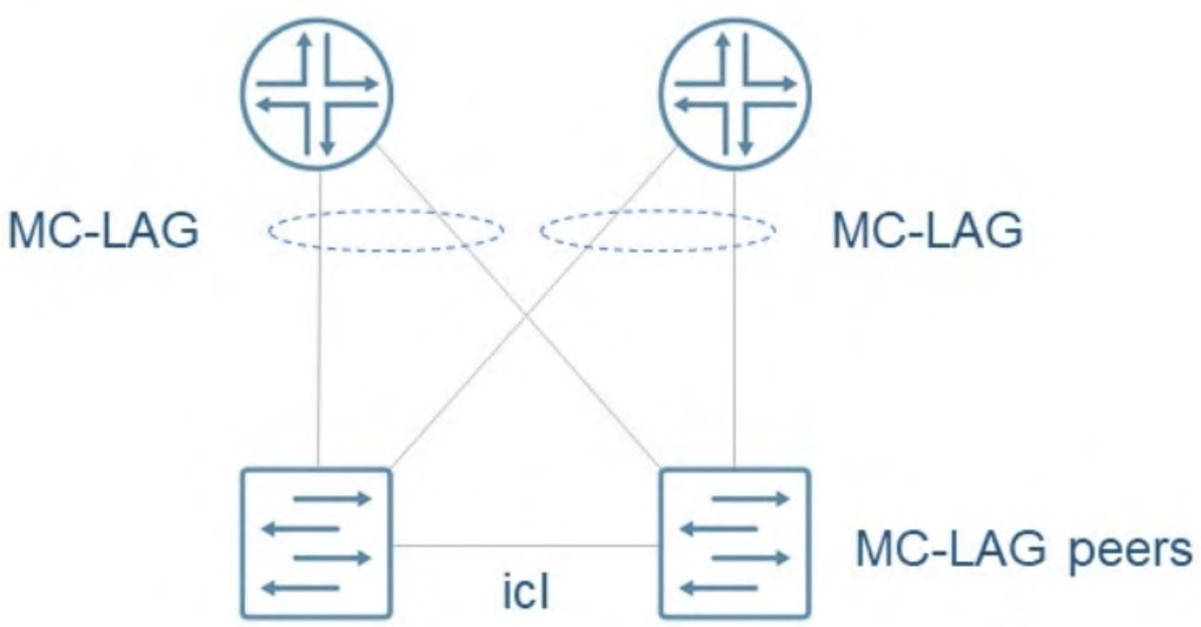

Both setups can be used to offer a high-available layer 2 underlay as well as a redundant layer 3 gateway.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Virtual Chassis Fabric

Junipers Virtual Chassis Fabric (VCF) technology can be used to construct a three stage Clos switching fabric. A VCF is made up of individual member switches that are placed into a spine and leaf architecture. The whole of the VCF is managed as a single device.

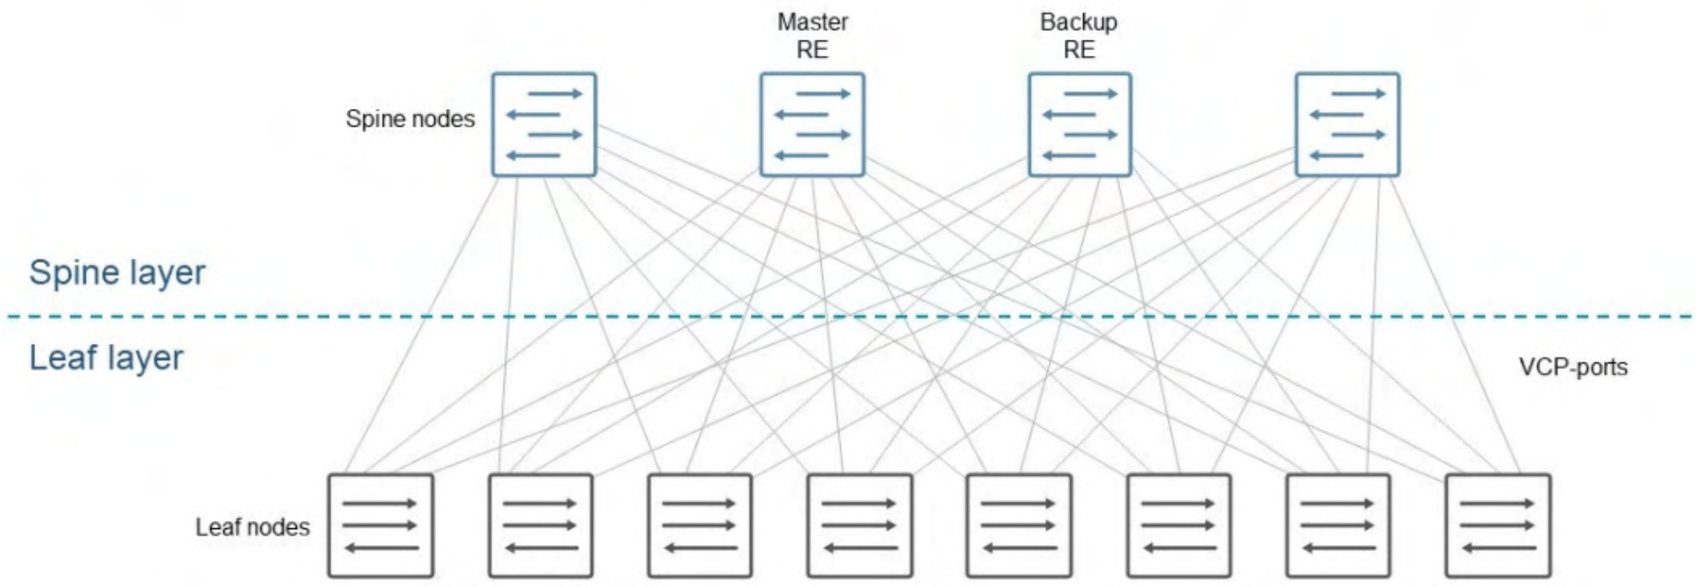

We can identify two layers in the VCF architecture; the **spine layer** and **the leaf layer**. The spine layer can contain up to four spine nodes. Every spine node should have 1 or more connections to every leaf node. The spines nodes can function as a Routing-Engine (RE) or as a Line Card (LC). Up to 16 leaf nodes can be connected in the leaf layer. Every leaf device functions as a LC.

Routing Engines

The spine layer can contain up to two devices that function as a RE. One RE is active and the other is standby. The active RE is called the master RE and the standby RE is called the backup RE. The whole of the VCF is controlled by the master RE. This means that that the master RE runs the control protocols and manages all the of the VCF member switches.

Juniper offers high-availability (HA) features to minimize the effects in case the master RE fails. These HA features include Graceful Routing Engine Switchover (GRES), Non-Stop Active Routing (NSR) and Non-Stop Bridging (NSB). GRES can be activated to preserve interface and kernel information on the backup RE. NSR will have the backup RE run the RPD whereas NSB will have the backup RE run the L2CPD. None of these HA features are active by default; they all require configuration.

Line cards

Devices that are operating in LC-mode run only subset of Junos. The ‘regular’ leaf nodes as well as spine nodes that were not selected as master or backup RE function in LC-mode.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

VCF management

Whenever you log in to a VCF using any of the members' console ports, you will be connected to the master RE. Every member switch inside a VCF runs virtual console software and will redirect all console traffic to the master RE. It is also possible to set up a vty session to another member switch from the master RE. This can be done through the use of the '**request session member x**' command.

The OoB interfaces on all the individual member switches of the VCF are automatically placed inside a management VLAN. The layer 3 interface tied to this management VLAN is called the 'vme' interface, or Virtual Management Ethernet interface. When you connect to this IP address, or any other IP address configured on the VCF, you will automatically be connected to the master RE.

Virtual Chassis Control Protocol

All switches inside the VCF run the Virtual Chassis Control Protocol (VCCP). The VCCP is a Juniper proprietary protocol that is based on IS-IS. Switches running VCCP exchange LSA-based discovery messages that enables them to discover the VCF topology. When devices are done building the topology, they run an SPF algorithm for every PFE. The result is a loop-free path between every PFE inside the VCF. When the switches run SPF, they can take into account multiple paths and will automatically load share traffic across the links.

VCF configuration options

There are three different provisioning options:

- **non-provisioned**: configure VCP and let the rest happen automatically.
- **pre-provisioned**: statically configure every member of the VCF by including the switch serial number, member-ID and role in the VCF configuration.
- **auto-provisioned**: pre-provision the RE switches only and have the other switches automatically join the VCF as LCs.

Virtual Chassis Ports

The connections between the spine and leaf nodes are called Virtual Chassis Port (VCP) connections. VCPs carry both control plane as well as forwarding plane traffic within the VCF.

By default, ports on QFX switches operate as 'normal' Ethernet ports. Any (non-channelized) QSFP+ or SFP+ port can be turned into a VCP. This can be done manually or automatically.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To manually change a port into a VCP, use the **'request virtual-chassis vc-port set'** command. To turn a VCP port into a normal port again, use the **'request virtual-chassis vc-port delete pic-slot x port x member x'** command.

To have a normal port change into a VCP automatically, make sure that the following requirements are met:

- One of the devices on either side of the link is part of a VCF setup
- The device(s) that are part of the VCF setup have their configuration mode set to **'auto-provisioned'** or **'pre-provisioned'**
- Interfaces on both end of the link are running LLDP
- Interfaces on both end of the link have not been configured as a VCP
- The device that is being added to the VCF is configured for fabric mode

When a network port is turned into a VCP, the change is permanent. A VCP is never converted (back) into a normal network interface automatically.

Another thing worth mentioning is that when multiple VCPs are active between two switches, they will automatically form a LAG.

Device modes:

Devices that participate in a VCF should be configured in **fabric mode**. By default, Juniper switches will be placed in virtual-chassis mode. To put a device in fabric mode, you will need to issue an operational command followed by a reboot. This can be done simultaneously by issuing the following command: **'request virtual-chassis mode fabric reboot'.**

To allow different types of QFX switches and EX4300 switches to become members of a VCF, you will need to run the devices in **mixed mode**. Keep in mind that only QFX5100 switches can serve as a spine node.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

LAG

Multiple physical links with matching properties (media, speed and duplex mode) can be combined into a single logical link. This logical link is called a Link Aggregation Group (LAG) and is referred to as an AE interface within JUNOS.



The use of link aggregation allows for increased redundancy and throughput. In general, JUNOS allows you to use the AE interfaces in the same way you would use a regular interface. Depending on the platform, an AE interface can be used for L2, L3, MPLS, etc.

A LAG can be created using the LACP protocol (IEEE 802.3ad) or through manual configuration. An advantage to using the LACP protocol is that the devices involved with the LAG must come to an agreement on the formation of the LAG. A statically configured LAG will form whenever a port is up, potentially black holing traffic.

Steps to configure LAG:

1: Configure the switch to support aggregate Ethernet interfaces:

```
set chassis aggregated-devices ethernet device-count 20
```

Note: numbering starts at 0, so when you want to configure **ae6**, you need to configure a **device-count** of **7**.

2: Assign physical interfaces to the LAG:

```
set interfaces xe-0/0/0 ether-options 802.3ad ae1
set interfaces xe-0/0/1 ether-options 802.3ad ae1
```

3: Configure the LAG, or AE interface itself. Parameters can include IPv4, IPv6, firewall filters, etc. The following is an example of an AE interface that is configured as a layer 2 trunk:

```
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 2000
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Note: LACP options (mode and minimum-links configuration for example) can be found under the **[edit interfaces ae3 aggregated-ether-options]** stanza.

MC-LAG

When devices on either side of a LAG suffer an outage, traffic is interrupted. When Virtual-Chassis (Fabric) technologies are used, a control plane issue in the master routing engine can impact the LAG. Even when child interfaces are spread out over several physical member switches.

A normal LAG lacks device redundancy and a potential risk that a VC(F) implies is that it is controlled by a single instance of JUNOS. This is where the Multi-Chassis Link Aggregation Group (MC-LAG), comes into play. As the name implies, MC-LAG enables you to construct a LAG using two different physical devices that run their own control plane operations.

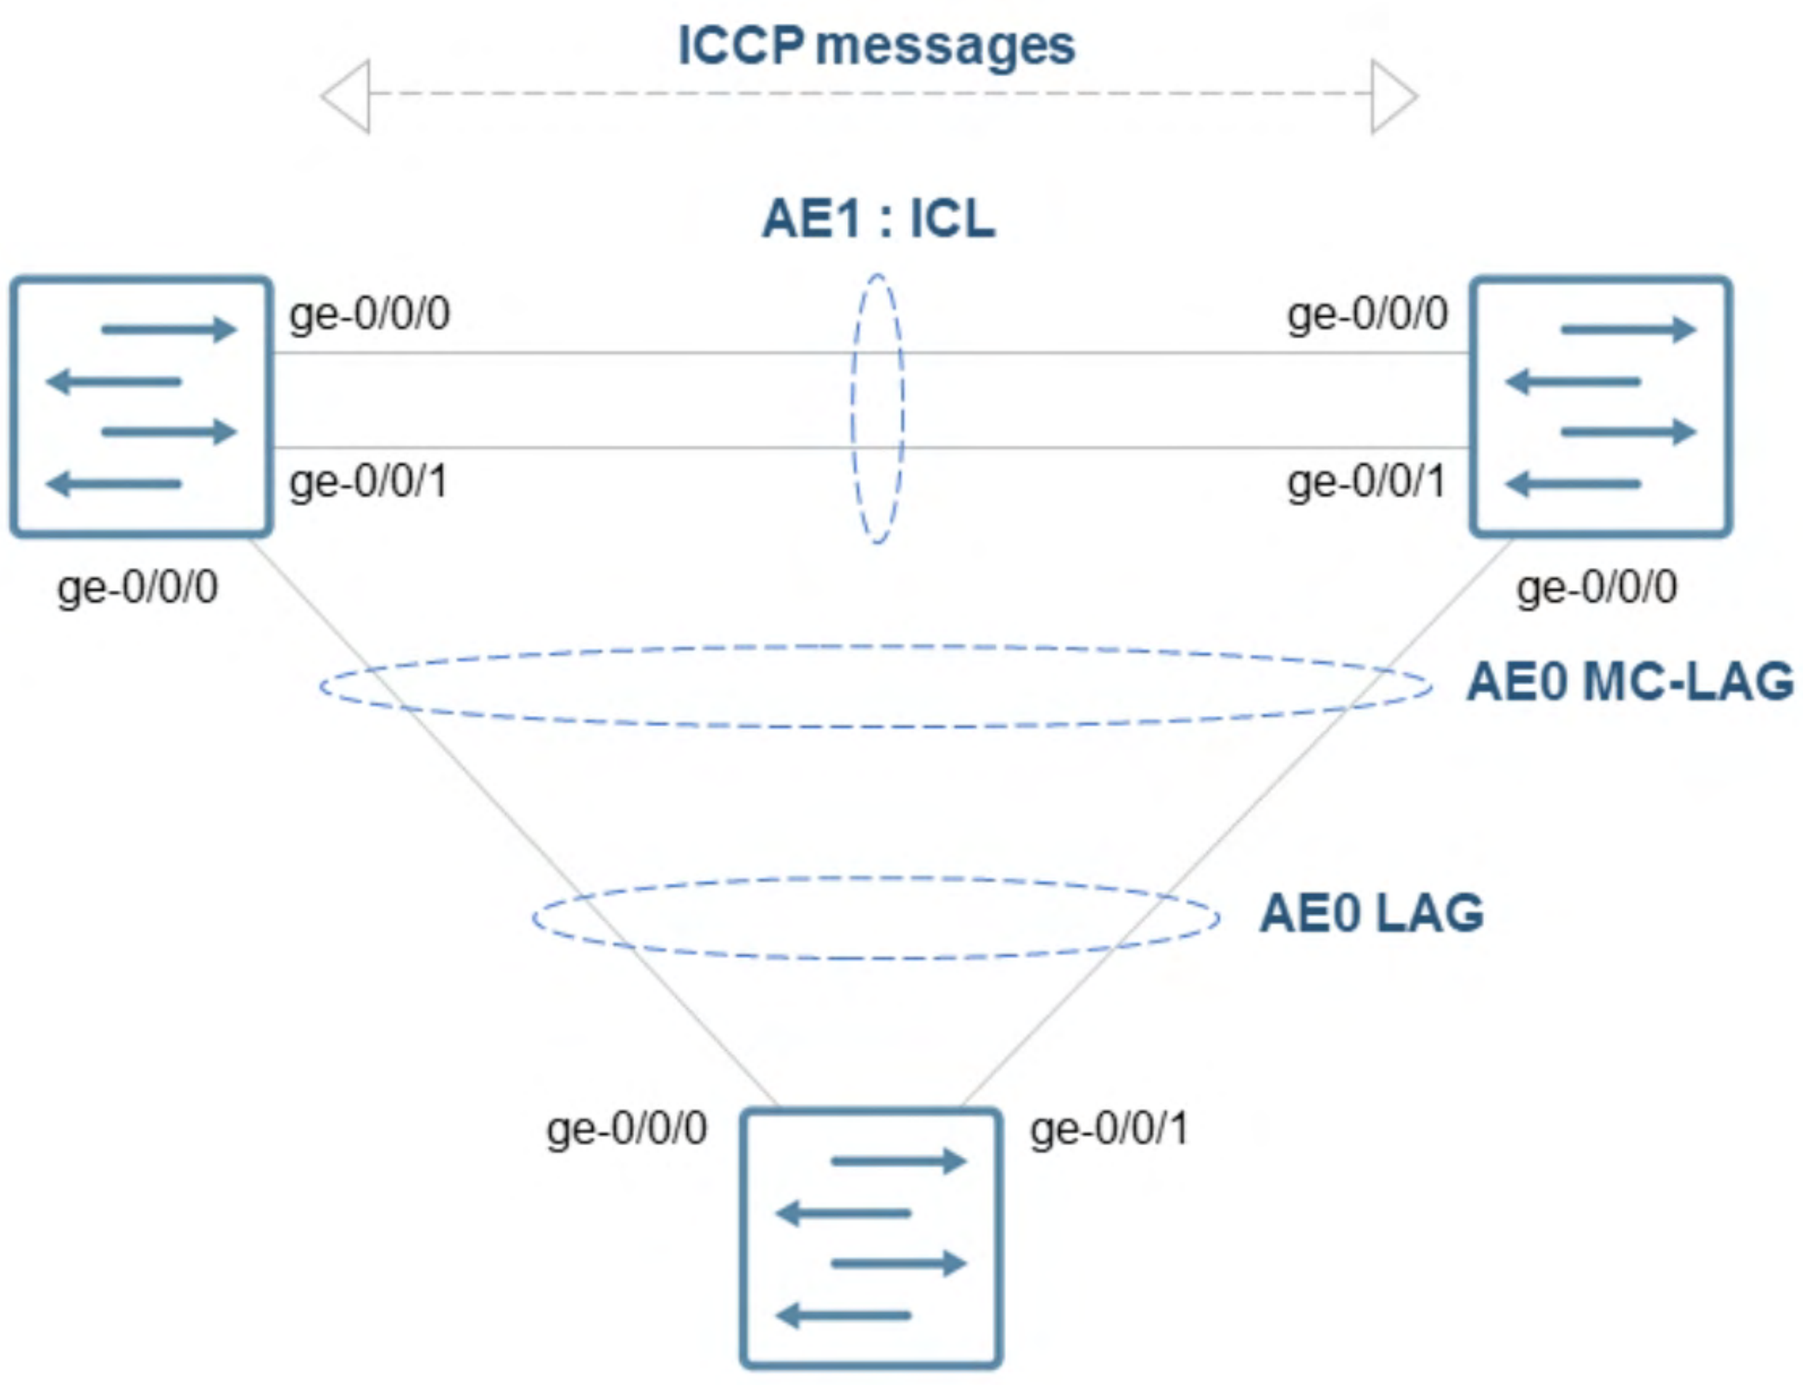

In the picture above, the top switches are collaborating to form an MC-LAG. To the bottom switch, the two MC-LAG peers appear as a single switch using 802.3ad.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

MC-LAG depends on the Inter-Chassis Control Protocol (ICCP) to replicate state between the different devices. Apart from replicating state information, the ICCP messages also exchange MC-LAG configuration parameters. When the protocol detects any faulty configuration, the MC-LAG will not be able to form. The ICCP protocol runs on top of TCP/IP and the ICCP messages are exchanged across the Inter-Chassis Link-Protection Link (ICL). For a more robust setup, the ICL itself can be configured as a LAG.

A fault-scenario for MC-LAG deployments is the 'split-brain' scenario. This can occur when the ICL link fails. To enable an MC-LAG setup to cope with such an event, 'backup-liveness-detection' can be used. A device can be configured to use an alternate link to determine if the ICCP peer is up or down. The OoB would be a good candidate link to consider for this purpose.

MC-LAG can operate in one of two modes: 'active-standby' or 'active-active'.

In an active-standby setup, only one of the MC-LAG peers is active. The active MC-LAG peer uses LACP to signal to the remote end that its link is ready to accept traffic. The inactive MC-LAG peer does the opposite. As soon as the active MC-LAG peer stops functioning, the inactive MC-LAG peer will signal to the remote end that it is ready to accept traffic.

In an active-active scenario, all the links are handling traffic. In some cases, this means that traffic needs to traverse both MC-LAG members. This traffic will be sent across the ICL-PL.

Apart from using the MC-LAG to switch traffic at layer 2, you can configure the MC-LAG peers to be the Layer 3 gateway for the devices connected to the MC-LAG. Two, mutually exclusive, options can facilitate this on Juniper devices. The first option is to configure both MC-LAG peers with VRRP and the second option is to use IRB MAC address synchronization.

When VRRP is used, both devices use VRRP to offer a virtual redundant gateway.

When IRB MAC address synchronization is used, both MC-LAG peers will install each other's MAC address in the forwarding table. When an MC-LAG peer experiences an outage, the peer device will continue to forward traffic that is destined to the MAC address of that MC-LAG peer. Devices connected to the MC-LAG peers can direct traffic towards either of the MAC gateway addresses. The MC-LAG peers will also synchronize ARP resolutions. This is done by replicating all ARP responses that it receives to the MC-LAG peer.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Steps to configure MC-LAG:

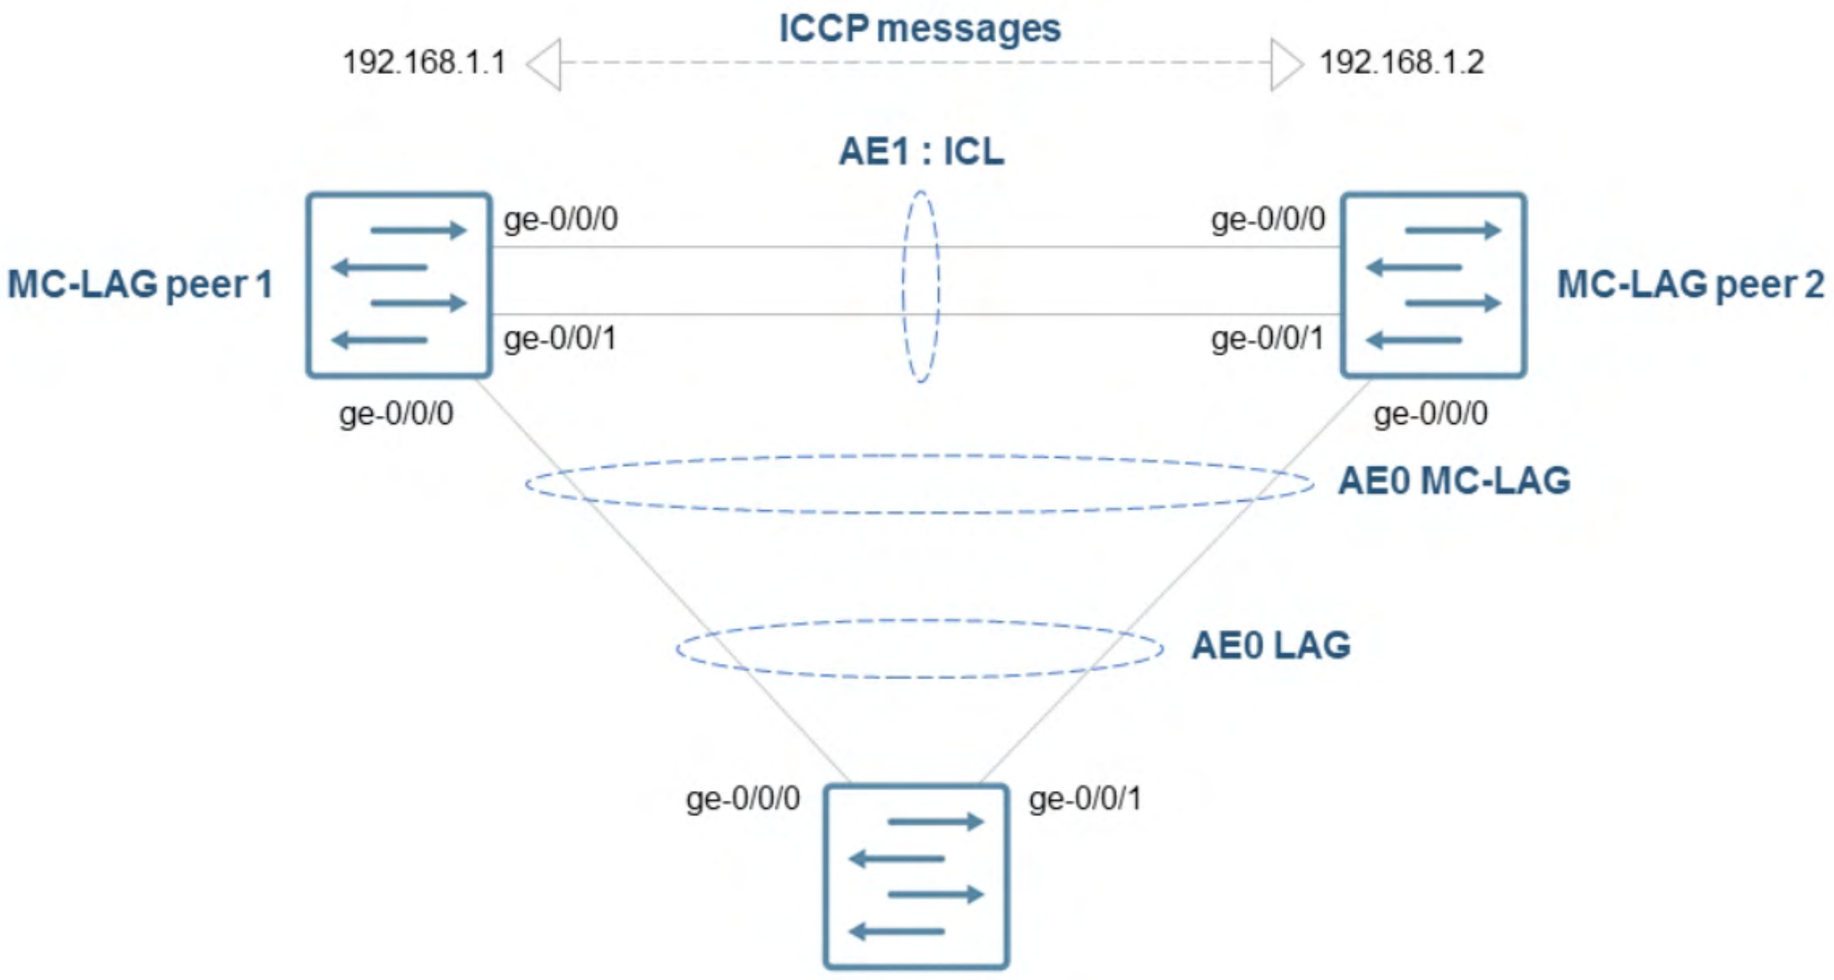

1: Configure both MC-LAG peers to support aggregate Ethernet interfaces:

```
set chassis aggregated-devices ethernet device-count 20
```

2: Assign a physical interface to the LAG on both MC-LAG peers:

```
set interfaces ge-0/0/0 ether-options 802.3ad ae1
```

2: The AE that will be multi-chassis needs to be configured with the same MC-AE identification number on both MC-LAG peers:

```
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
```

3: Specify a unique chassis ID on both MC-LAG peers, usable values are '0' or '1':

MC-LAG peer 1:

```
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
```

MC-LAG peer 2:

```
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
```

4: Specify the same MC-LAG mode on both MC-LAG peers:

```
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

5: Configure the mc-ae status control on both MC-LAG peers. When the ICL fails, one MC-LAG peer should become active and the other should become standby:

MC-LAG peer 1:

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
```

MC-LAG peer 2:

```
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
```

6: Activate LACP and specify the same LACP system ID on both MC-LAG peers:

```
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:01
```

Note: by configuring the same LACP ID on both switches, the remote end will see both MC-LAG participants as a single device.

7: Specify the same LACP admin-key number on both MC-LAG peers:

```
set interfaces ae0 aggregated-ether-options lacp admin-key 4
```

8: Configure ICCP on both MC-LAG peers:

MC-LAG peer 1:

```
set protocols iccp local-ip-addr 192.168.1.1
set protocols iccp peer 192.168.1.2 backup-liveness-detection backup-peer-ip 10.0.0.1
set protocols iccp peer 192.168.1.2 liveness-detection minimum-interval 1000
```

MC-LAG peer 2:

```
set protocols iccp local-ip-addr 192.168.1.2
set protocols iccp peer 192.168.1.1 backup-liveness-detection backup-peer-ip 10.0.0.2
set protocols iccp peer 192.168.1.1 liveness-detection minimum-interval 1000
```

The 'backup-liveness-detection backup-peer-ip' is optional. It can be configured to have the MC-LAG peers check whether or not the MC-LAG peer is up or down by exchanging keepalive messages over another (OoB) link.

9: Specify a service ID on both MC-LAG peers:

```
set switch-options service-id 1
```

10: Configure multi-chassis protection link between the MC-LAG peers:

MC-LAG peer 1:

```
set multi-chassis multi-chassis-protection 192.168.1.2 interface ae1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

MC-LAG peer 2:

```
set multi-chassis multi-chassis-protection 192.168.1.1 interface ae1
```

With this configuration command, you reference the remote peer as well as the physical link that will be used to reach the remote peer.

11. On both MC-LAG peers, configure the MC-AE interface with the desired characteristics:

```
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 2000
```

Note: At this moment, the QFX5100 only supports active-active mode. The previous configuration example was done on QFX. This configuration differs slightly from the configuration on MX.

Additional tips:

- Do not forget to specify the child interfaces that belong to the MC-AE interface
- The MC-LAG will not become active if there is no unit configuration for the AE
- When the interface of an active-active MC-LAG is enabled for Ethernet switching, the VLAN needs to be added to the ICL as well
- You can reuse most settings for a second MC-AE between MC-LAG peers. Only the system ID, admin key and mc-ae-id need to differ on a per MC-AE basis
- Certain MC-LAG deployments involving active routing-protocols between the MC-LAG peers may require static ARP entries
- If the remote end does not support LACP, use the '**force-up**' option to make sure that the MC-AE interfaces become operational:

```
set interfaces xe-0/0/12 ether-options 802.3ad ae0 lacp force-up
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: MC-LAG

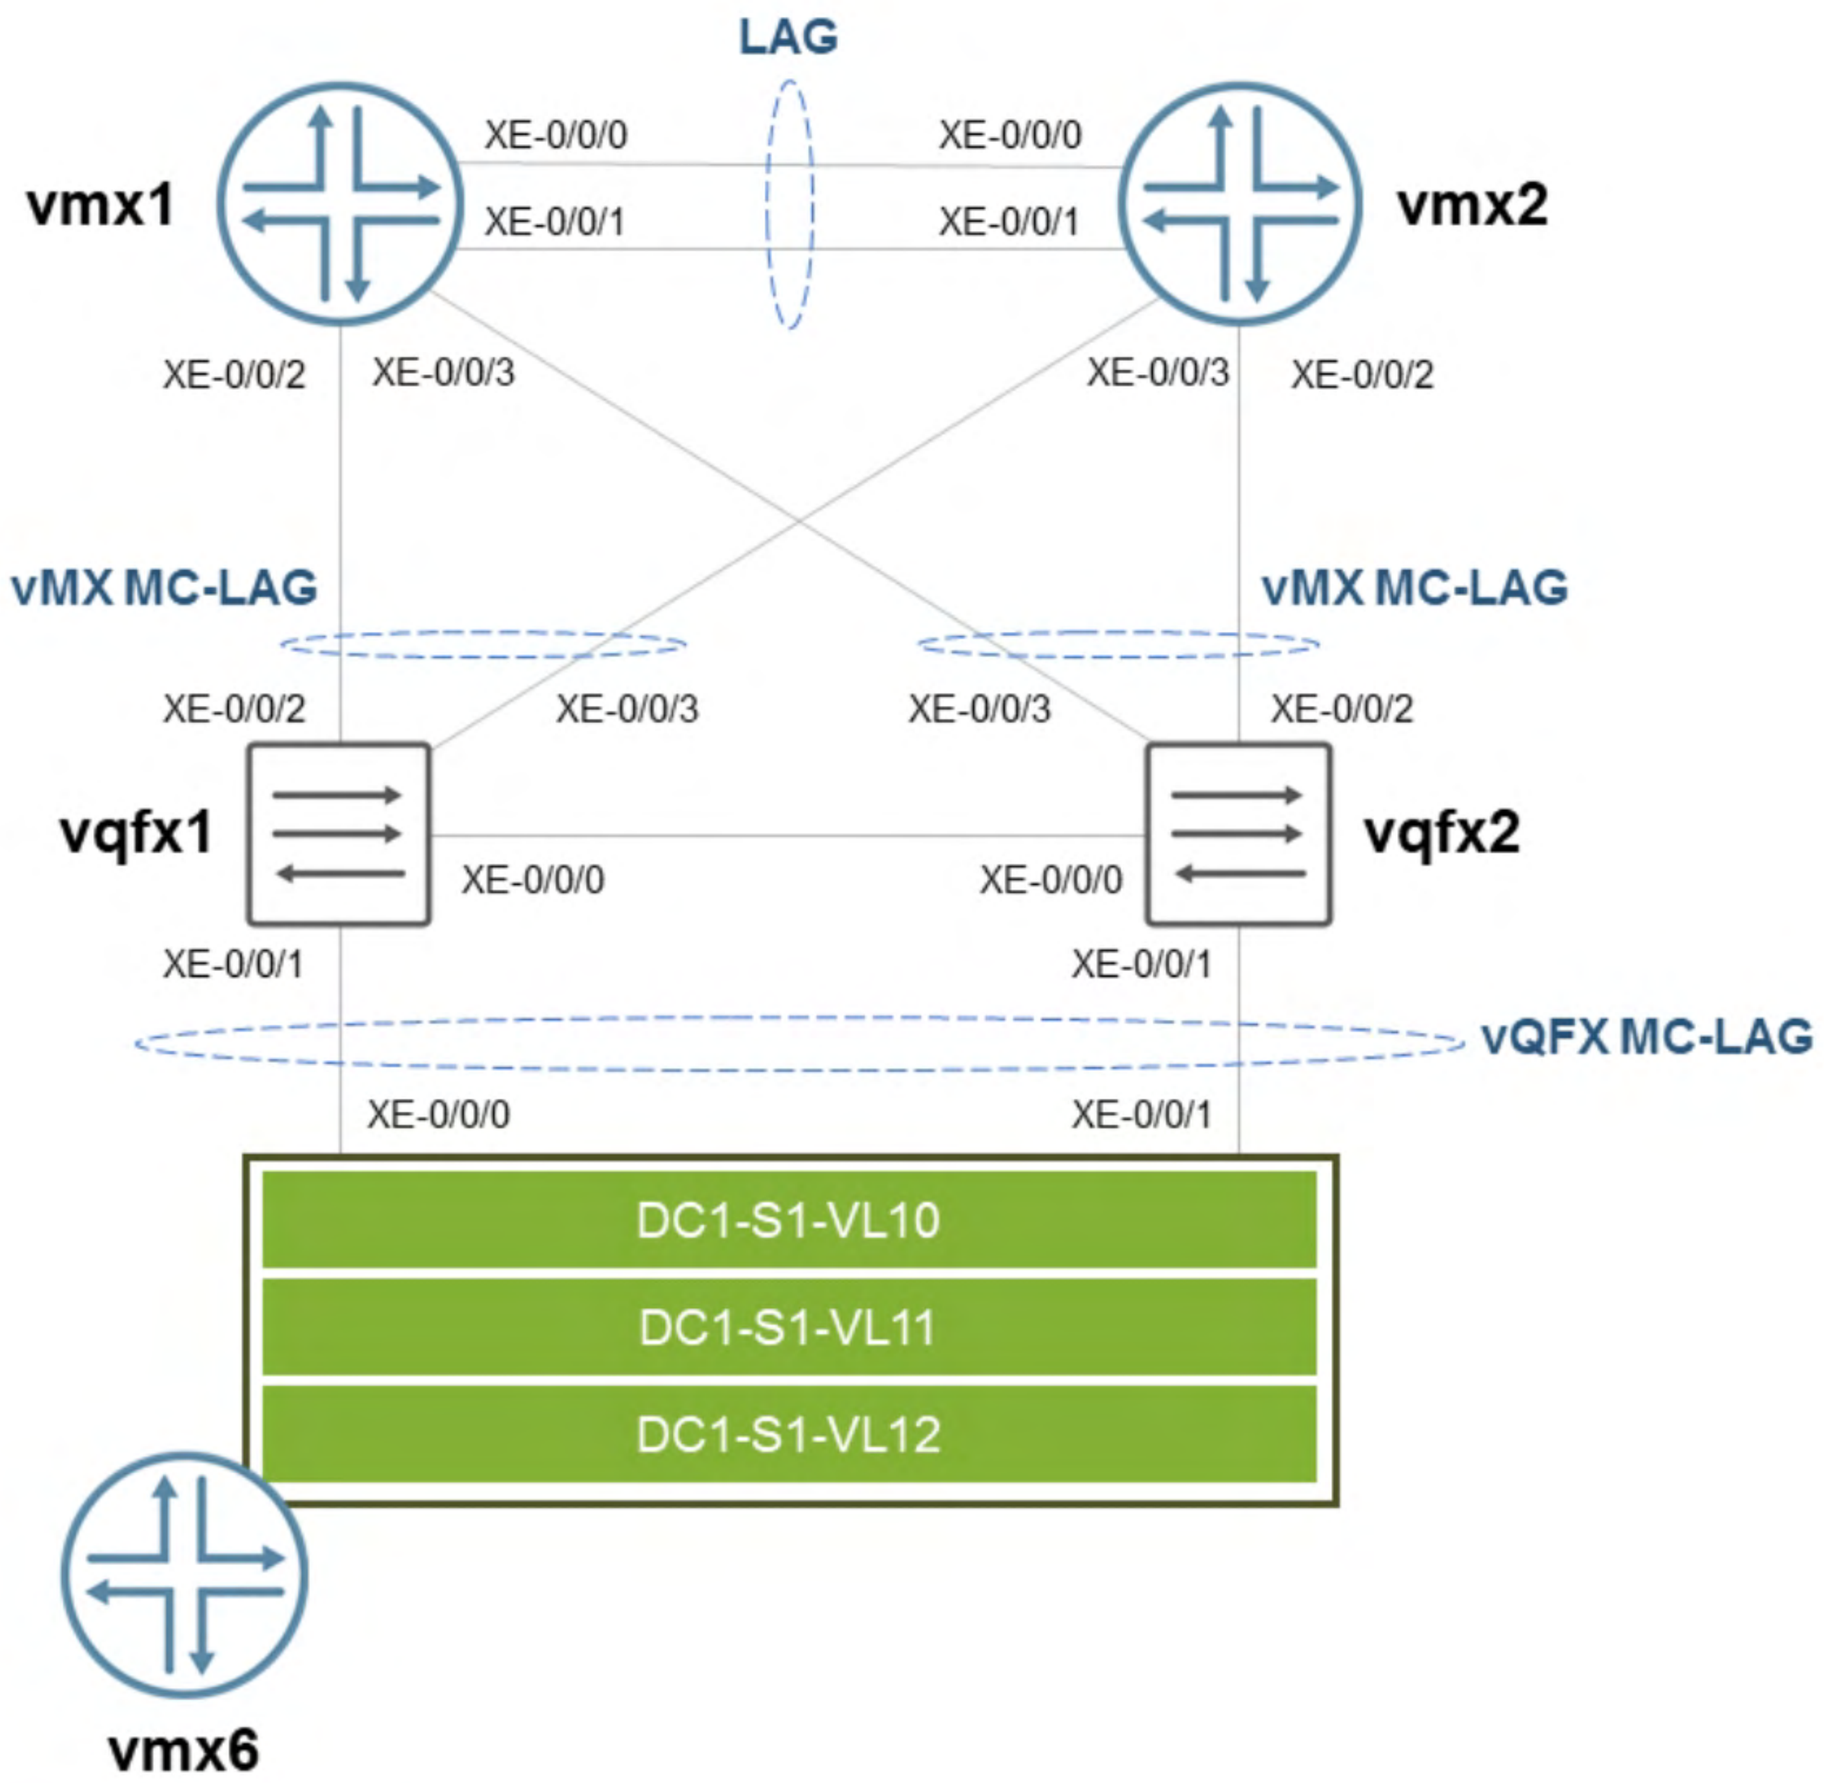

Figure 1 - Physical topology

Note: starting topology is configured with IP-addressing only.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.1: ToR Configuration

- Configure VLAN 100 on interface xe-0/0/0 on both vQFX1 as well as vQFX2.
- Configure an IRB interface for the VLAN. Use the following IP addresses:

vQFX1	*192.168.100.1/30*
vQFX2	*192.168.100.2/30*

- Traffic between the switches should be tagged.

Task 1.2: MC-LAG Peers: Server-Facing Configuration

- Configure vQFX1 and vQFX2 as MC-LAG peers.
- Use IRB 100 to source the ICCP session. Use 'jncie' as the MD5 authentication key.
- Set the mode to ‘active-active’ and determine the other configuration parameters yourself.
- Use the OoB network for additional checks between the MC-LAG peers, enabling the network to handle split-brain scenarios.
- During a split brain, vQFX2 should become the inactive MC-LAG peer.
- Enable the MC-LAG interface for VLANs 10, 11 and 12.

Task 1.3: MC-LAG Peers: Core-Facing Configuration

- Configure an AE uplink towards the vMX routers.
- Use interfaces xe-0/0/2 and xe-0/0/3 on both vQFX devices.
- Enable the uplink for the server VLANs.
- Make sure that communications inside the server VLANs can tolerate the loss of an uplink.

Task 1.4: vMX LAG Configuration

- Configure VLAN 101 on both vMX1 as well as vMX2.
- Configure an IRB interface for the VLAN. Use the following IP addresses:

vMX1	192.168.101.1/30
vMX2	192.168.101.2/30

- Traffic between the routers should be tagged.
- The links between the vMX routers should be bundled into a LAG.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.5: vMX MC-LAG Configuration

- Configure MC-LAGs on the vMX routers.
- The MC-LAGs should connect to the LAGs that were configured in task 1.3.
- Use the OoB network for additional checks between the MC-LAG peers.
- During a split brain, vMX2 should become the inactive MC-LAG peer.
- Enable the server VLANs on the MC-AE interface.

Task 1.6: vMX MC-LAG Gateway Configuration

- Configure the following gateways on the vMX routers:

IRB / VLAN	Gateway address
IRB10	10.0.10.254/24
IRB11	10.0.11.254/24
IRB12	10.0.12.254/24

- Use VRRP for VLAN 10 and VLAN 11.
- The vMX1 router should be VRRP master for VLAN 10 and the vMX2 router should be the master for VLAN 11.
- Make sure that you can ping the virtual-router.
- Authenticate the messages using password 'jncie'.
- Use an anycast IRB for VLAN 12.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: Multistage MC-LAG

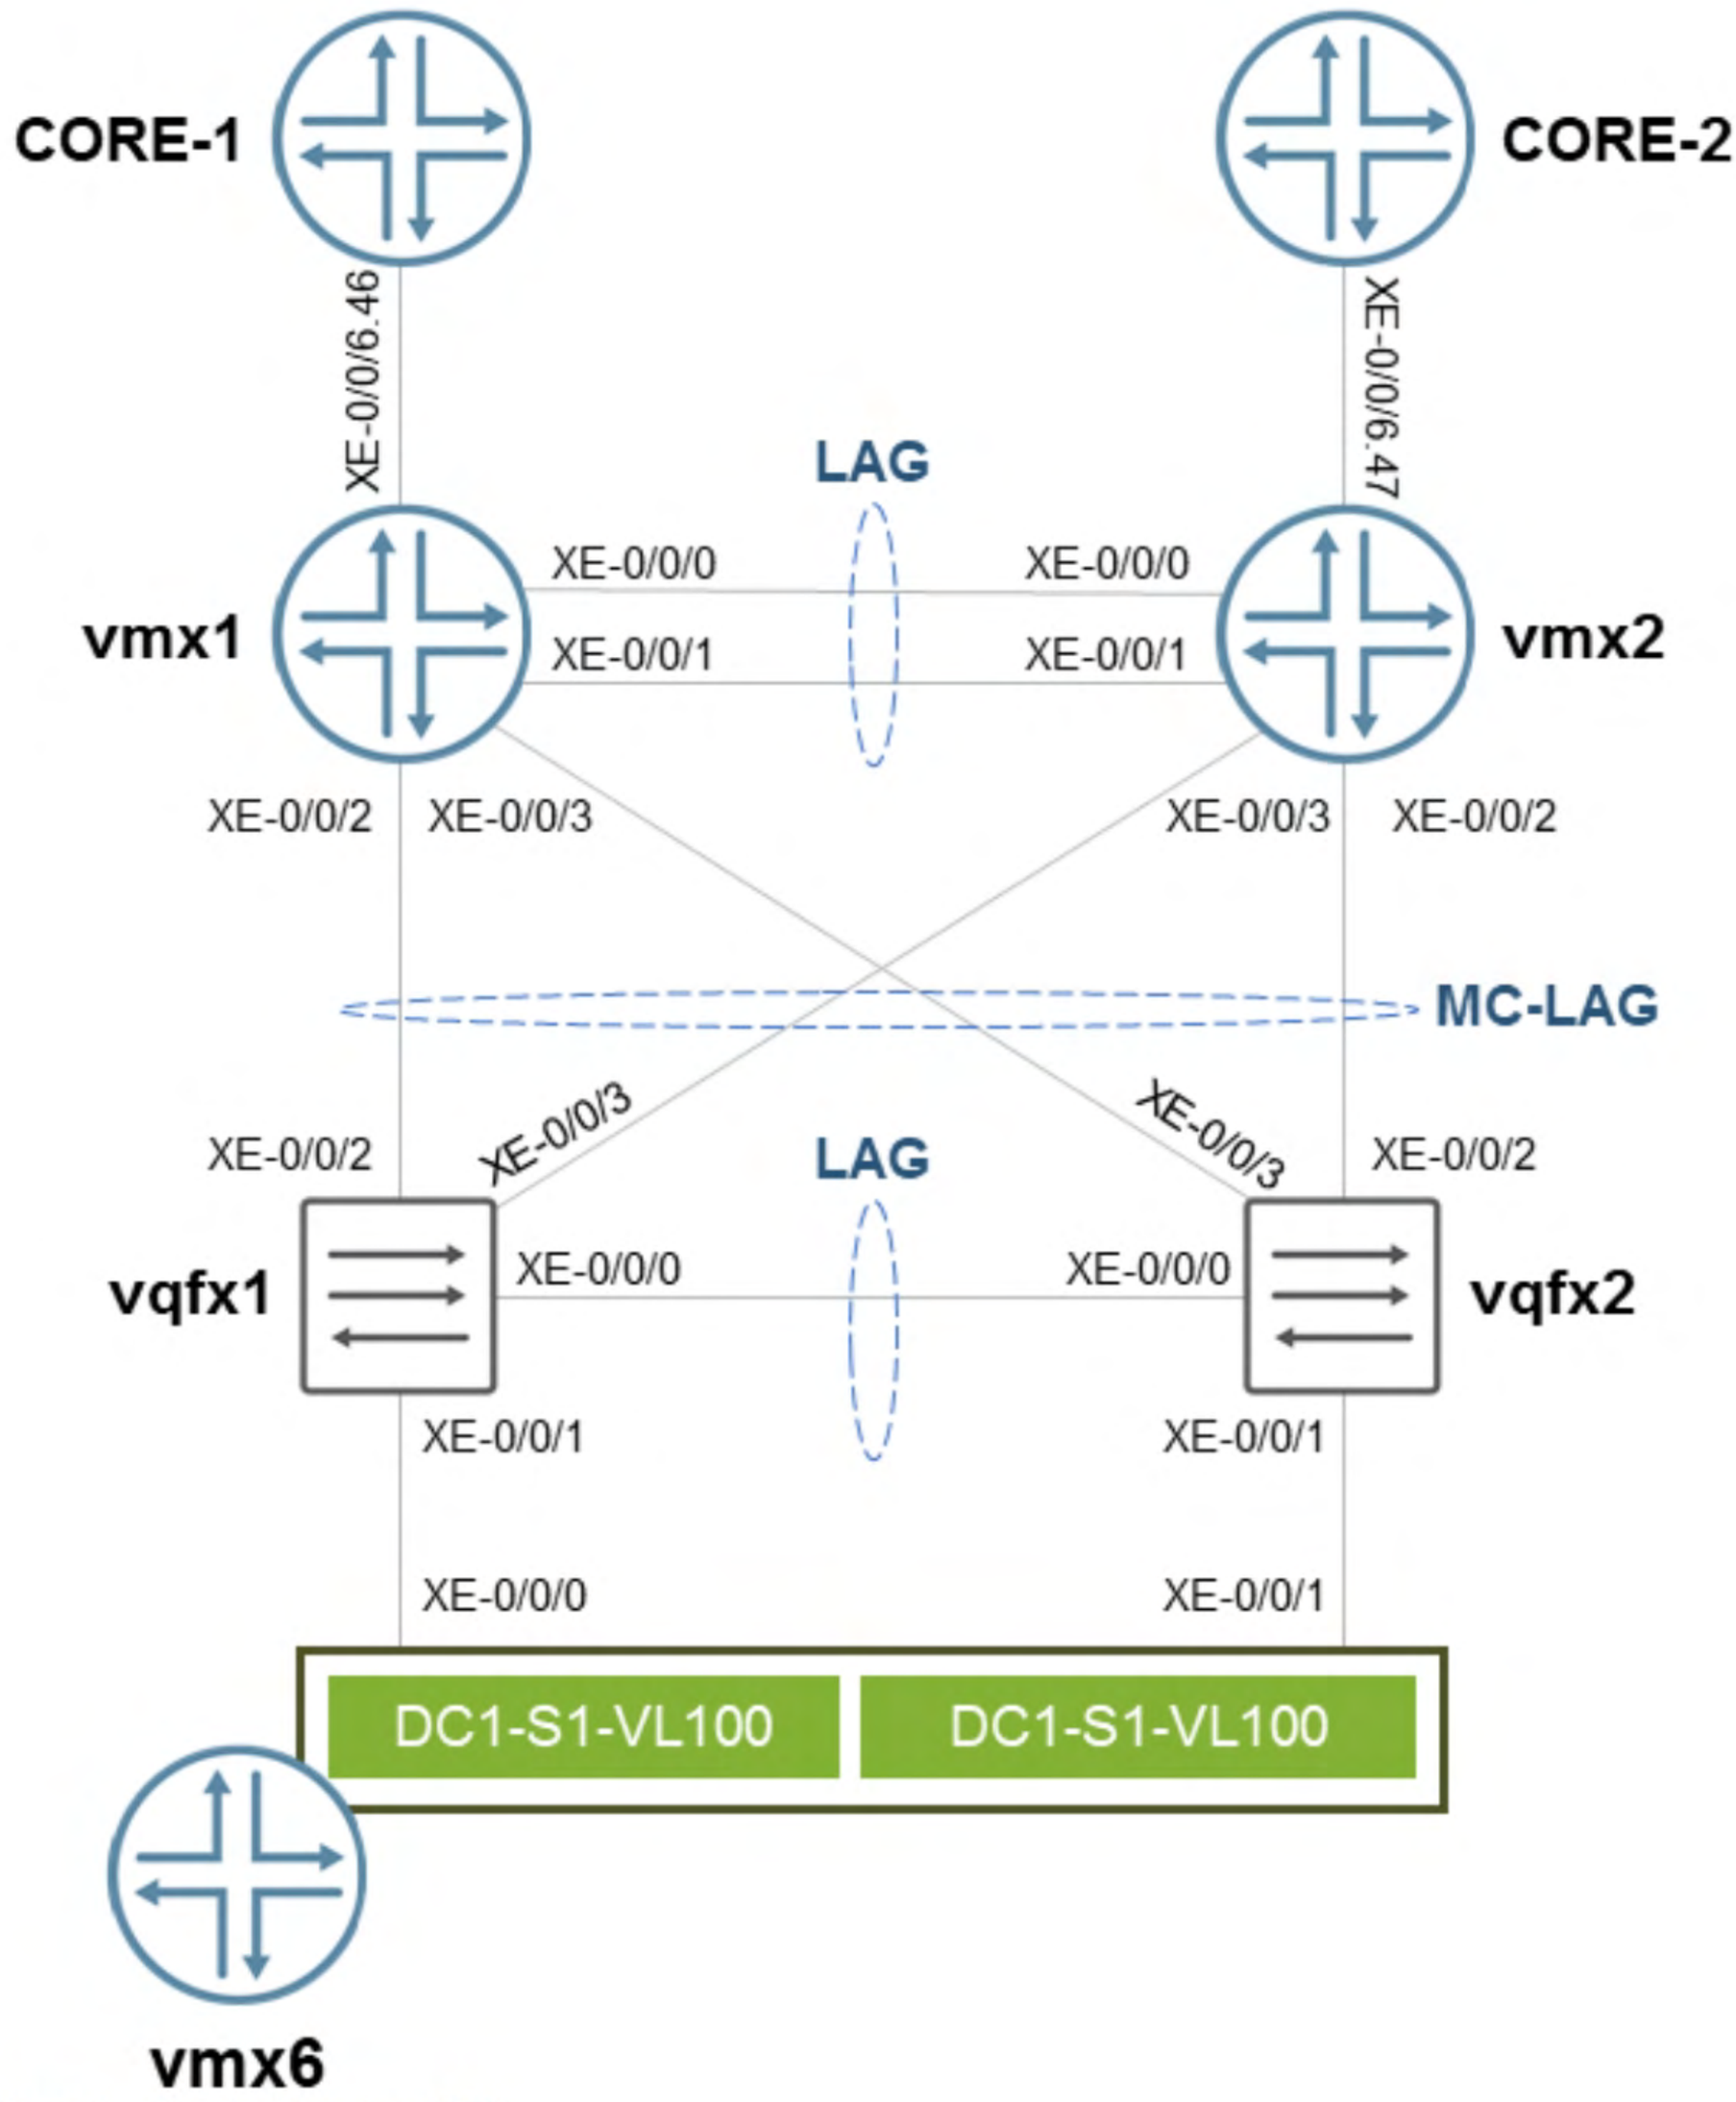

Figure 2 - Physical topology DC1

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

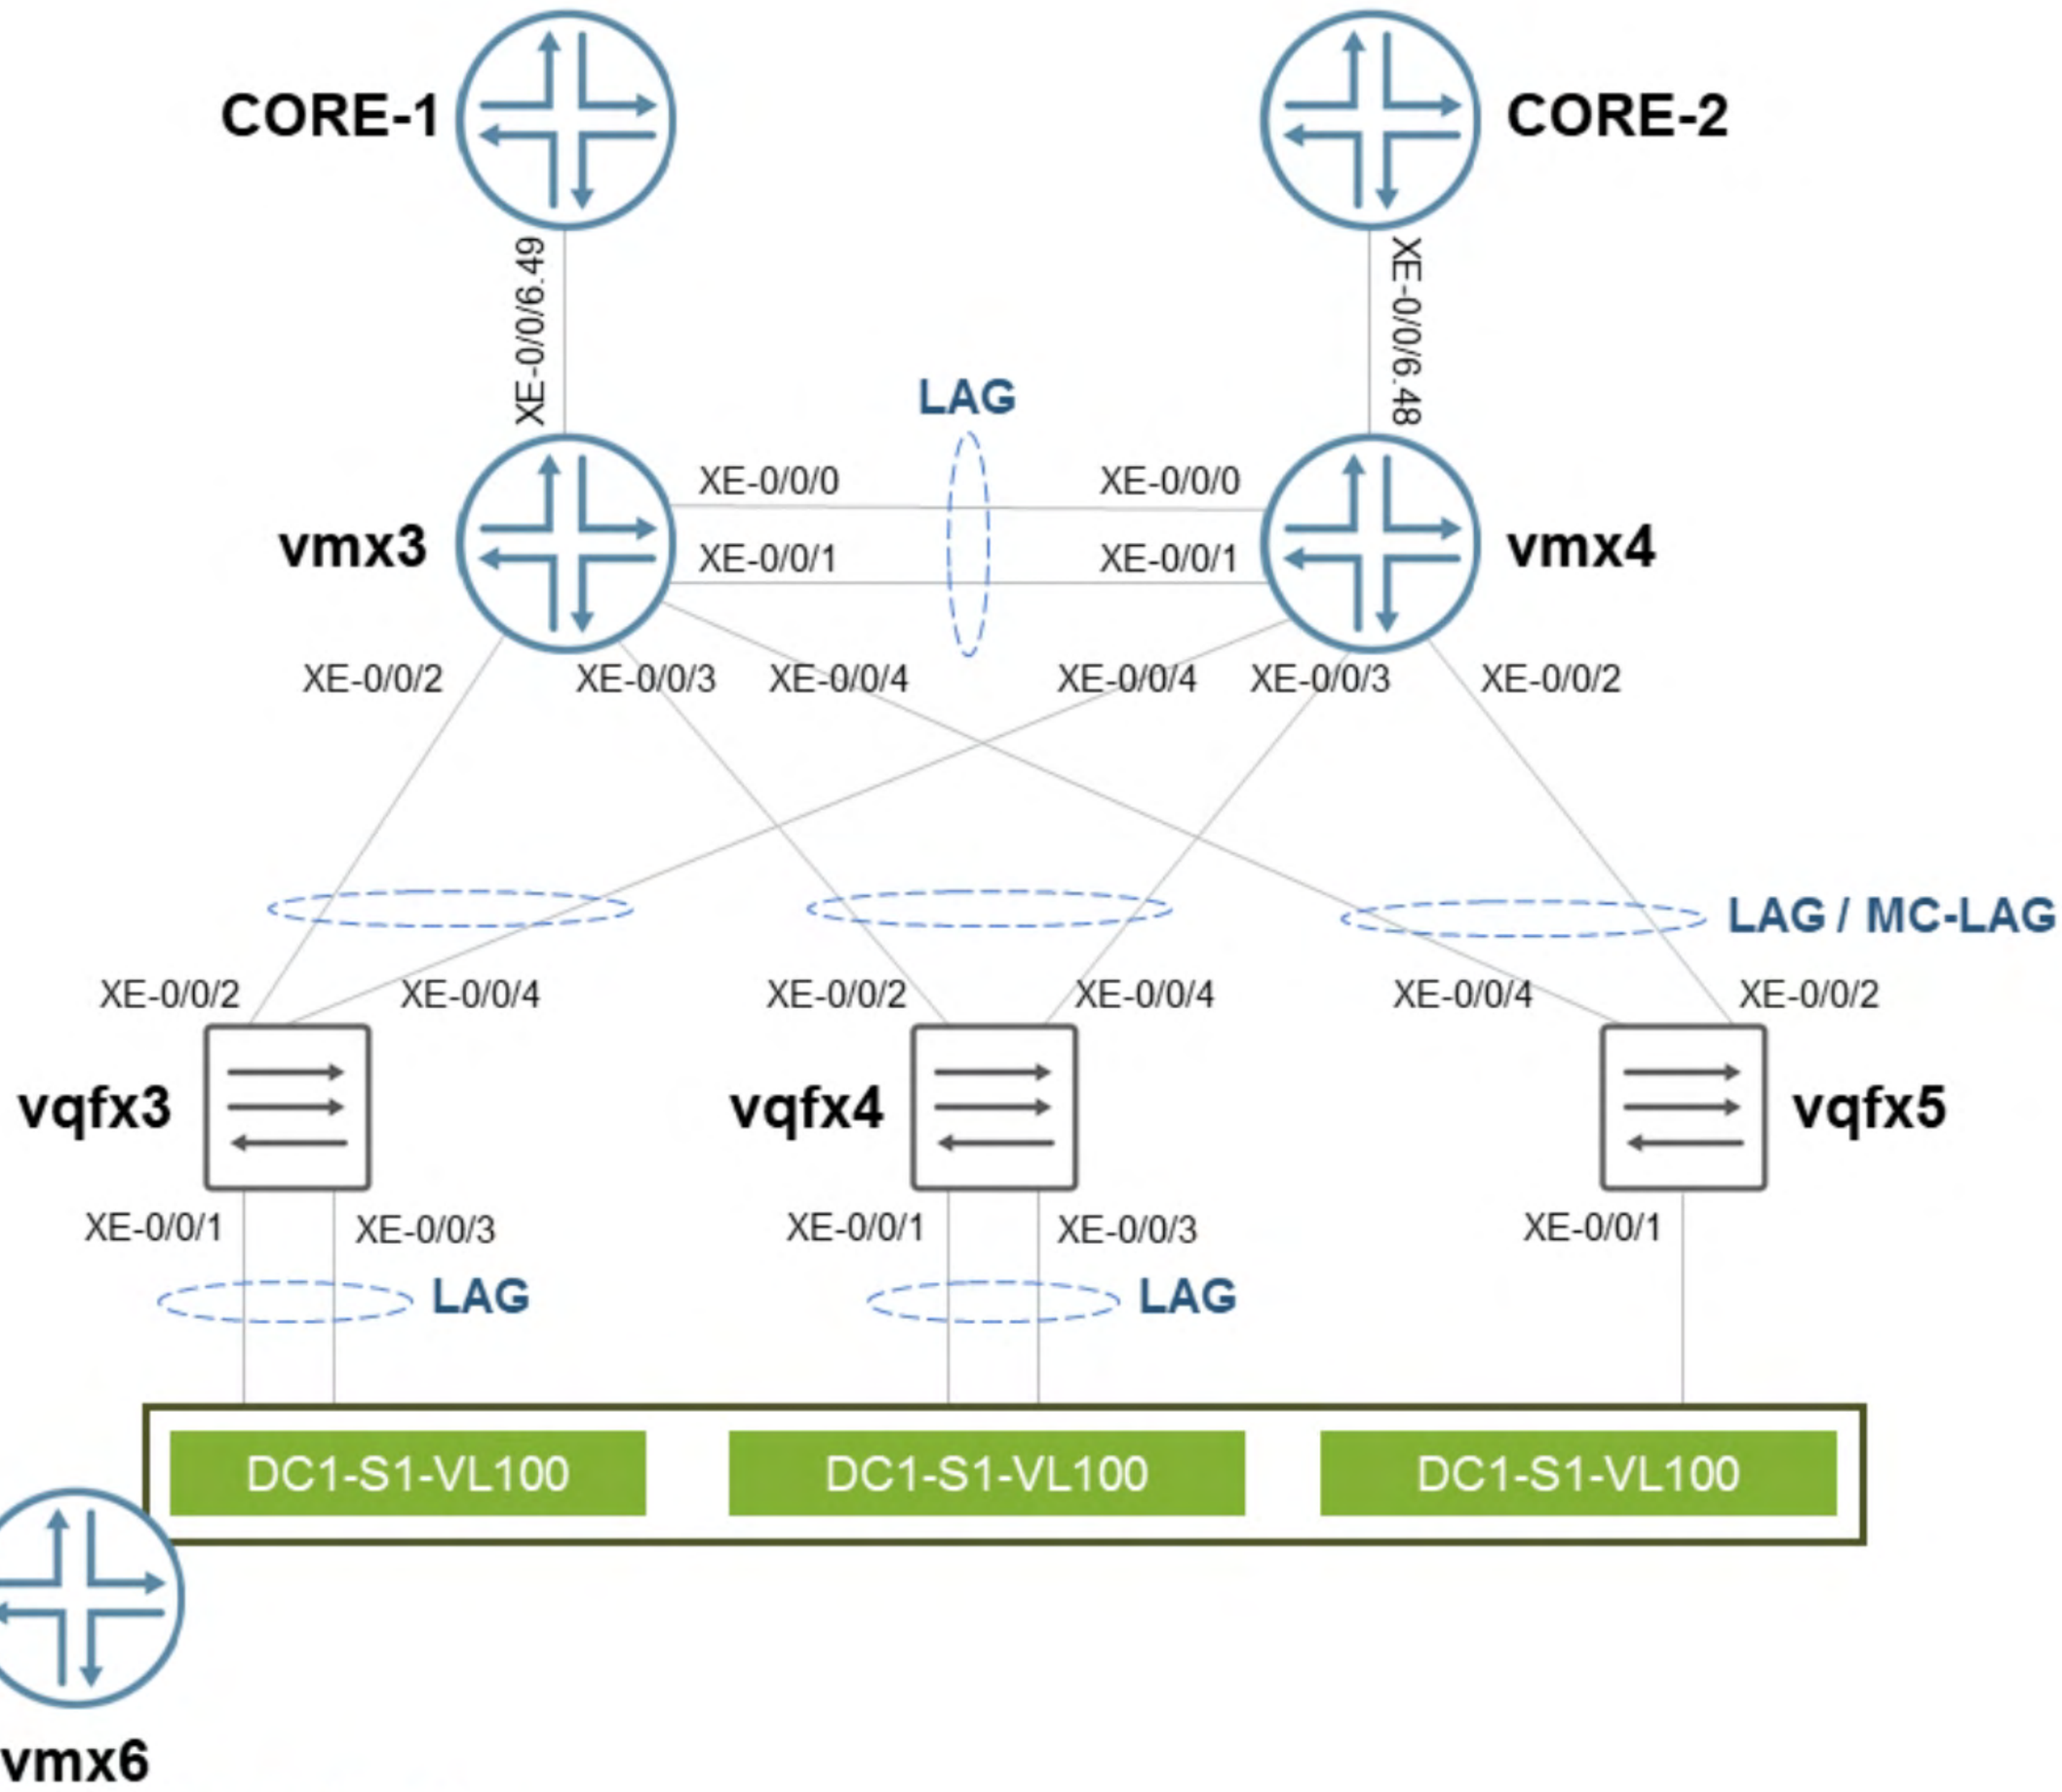

Figure 2 - Physical topology DC2

Note: starting topology is configured with IP-addressing only.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.7: DC1 LAG

- Configure a LAG between both vMX1 and vMX2 as well as between vQFX1 and vQFX2.
- Configure an IRB interface for the VLAN. Use the following IP addresses:

Device	IP address	IPv6	VLAN
vQFX1	192.168.44.1/30		44
vQFX2	192.168.44.2/30		44
vMX1	192.168.45.0/31	Link-Local	45
vMX2	192.168.45.1/31	Link-Local	45

- Traffic between the devices should be tagged.

Task 1.8: DC1 Multistage MC-LAG

- The interfaces between the vMX and vQFX devices in DC1 should be part of an MC-LAG interface.
- Use the IP addresses that were configured in the previous task to source ICCP sessions.
- The vMX routers should only forward traffic to the vQFX devices when there is a minimum of 20Gbit is available towards the vQFX.
- Enable the MC-LAG and LAG interfaces for VLAN 100.
- Enable the server-facing interfaces for VLAN 100.
- All interfaces should be configured as trunk links.

Task 1.9: DC1 Gateway and OSPF Configuration

- Configure 192.168.1.254/24 as the IPv4 gateway address on the vMX devices for VLAN 100.
- Configure 2001:100:0:100::254 as the IPv6 default gateway for VLAN 100.
- For IPv4, configure an anycast gateway and for IPv6 configure VRRP. Allow hosts to ping the IPv6 default gateway.
- Enable the vMX uplinks for OSPF and OSPFv3 and advertise all IRB and lo0 interfaces.
- Ensure that the loss of a single link between the vQFX and vQFX devices will not impact reachability.
- OSPF links carrying neighbor relationships should be configured as point-to-point. It should not be possible to establish OSPF neighbor adjacencies across server facing interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.10: DC1 Security

- Configure storm control to drop all BUM traffic over 5Mbit on the server-facing links.
- Allow only 1Gbit of traffic for VLAN 100 on the vQFX devices. Ensure that VLANs added in the future will not have their bandwidth restricted.
- Allow for a maximum of 20 MAC addresses to be learned in VLAN 100. Traffic for additional hosts should be dropped.
- Limit the movement of MAC addresses to 2 moves per second.
- Cap the cumulative of IPv4 and IPv6 traffic for VLAN 100 on the vMX routers to 1Gbit.

Task 1.11: DC2 LAG and MC-LAG Configuration

- Configure the vQFX uplinks towards the vMX devices as LAG interfaces.
- Configure the downlinks from the vMX towards the vQFX as MC-LAGs.
- Enable the LAG and MC-LAG interfaces for VLAN 100.
- Enable the server-facing interfaces on the vQFX for VLAN 100.
- All links should transport tagged traffic.
- On vQFX3 and vQFX4, configure the server-facing interfaces as a LAG interface.

Task 1.12: DC2 Gateway, VRRP and OSPF Configuration

- Configure the following gateways on the vMX routers:

IRB / VLAN	**Gateway address IPv4**	**Gateway address IPv6**
IRB100	192.168.100.254	2001:100:100:100::254

- Use VRRP and ensure that vMX3 is the active router.
- Make sure that you can ping the virtual-router and authenticate the VRRP-message for IPv4 with the MD5 hash 'jncie'.
- Ensure that vMX3 and vMX4 advertise the IRB as well as the loopback interface to the core.
- It should not be possible to form OSPF neighbor relationships across IRB 100.
- OSPF links carrying neighbor relationships should be configured as point-to-point. Use OSPFv2 for IPv4 and OSPFv3 for IPv6.

Task 1.13: MC-LAG Verification

- Perform ping from routing-instance DC1-S1-VL100 towards 192.168.100.1. Along the path of the packet, find:
    - all relevant destination MAC entries.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

  - the ARP entries that exist as a result of this traffic.
  - the routing table entries required to get the packet from DC1 to DC2.
  - the interfaces that the traffic traverses.

- Perform a ping from routing-instance DC2-S3-VL100 towards 2001:100:0:100::1. Along the path of the packet, find:

  - all relevant destination MAC entries.
  - the NDP entries that exist as a result of this traffic.
  - the routing table entries required to get the packet from DC1 to DC2.
  - the interfaces that the traffic traverses.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 3: VCF

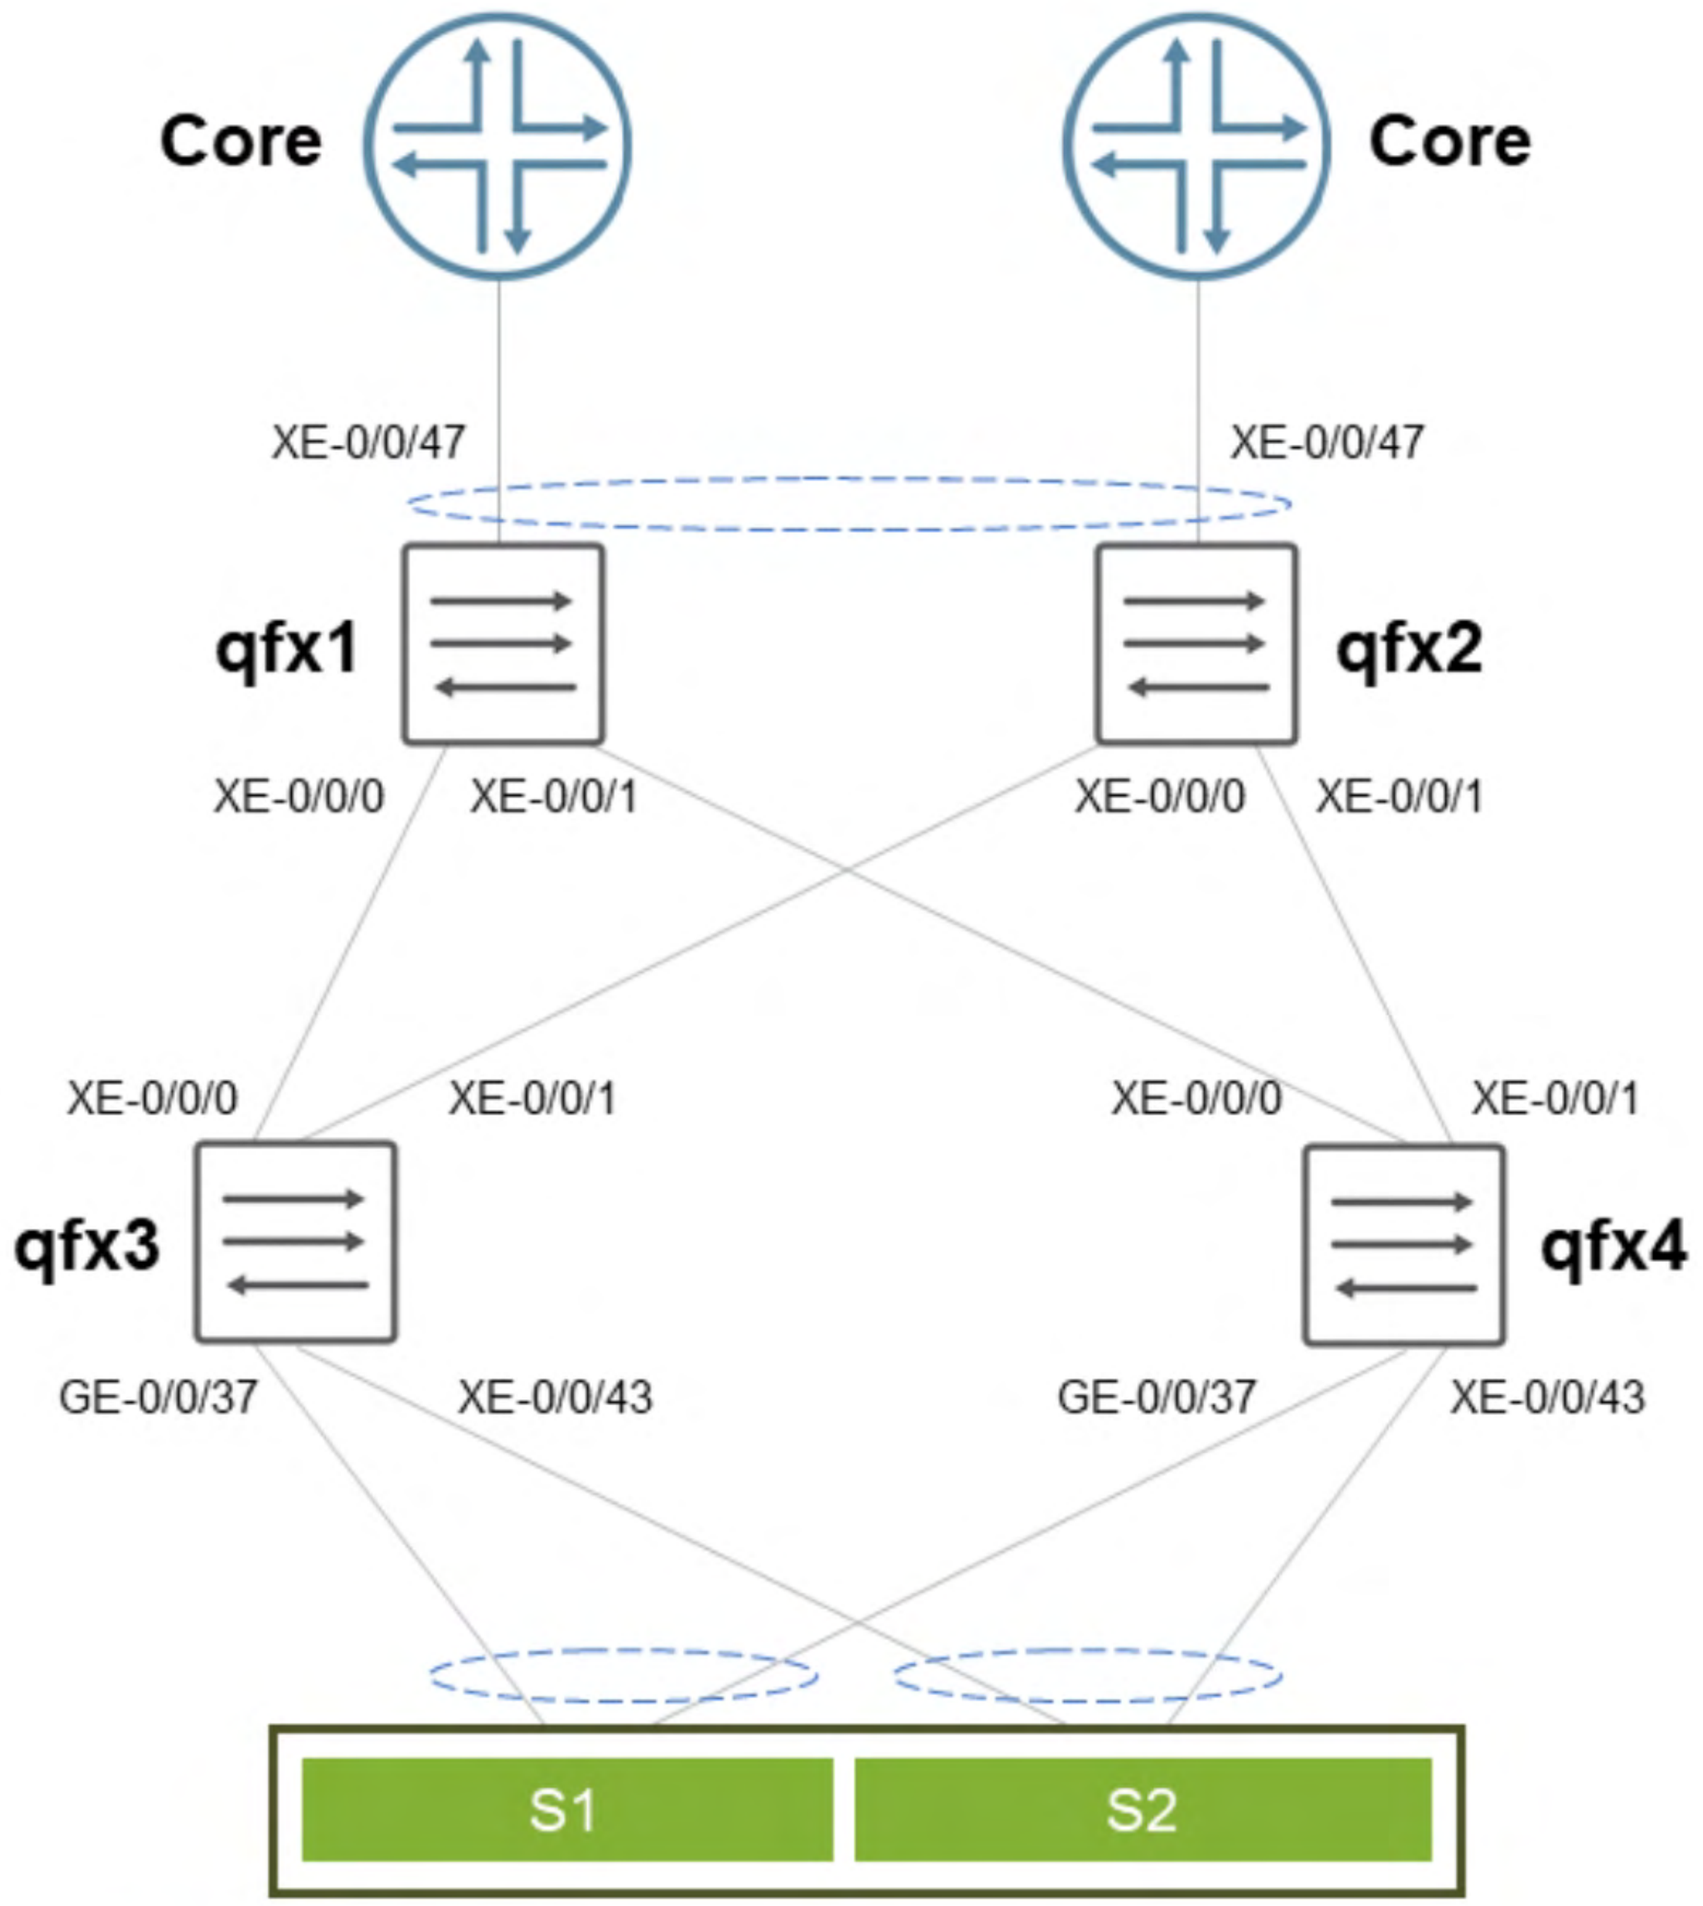

Figure 3- Physical topology

Note: This assignment is not included in the Juniper Networks' rack rental topology. The rack rental only includes vMX and vQFX devices, and unfortunately, vQFX does not support VCF.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Console access to every switch is required for these assignments.

Revert the QFX switches to factory default before you start this lab.

Task 1.14: VCF Configuration

- Configure the QFX switches to form a single VCF with the hostname 'jncie-vcf'.
- QFX1 should be the active master switch and QFX2 should be the backup switch.
- Configuration must determine the role of every device.

Task 1.15: VCF Features

- When the master fails, the backup should continue forwarding layer 3 as well as layer 2 traffic immediately.
- The MDT should be rooted at the spine node.
- The root user should be able to manage the VCF on 10.10.20.120/24 using SSH.

Task 1.16: Deconstructing the VCF

- Deconstruct the VCF.
- Return all switches to the default configuration and mode.
- No VC or VCF should be able to form.

Task 1.17: Autoprovisioning a VCF

- Configure an autoprovisioned VCF with hostname 'jncie-vcf'
- The qfx1 and qfx2 devices should operate as spines.
- Enable the VCF for future additions of EX4300 switches.
- The root user should be able to connect to the device on 10.10.20.120/24 using SSH.
- Enable GRES, NSB and NSR.

Task 1.18: VCF VLAN and Interface Configuration

- Enable all core and server-links for VLAN 20 and VLAN 30.
- VLAN 20 should be the native VLAN on every link.
- Ensure that all current as well as future VLANs are limited to 500 MAC addresses. The packets that are violating this setting should be dropped.
- None of the VLANs should not allow for MAC flapping to exceed 2 per second. Packets that are violating this setting should be dropped.
- The default VLAN should not be allowed to forward any packets.
- Use the following interface numbering:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

S1	ae18
S2	ae19
core	ae20

Chapter 2: Clos IP Fabric

The three stage Clos Layer 3 IP Fabric is becoming a popular design choice for the modern datacenter. The general topology of a three stage Clos IP-Fabric network is as follows:

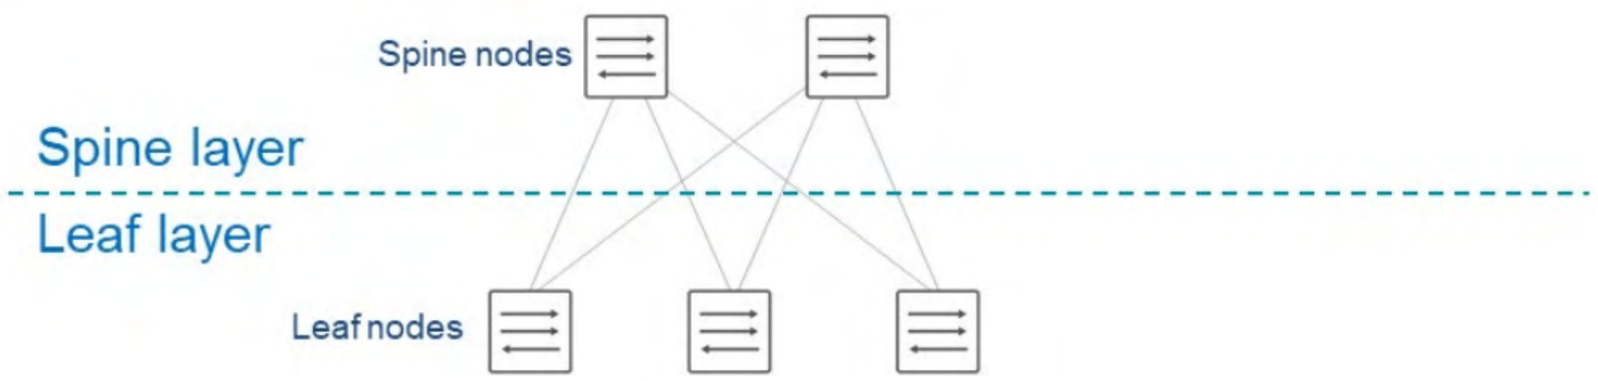

There are several characteristics to a Clos IP-Fabric that can explain this increase in popularity. First of all, the IP-Fabric design is extremely scalable and allows a network to grow into the tens of thousands of ports. Because an IP-fabric scales-out linearly by adjusting the width of the spine, the growth is something that is easily managed. Furthermore, unlike with a VCF for example, an IP-fabric does not rely on any proprietary protocol that imposes a vendor lock-in. The devices that are a part of the IP-fabric only need to support a common routing protocol. This offers a tremendous amount of flexibility because it means that the IP-fabric can be designed with different devices from multiple vendors. In addition to all the aforementioned examples, an IP-fabric can also serve as an underlay network for an EVPN/VXLAN overlay.

The IP-Fabric can be build using EBGP or by using IBGP together with an IGP. Both strategies allow for ECMP between the different members of the IP-Fabric. When EBGP is chosen as a routing protocol within an IP-fabric, it will be able to scale the most. An additional benefit to choosing EBGP is the possibility to use policies to tag, steer and control traffic flows over the IP-Fabric.

IP-Fabric using EBGP.

When an IP-Fabric is designed using EBGP, every node is configured with its own (usually private) autonomous system:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

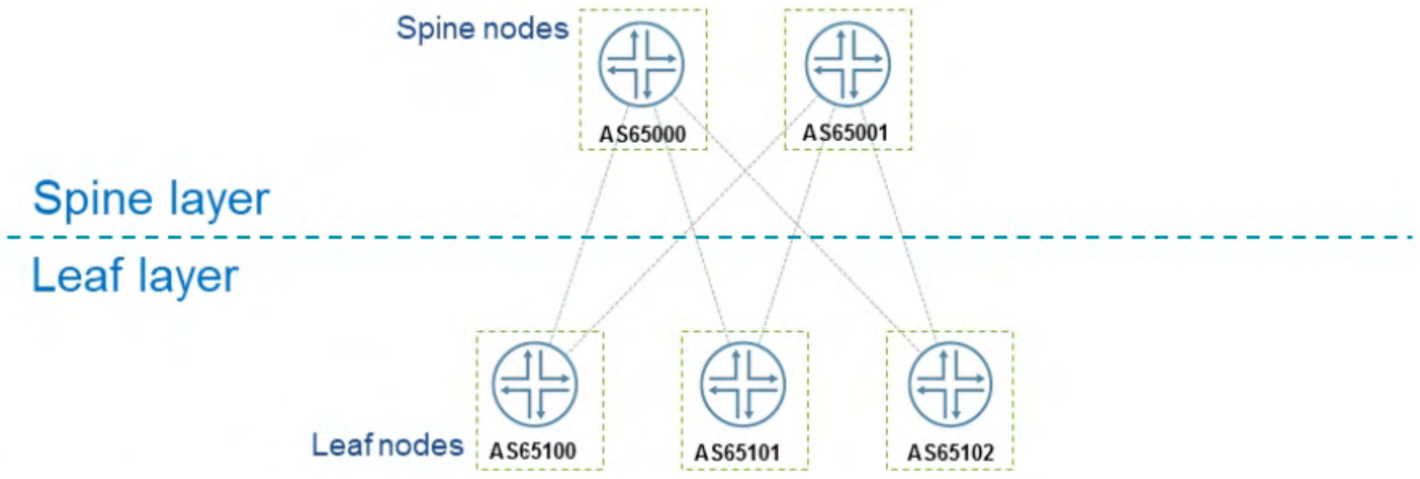

No IGP is required. The devices that participate in the IP-Fabric set up EBGP peering sessions over the physical links.

For management and monitoring purposes, all the leaf and spine nodes can be configured to advertise their loopback IP to the rest of the IP-Fabric. To provide other devices with IP-connectivity, the leaf nodes will need to be configured with a policy that advertises the prefixes used on the server-facing interfaces.

By default, Junos devices will not perform load-balancing. To make sure that spine and leaf nodes start to load share traffic across all of the available links, configure a load-balancing policy and apply that to the forwarding-table. In addition to this, Junos needs to be configured to consider multiple EBGP paths for load-balancing. This can be done by activating the configuration knob ‘multipath multiple-as’.

After turning on load-sharing across multiple links, extra caution is needed when links of different speeds are used in the network. Because of the fact that BGP does not factor in the link-speed, traffic can easily be black-holed when links are stressed with high traffic loads.

IP-Fabric using IBGP

An IP-Fabric can also be build using IBGP. In this case, all of the BGP speakers are configured with the same autonomous system.

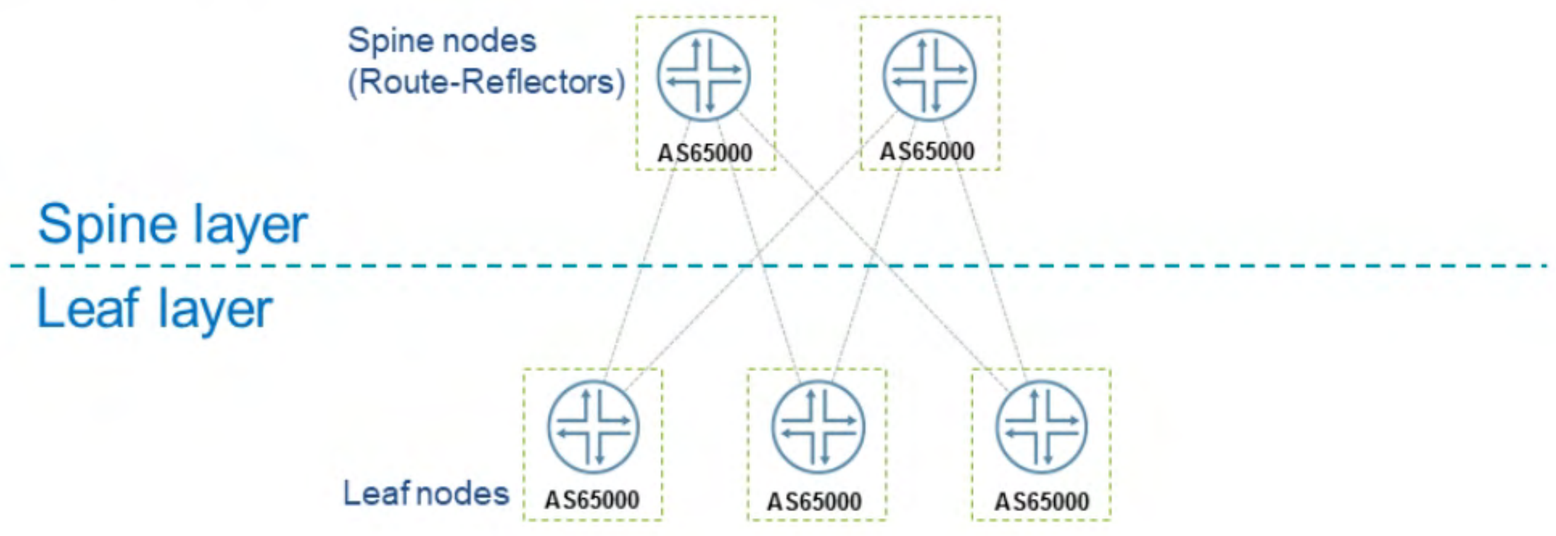

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The IBGP session are sourced from the loopback interface and the next-hop in BGP route-advertisements is typically set to the loopback IP address as well. For the IP-connectivity between the loopback interfaces, the IBGP based IP-fabric relies on an IGP.

By default, IBGP learned routes are not advertised across IBGP sessions. As a consequence, a full-mesh of IBGP sessions is required inside an IBGP based IP-fabric. Alternatively, either of the two BGP scaling mechanism can be used. Even though route-reflection and confederations both work the best and most practical fit would be to go for route-reflection. In a three stage Clos IP-Fabric, the spine switches or routers are good candidates to consider for the route-reflector role.

To achieve load sharing across multiple links inside an IBGP based IP-Fabric, several requirements have to be met. First of all, a load-balancing policy needs to be configured and applied to the forwarding table. Second, the IGP needs to have equal-cost paths towards the IBGP next-hop. Third, the BGP sessions need to be configured for multipath.

When route-reflection is used, you need to consider the fact that an IBGP peer only advertises the best route. This means that when the route-reflector learns multiple paths, only the best path will be reflected to the route-reflector clients. Through the use of ‘ADD-PATH’, IBGP peers are able to advertise and receive multiple best routes. The ‘ADD-PATH’ capability needs to be configured on all IBGP speakers.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: EBGP-based Clos IP Fabric

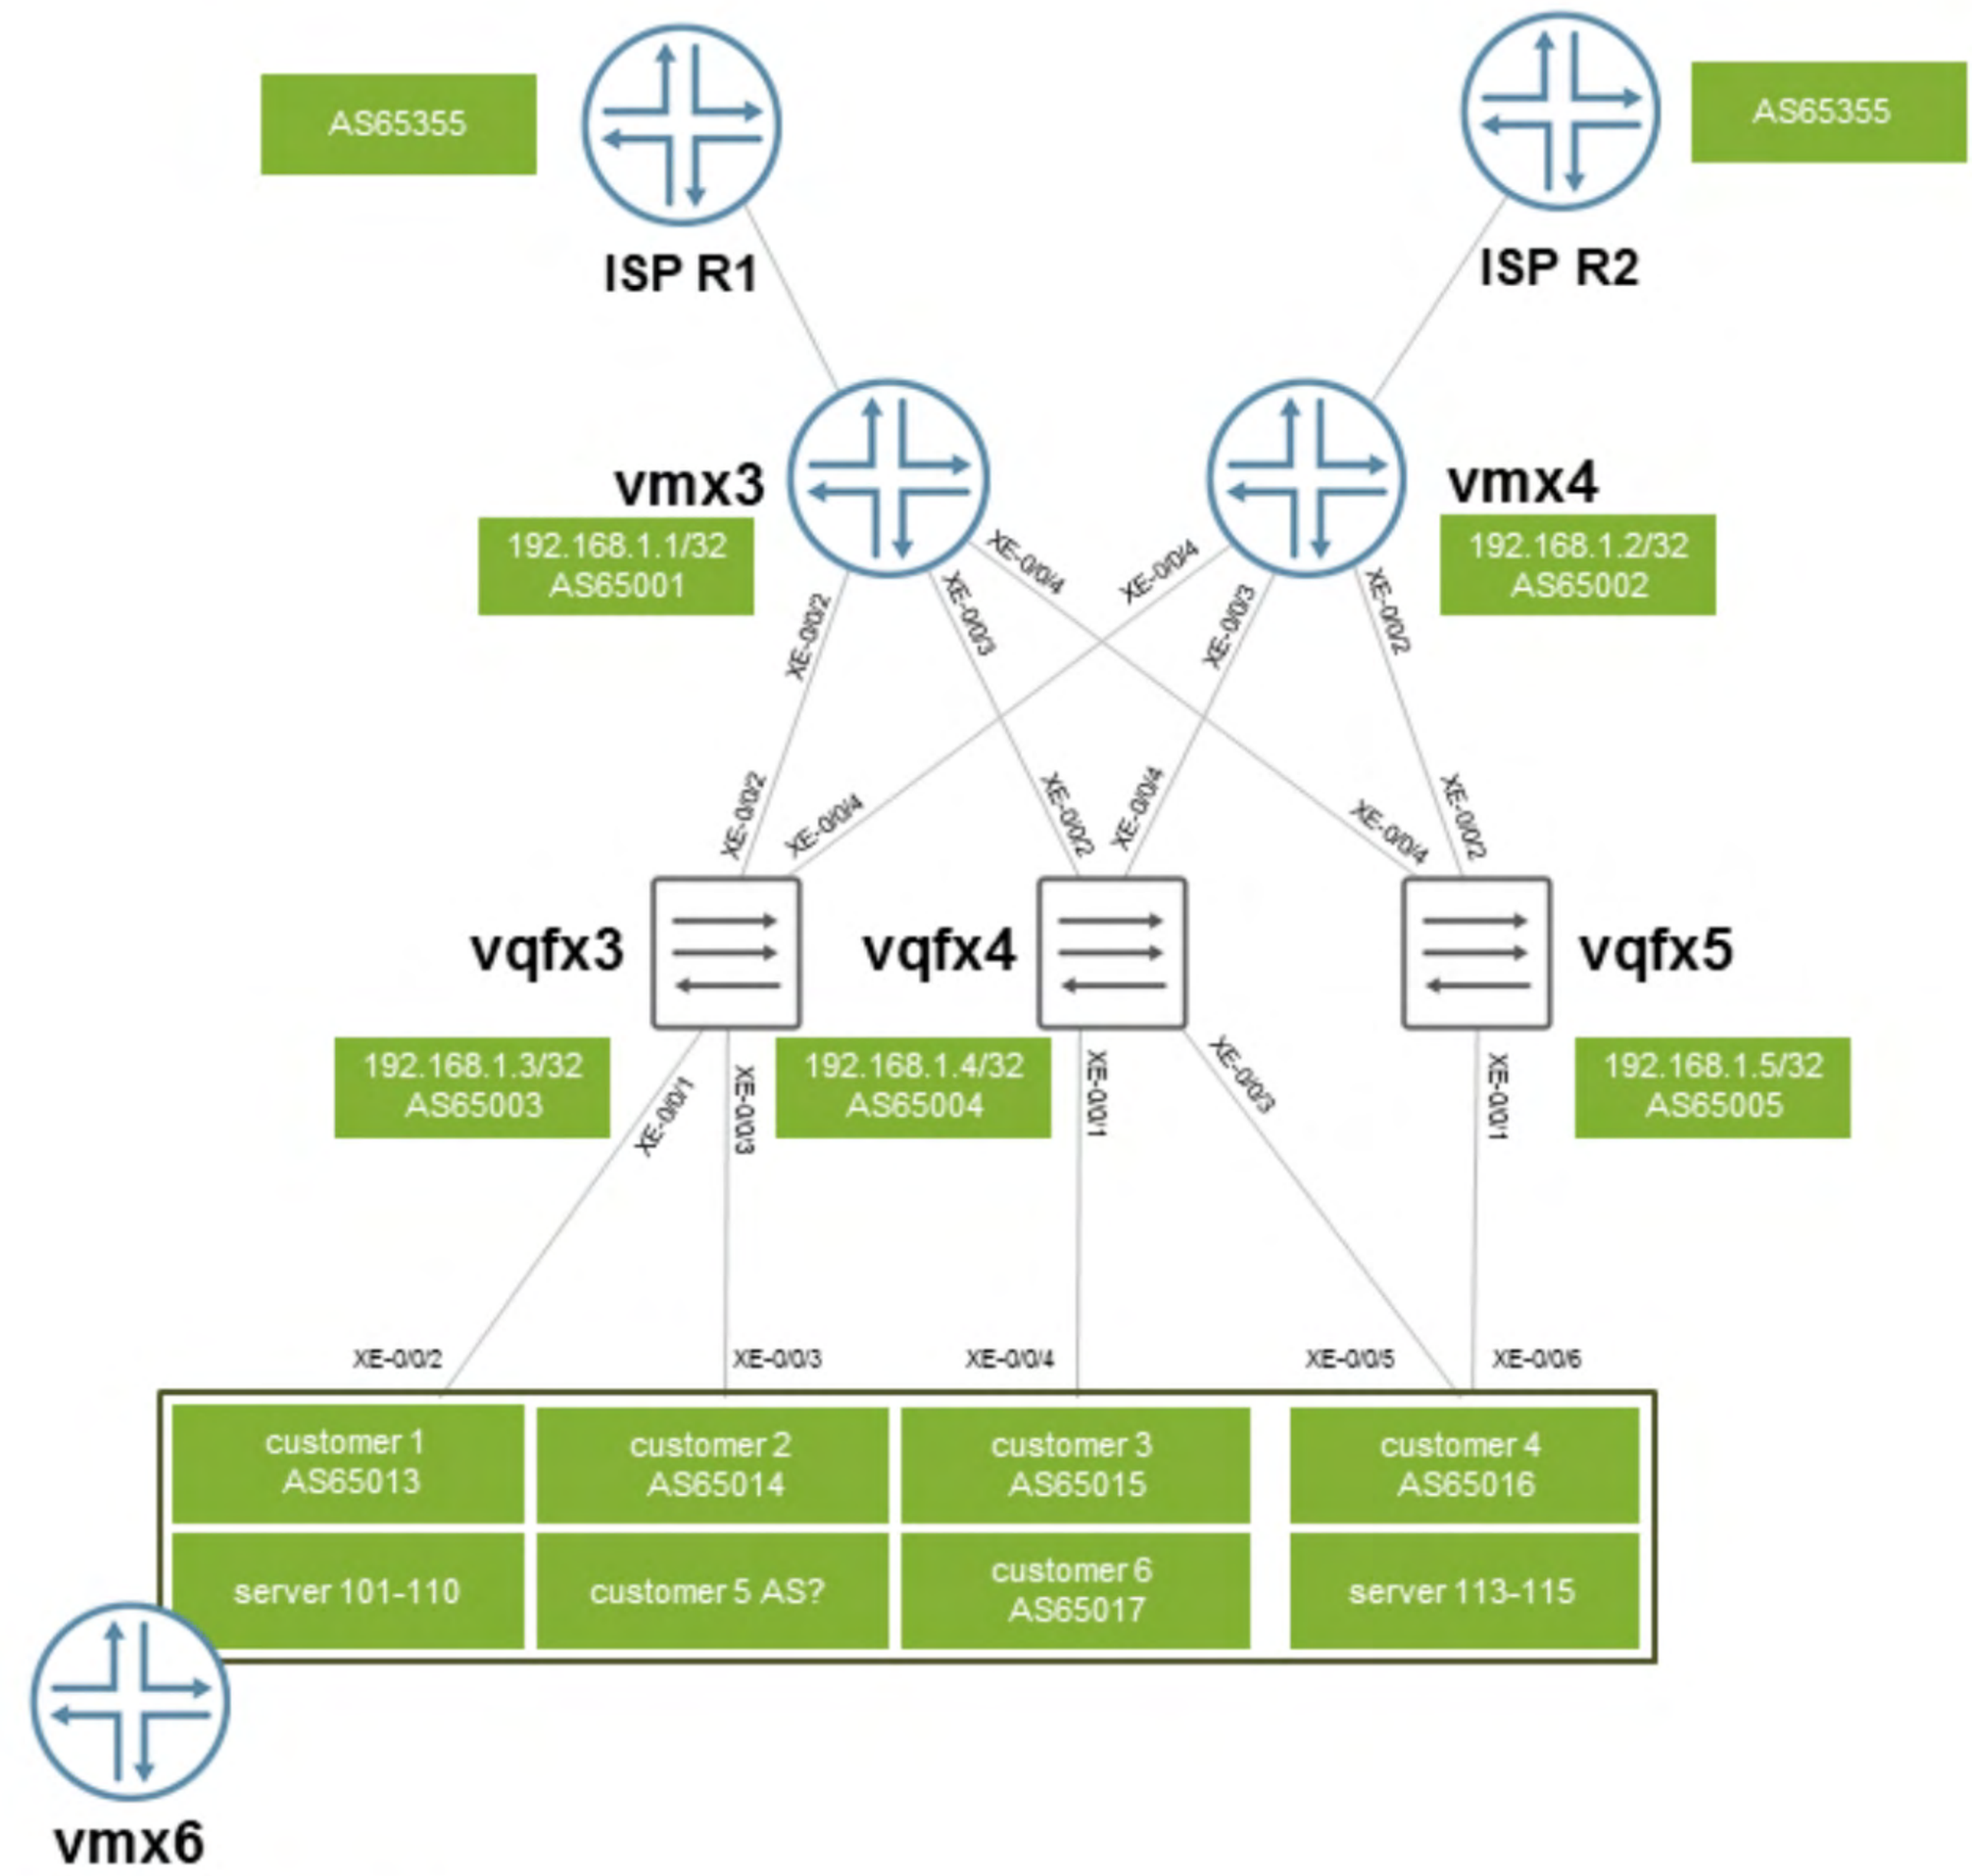

Figure 4 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

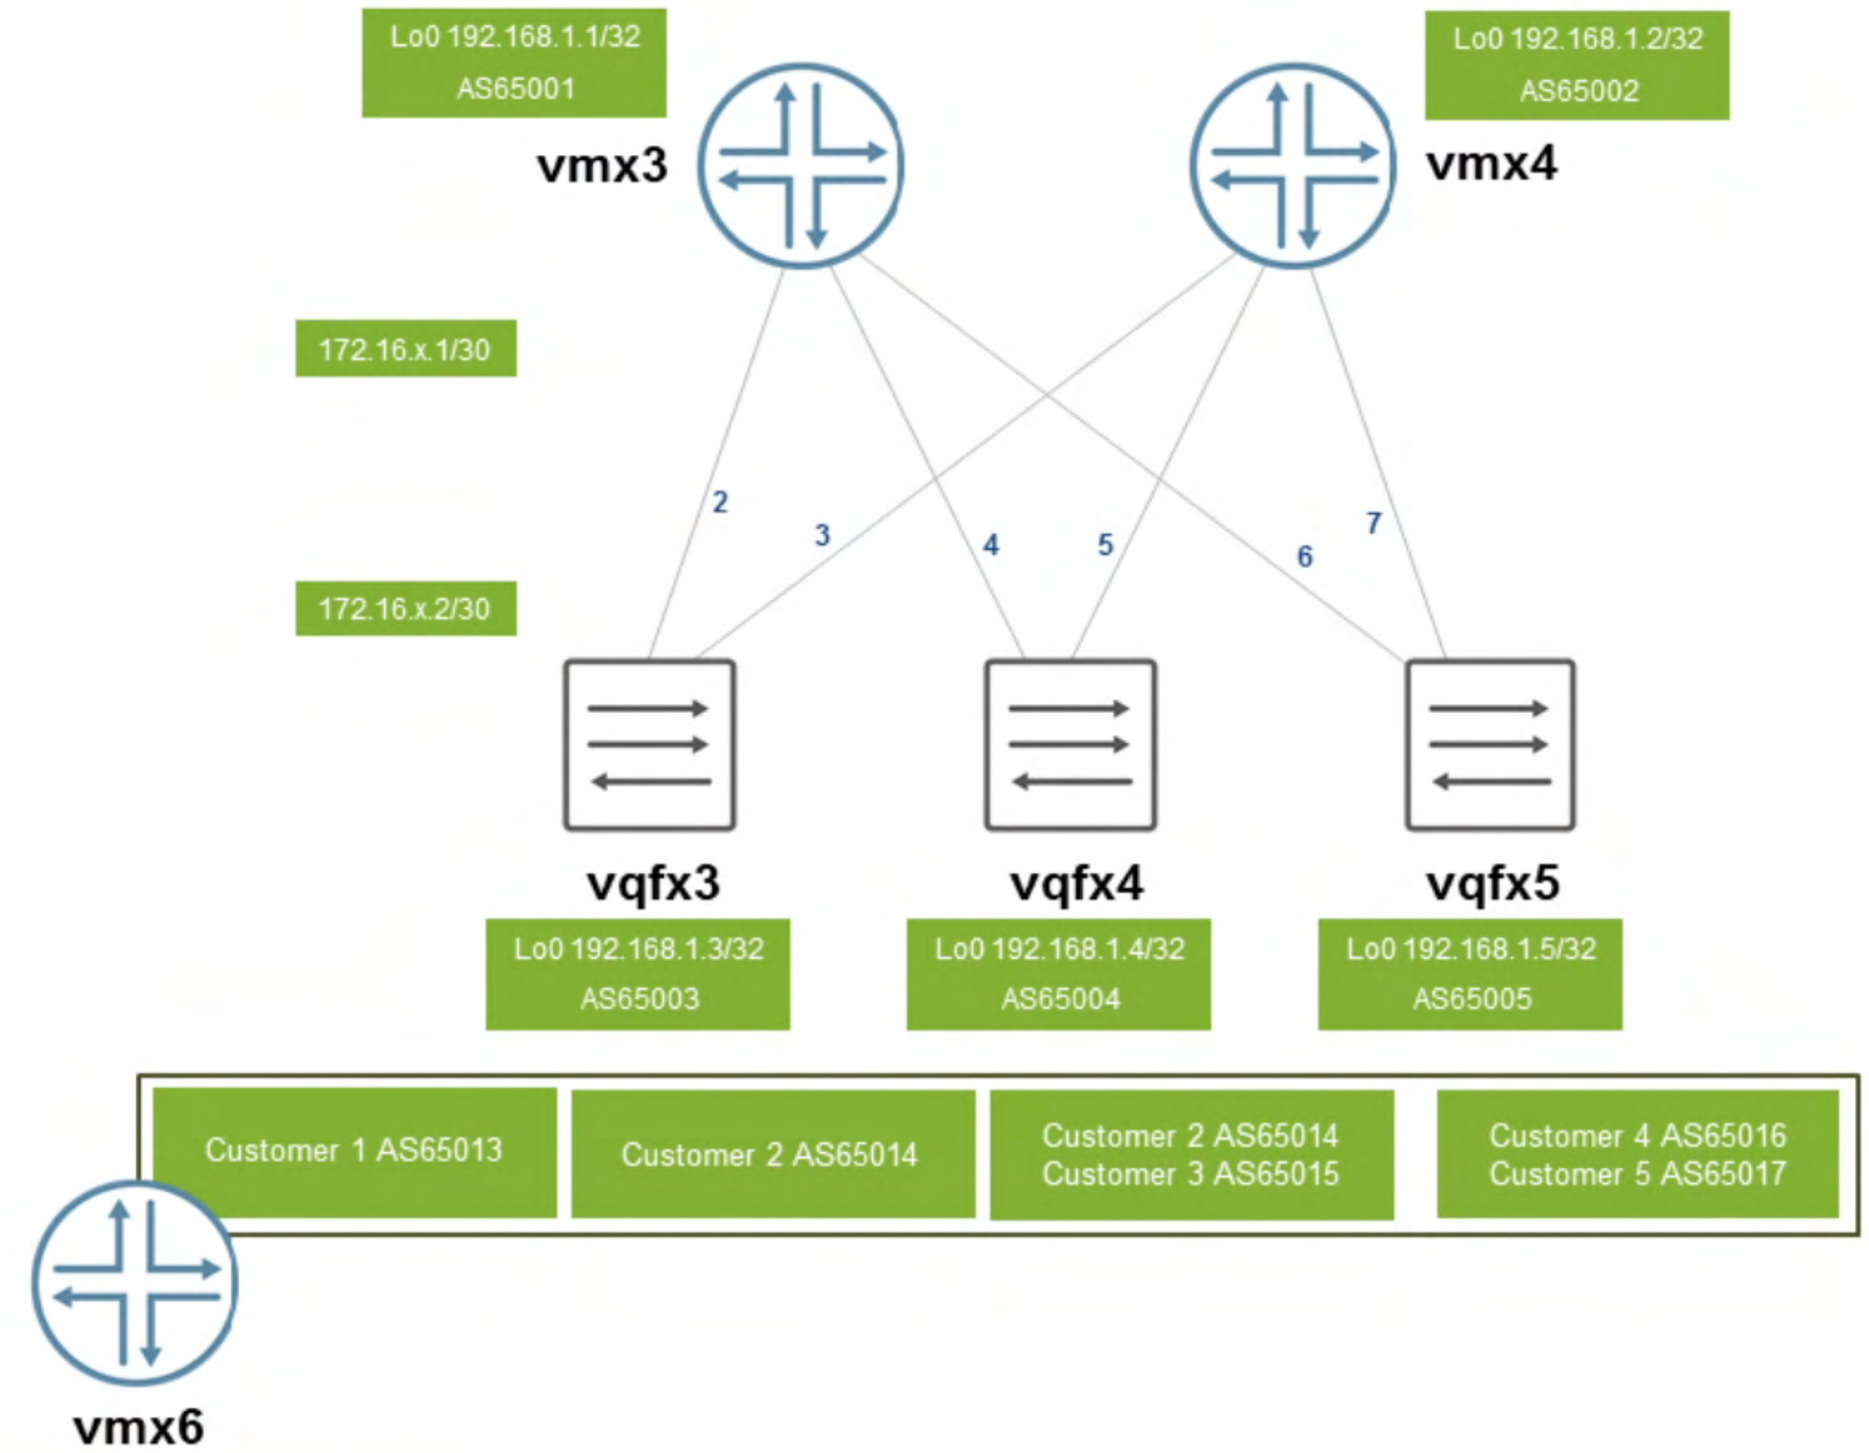

Figure 2 - Logical topology

Note: The starting topology is configured with IP-addressing only.

Task 2.1: Core Network BGP Configuration 1

- Configure BGP according to the information provided in the logical topology.
- Ensure IP-connectivity between all of the loopback IP addresses.
- Make sure all BGP sessions are 'Established' and that sessions state changes are logged to syslog.
- Write all BGP-related syslog messages to a file called ‘bgp.log’.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.2: Core Network BGP Configuration 2

- Make sure that BGP can detect the loss of a neighbor within 15 seconds. Do not alter the hold time anywhere.
- Activate per-flow load-balancing between the leaf switches and the spine routers.
- Authenticate the BGP sessions using the MD5-hased password 'jncie'

Task 2.3: ISP Uplink Configuration

- Configure EBGP between the ISP and both vMX3 as well as vMX4.
- Make sure that the keepalive is 3 seconds.
- Present the IP-Fabric to both ISP routers as AS65500.

Task 2.4: IP-Fabric BGP Customer Configuration

- Configure EBGP between the IP-Fabric and all of the customers.
- The IP address customer 3 is using to set up peering is unknown. Use monitoring tools to obtain what address the BGP session is sourced with.
- The AS number of Customer 5 is unknown, you need to find out what AS the customer is using.
- Customer 6 is configured to peer with the wrong AS number. Make sure the peering is established and make sure this wrong AS number will not be added to the AS-path when prefixes from that customer are advertise to other devices.
- All sessions need to be authenticated using the password 'jncie'.

AS	IP	AS
Customer 1	172.16.200.2	65013
Customer 2	172.16.200.6	65014
Customer 3	?	65015
Customer 4 (vQFX4)	172.16.200.14	65016
Customer 4 (vQFX5)	172.16.200.18	65016
Customer 5	172.16.200.22	?
Customer 6	172.16.200.26	65017

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.5: Server 100-110 Internet Feed.

- Server 100-110 have their gateway configured on vQFX3. Advertise the subnets in which these servers are placed.
- Make sure that vQFX3 will advertise any additional subnet that is configured for the same interface without any additional configuration.

Task 2.6: Server 113-115 Configuration

- Configure VRRP and authenticate the control messages with the 'jncie' password.
- You can only have VRRP exchange messages inside VLAN 113.
- As long as vQFX5 has a default route, vQFX4 should not be the primary router.
- The rest of the IP-Fabric should prefer vQFX5 to reach the servers.
- Make sure the interface of the vQFX will remain down for at least 45 seconds after the device boots.

Task 2.7: IP-Fabric Customer Policy Configuration

- Make sure that customer 1 has a keepalive interval of 10 seconds.
- The session with customer 1 should be torn down for 5 minutes when more than 10 prefixes are announced.
- Make sure that flapping routes from customer-2 are dampened.

Task 2.8: IP-Fabric Customer Policy Configuration

- The customers should only receive a default route.
- Do not accept any default route.
- Prefixes advertised by the customers should be identifiable by a unique community.

Task 2.9: ISP Policy Configuration

- Advertise a single summary IPv4 route (10.0.0.0/8) that aggregates the datacenter range to the ISPs.
- The ISP only allows you to use your own AS.
- Do not accept any other route then the default route from the ISP.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: IBGP-based Clos IP Fabric

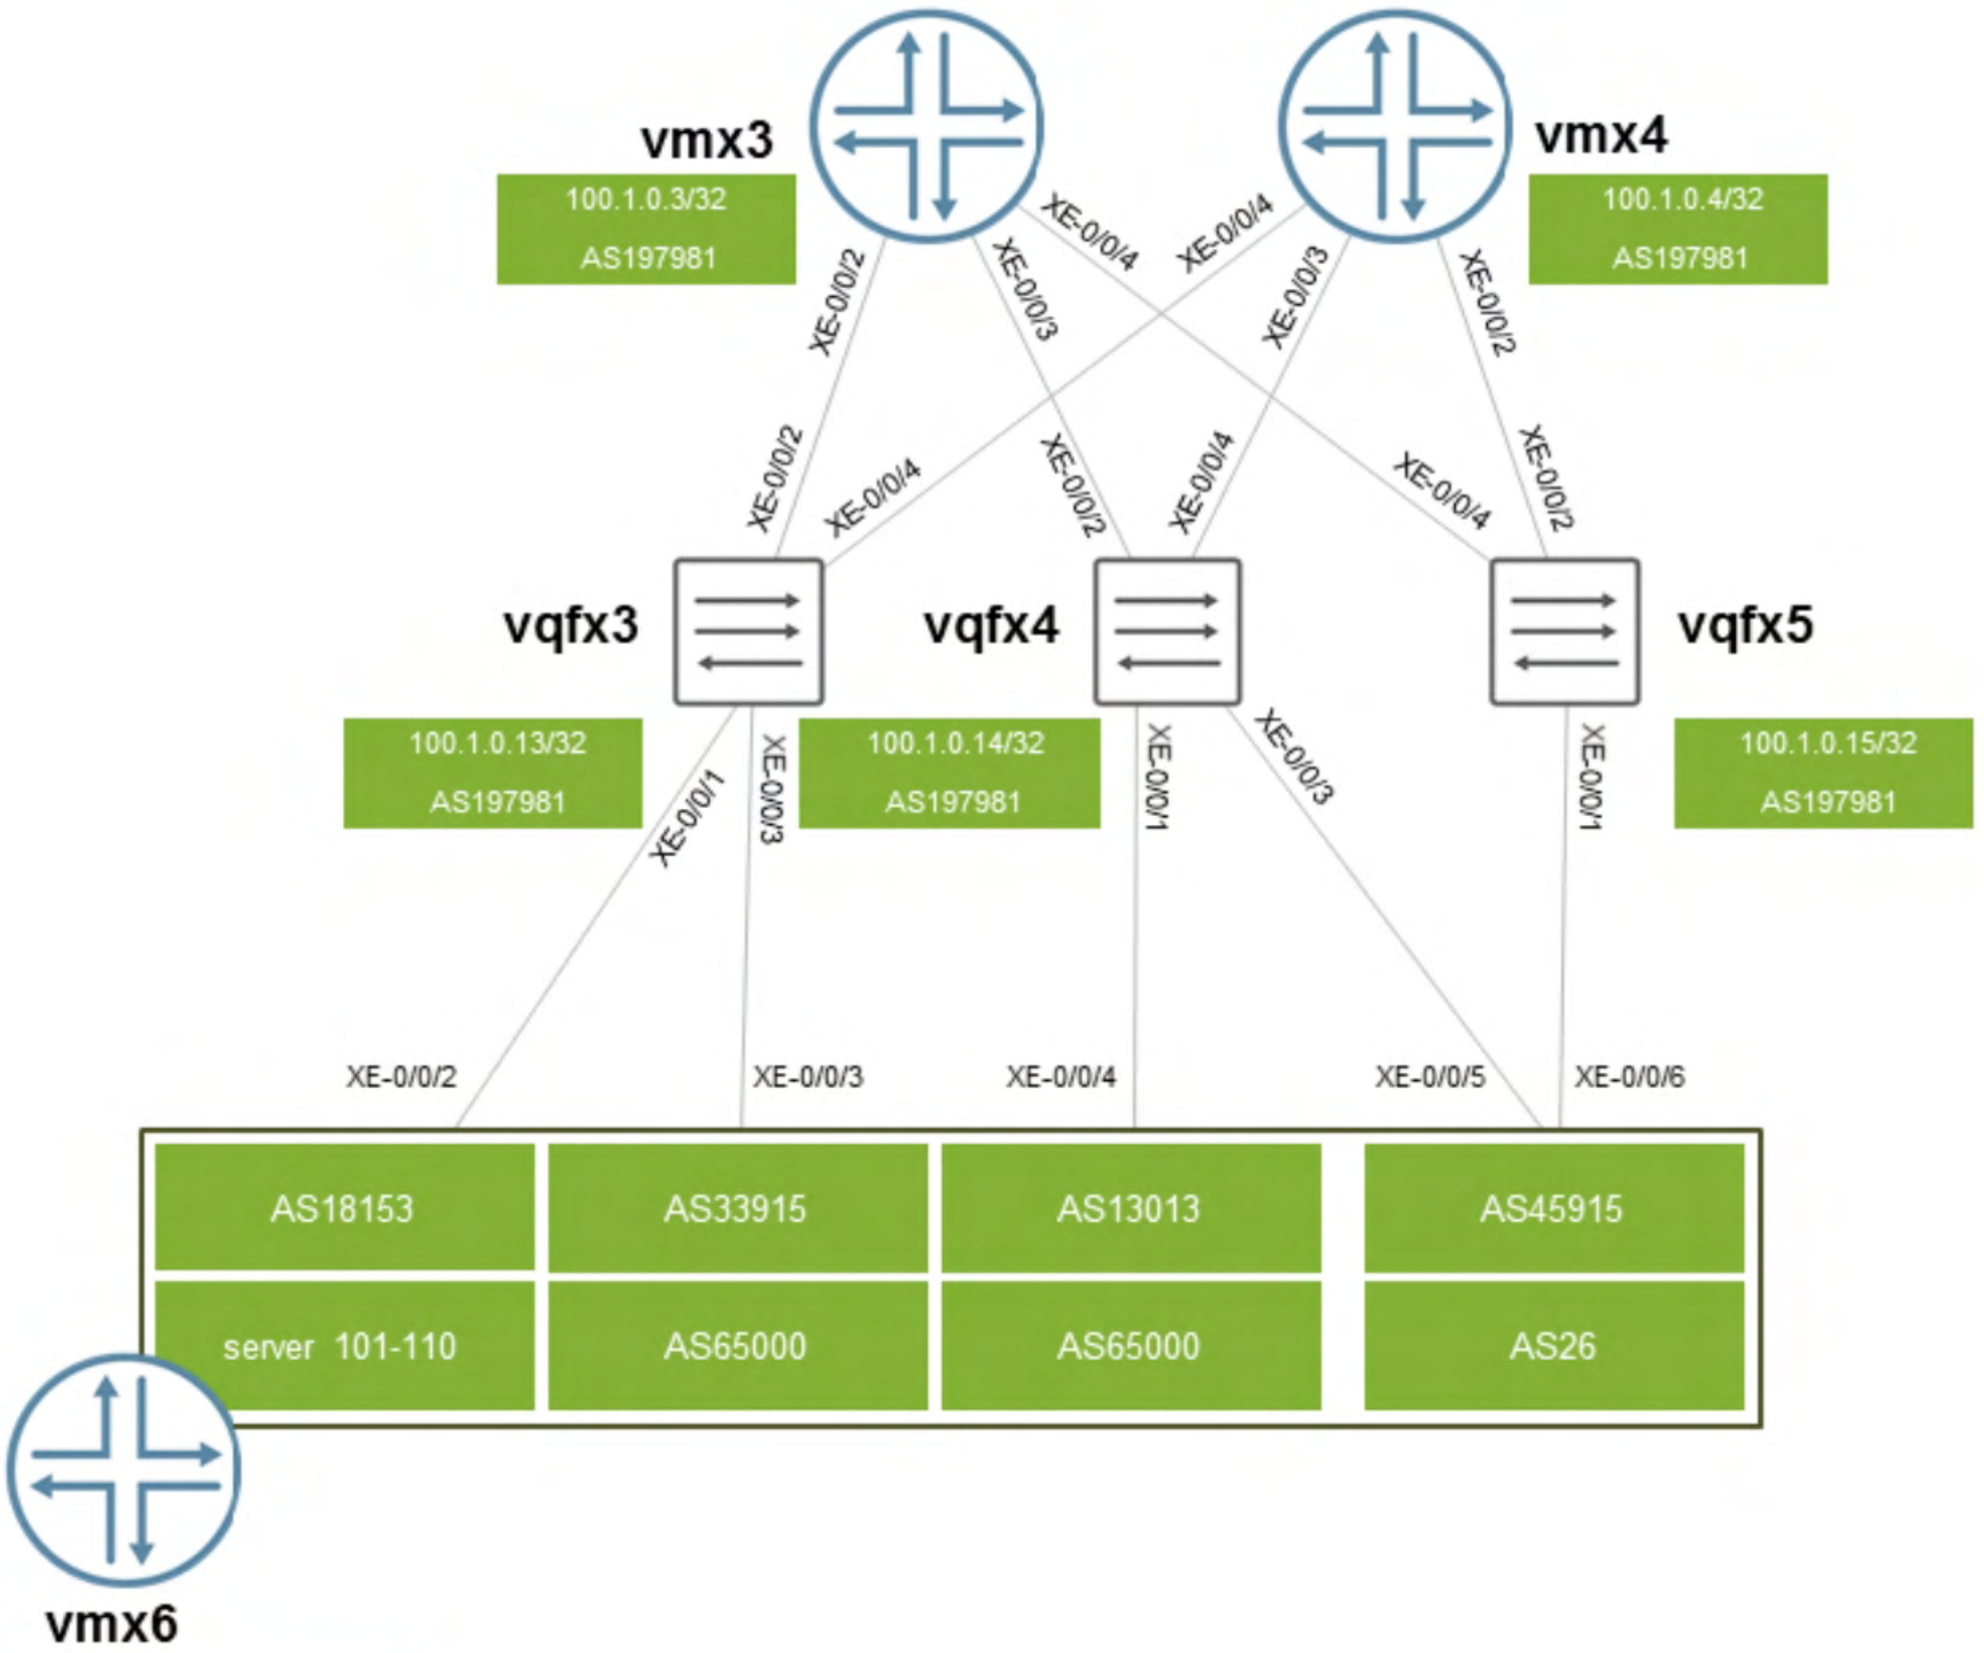

Figure 5 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

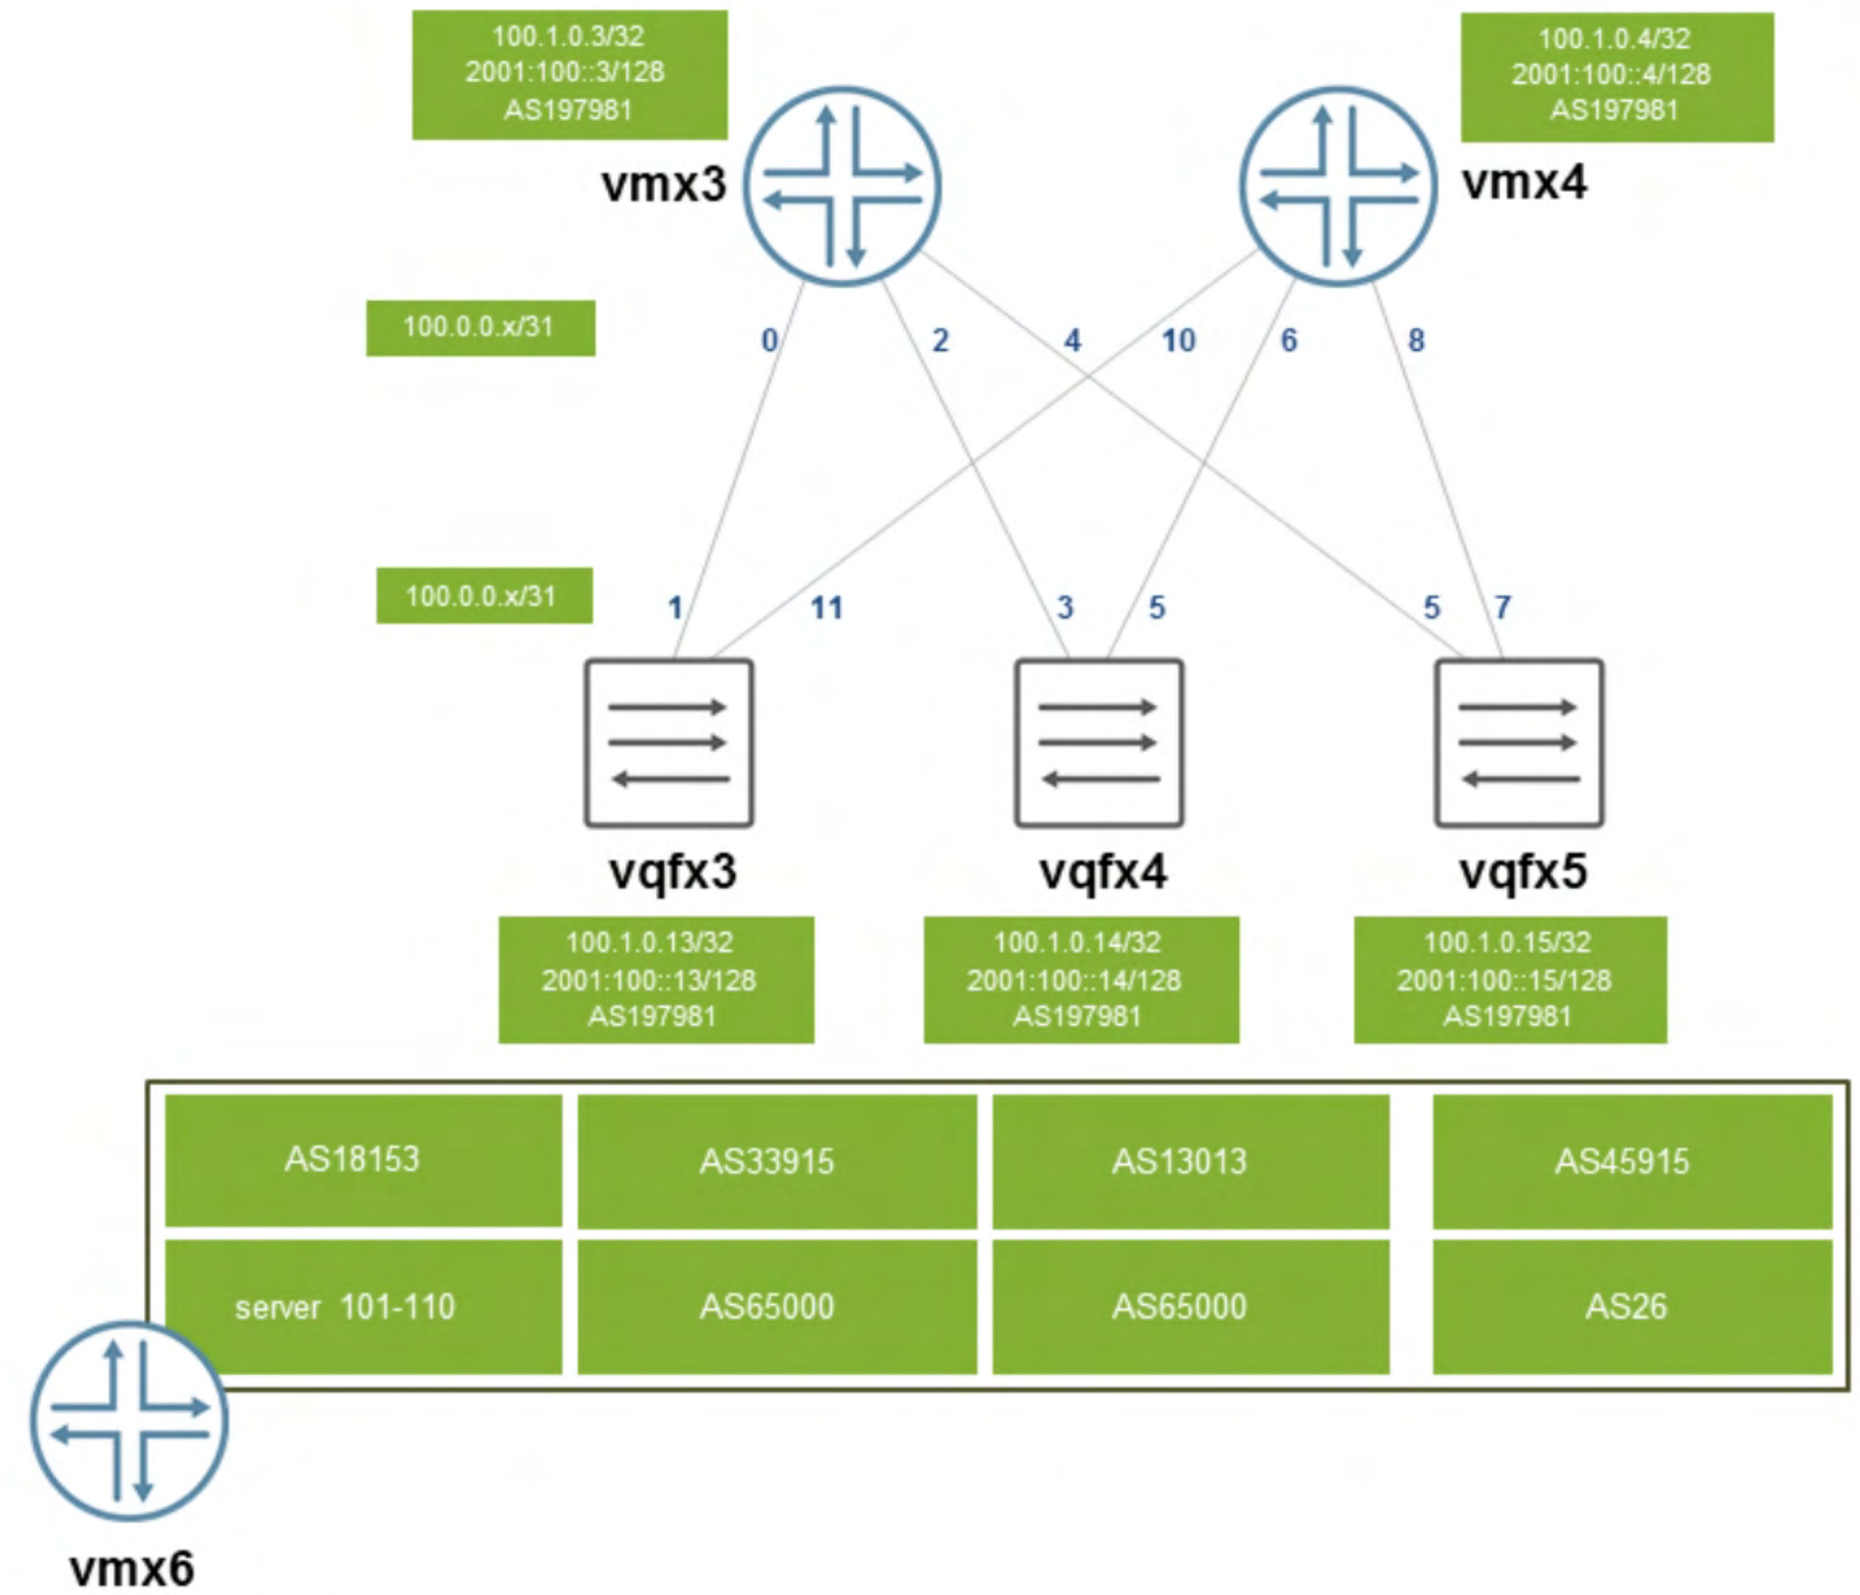

Figure 2 - Logical topology

Note: starting topology is configured with IP-addressing only.

Task 2.10: IGP Configuration

- Place the nodes that make up the IP-Fabric into a single area OSPFv3.

Task 2.11: IGP Configuration

- Links actively running OSPFv3 should be active for IPv4 as well.
- Have OSPFv3 advertise the IPv4 loopback address on every node.

Task 2.12: IGP Configuration

- Ensure the utilization of equal-cost links.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Make sure OSPFv3 calculates path cost using a reference bandwidth of 100Gbit.

Task 2.13: IPv4 IBGP Configuration

- Configure an IBGP session between the vMX routers and the vQFX switches.
- vQFX switches are not allowed to peer with each other.
- Authenticate all IBGP sessions using password 'jncie'.

Task 2.14: IPv6 IBGP Configuration

- Configure an IBGP session between the vMX routers and the vQFX switches.
- vQFX switches are not allowed to peer with each other.
- Authenticate all IBGP sessions using password 'jncie'.

Task 2.15: Server 101-110

- Advertise the prefixes for server 101-110 into BGP
- Make sure you can identify the routes to these servers anywhere inside the IP-fabric.

Task 2.16: EBGP Configuration

- Configure EBGP sessions with the autonomous systems displayed in the topology.
- Make sure that the received routes are correctly distributed throughout the IP-Fabric.

AS	Neighboring IPv4	Neighboring IPv6
18153	55.0.0.1	
33915	33.0.0.1	
65000	100.2.0.2	2001:100:2::1
65000	100.2.0.6	2001:100:2::3
13013	13.0.0.2	
45915	8.0.0.1	
26	26.0.0.2	2001:26::1
45915	8.0.0.5	
26	26.0.0.6	2001:26::3

Task 2.17: Local AS Advertisement

- Advertise the IP ranges for your AS to all autonomous systems.
    - For IPV4, advertise 100.0.0/8
    - For IPv6, advertise 2001:100::/32

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Make sure that leaf switches will not advertise any of these ranges when the connections to the spines are lost.

Task 2.18: BGP RIB Configuration

- Make sure that redundant address information from multihomed autonomous systems is available in the RIB as well as the FIB on every device.

Task 2.19: AS65000 Customers

- Make sure that the two different customers that are using AS65000 receive all routing information.
- When prefixes that originate from AS65000 are advertised to other public ASes, AS65000 should not appear in the AS-path.

Task 2.20: IP-Transit Policy Configuration

- Prefer the IP-transit that AS197981 is multihomed to for inbound as well as outbound traffic.
- Only advertise a default route to all non-transit networks.
- Do not alter the local-preference anywhere and only apply policies on vQFX3.
- AS33915 should only receive a default route and the route for your AS.

Chapter 3: Controllerless Overlay

Instead of using an IP-Fabric to deliver layer 3 services, the IP-Fabric can also be leveraged to serve as a layer 3 Clos underlay network. A controllerless overlay can be configured on top of this underlay and this controllerless overlay network can be used to offer both layer 2 as well as layer 3 services.

A controllerless overlay uses VXLAN encapsulation in the data plane and EVPN signaling in the control plane. Some of the benefits to using EVPN in combination with VXLAN include the following:

- clean separation of the control plane and the data plane.
- active-active multihoming for both layer 2 as well as layer 3 traffic.
- integrated L2 and L3 solution.
- MAC-mobility across different data centers.
- optimal routing between the data centers for both layer 2 as well as layer 3.

VXLAN

Virtual Extensible LAN, or VXLAN, is defined in RFC 7348. In this RFC, VXLAN is described as the data plane operation of a layer 2 VPN. With VXLAN, Ethernet frames are

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

encapsulated into IP packets. VXLAN can fulfill the data plane role for several control plane strategies.

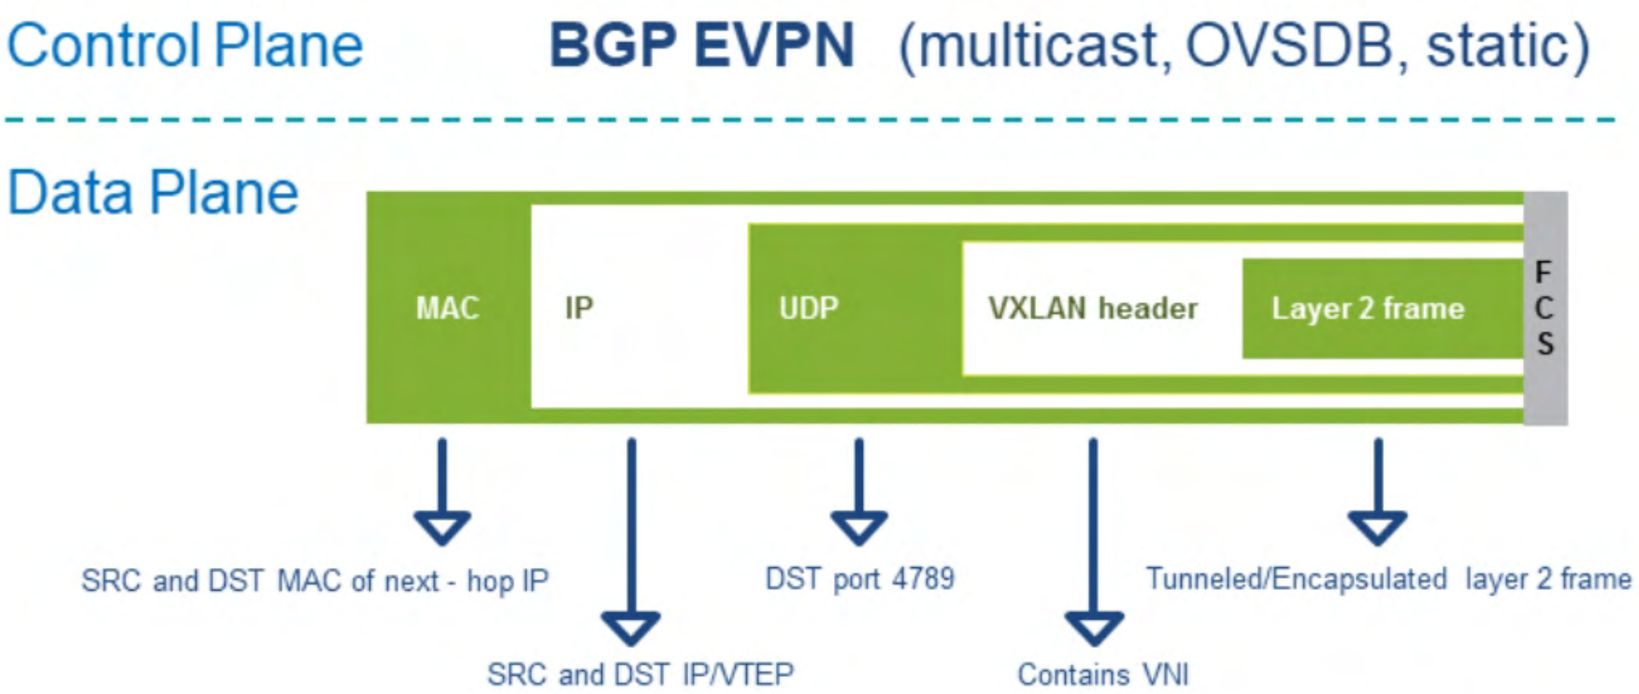

The encapsulation and decapsulation of VXLAN packets are performed by a VTEP, or Virtual Tunnel End-Point. By default, the original VLAN tag (if any) is stripped from incoming layer 2 frames upon encapsulation. After this, a VXLAN header is added to the original layer 2 frame. The VXLAN header contains a 24-bit VXLAN Network Identifier, or VNI. This is used to identify what VXLAN segment the frame belongs to. After the VXLAN header, a UDP header is added. By default, the destination port will be set to 4789. The UDP packet is encapsulated in an IP packet. The source and destination IP addresses used are the source and destination VTEP IP addresses.

EVPN

Control plane signaling inside a controllerless overlay is taken care of by EVPN. Understanding EVPN is key to understanding how to configure, troubleshoot and operate a controllerless overlay.



EVPN is an NLRI carried in MP-BGP. The AFI is 25 (L2VPN) and the SAFI is 70 (EVPN). Before BGP speakers can exchange EVPN routes, both peers have to be configured with the proper address family. On a Juniper device, this means the BGP session has to be configured with the **evpn signaling** family.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

EVPN uses 5 different route types to exchange routing information. Type 1 to 4 are discussed in RFC7432. Type 5 is discussed in 'draft-rabadan-l2vpn-evpn-prefix-advertisement-01'. We briefly describe each route type below:

Type 1: Ethernet Auto-Discovery route: Required for multihomed CEs. Network-wide messages used to perform fast withdraw, split horizon and multipath.

Type 2: MAC/IP advertisement route: A per VLAN route that is used to advertise IP and/or MAC addresses.

Type 3: Inclusive Multicast Ethernet Tag route: used to signal to other PEs that they can set up a path for BUM traffic towards the local PE for a specific VLAN or ESI.

Type 4: Ethernet Segment route: PEs connected to the same ESI can discover each other using this route-type. This route type enables multihoming and it is used for the DF election.

Type 5: IP Prefix route: an IP-prefix route that can be used for layer 3-connecity in an EVPN environment.

EVPN routes are associated with one or more target communities. These target communities are BGP extended communities. These target communities are used to determine to what VXLAN an EVPN route belongs. The PE attaches a route-target on export and evaluates any attached route-target(s) on import. Through the use of export policies, you can determine not only what routes are advertised, but also what communities are attached to these routes. You can set a single route-target for all routes advertised by a PE, explicitly configure a route-target per VXLAN or have the PE generate them automatically. Through the use of an import policy, you can make the PE import the route and decide into what VRF the route should be copied to or with what VXLAN the route should be associated with.

Additionally, every route has a route-distinguisher (RD) prepended to it. The RD is an 8-octet field and needs to be configured on every device that participates in the overlay network. On QFX, the RD is configured under the [switch-options] stanza. On MX, the RD is configured either under the [routing-options] stanza or under the [routing-instance *routing-instance-name* route-distinguisher] stanza.

Additional explanations on route-targets and route-distinguishers can be found in chapter 4, where Datacenter Interconnects and L3VPNs are discussed.

Route Type 1: Ethernet Auto-Discovery Route

EVPN type 1 routes are required when a CE is multihomed. They are network-wide messages. In the case of a controllerless overlay, this means that the route is send to all

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

other devices involved in EVPN signaling. The type 1 route is used to perform fast withdraw, split horizon and multipath.

Fast withdraw.

When the site-facing interface of a PE that is connected to a multihomed segment fails, the PE will withdraw the type 1 route associated with that ES. This signals all of the remote PEs that the ES is no longer reachable via that PE. This will immediately cause for all remote PEs to update their forwarding tables instead of waiting for the PE to withdraw all the individual MAC/IP routes associated with the ES. This is what is referred to as 'Fast withdraw'. After withdrawing the type 1 route for an ES, the PE will still withdraw the individual MAC/IP routes tied to the ES.

Split horizon.

EVPN introduced a new transitive extended community called the ESI label. PE devices connected to the same multihomed ES use this community to filter BUM traffic. The non-DF router will not forward BUM traffic to the CE connected to the multihomed ES. The split horizon process also leans on the type 4 route, that is used in the DF election.

Multipath.

A PE can indicate to other PEs whether or not the traffic can be load-shared across multiple links of an ES. The PE can signal this through a flag in the ESI community. When the flag is set to 0, remote PEs can use all of the links associated with the ES for traffic forwarding. When the flag is set to 1, only one of the links associated with the ES can be used for traffic forwarding.

Route Type 2: MAC/IP Advertisement Route

Every locally learned MAC address is advertised to the other PEs using a type 2 route. Optionally, the PE can use this route type to send both the MAC address and the corresponding IP address.

Every type 2 route has a route-target attached to it. Receiving PEs can use the route-target to decide whether or not the route should be imported and into what VRF table the route should be installed. Control over this process is exercised through the configuration of an import policy.

Inherent to the type 2 route is the fact that it is a per-VLAN (or per-VNI) route.

Route Type 3: Inclusive Multicast Ethernet Tag Route

A type 3 route effectively sets up a per-VNI path for BUM traffic between PE devices on a per-VLAN / per-ESI basis. A PE will advertise this route for every configured VNI. Same as

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

with the type 2 routes, every advertised type 3 route will have a route-target attached to it. Receiving PEs use these to decide on discarding the route or adding it to the appropriate VRF tables.

The type 3 route also carries a Provider Multicast Service Interface (PMSI) Tunnel attribute, indicating how BUM traffic should be sent. Options include ingress replication, PIM, etc.

Route Type 4: Ethernet Segment Route

When a PE is part of a multihomed ES, the interface is configured with an ESI. This ESI is used to determine the 'es-import-target', which is the ES-import extended community. This community is used in type 4 routes.

PE devices connected to the same multihomed ES can discover each other through the exchange of these type 4 routes. The type 4 route is sent to all other PE routers, but only PE routers serving interfaces configured with the same ESI value will import this route.

The ESI has to be unique for every multihomed segment in the network.

The route helps with the split-horizon rule and the DF-election.

Split horizon rule and DF-election.

When a PE receives a BUM packet from a locally attached device, the packet is flooded out all ports in the same VLAN and the packet is flooded to all remote PEs on which the VLAN (or VXLAN) is active. The default EVPN split horizon rule makes sure that BUM traffic is not flooded to the originator by not sending it out the port it was received on.

When a connection is multihomed, the same must apply to the PEs connected to the multihomed segment. Instead of having a PE flood BUM traffic out locally and towards the remote PE that is connected to the same multihomed ES, only one PE will flood the BUM traffic. This additional split horizon rule prevents the BUM traffic endlessly looping between the attached CE and multiple PEs serving that CE. The PE in charge of flooding is the DF.

The elected DF is the only PE that will forward BUM traffic from the core towards the multihomed CE device. The PE that is not elected to be the DF will not forward BUM traffic to the CE.

The DF election factors in the IP address of the PE and the value of the Ether tag value. This way, the DF election can occur on a per-Ethernet tag basis.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Controllerless Overlay.

In a controllerless overlay network architecture, a 'standard' IP-Fabric usually serves as an underlay network. Let's consider the following EBGP based underlay:

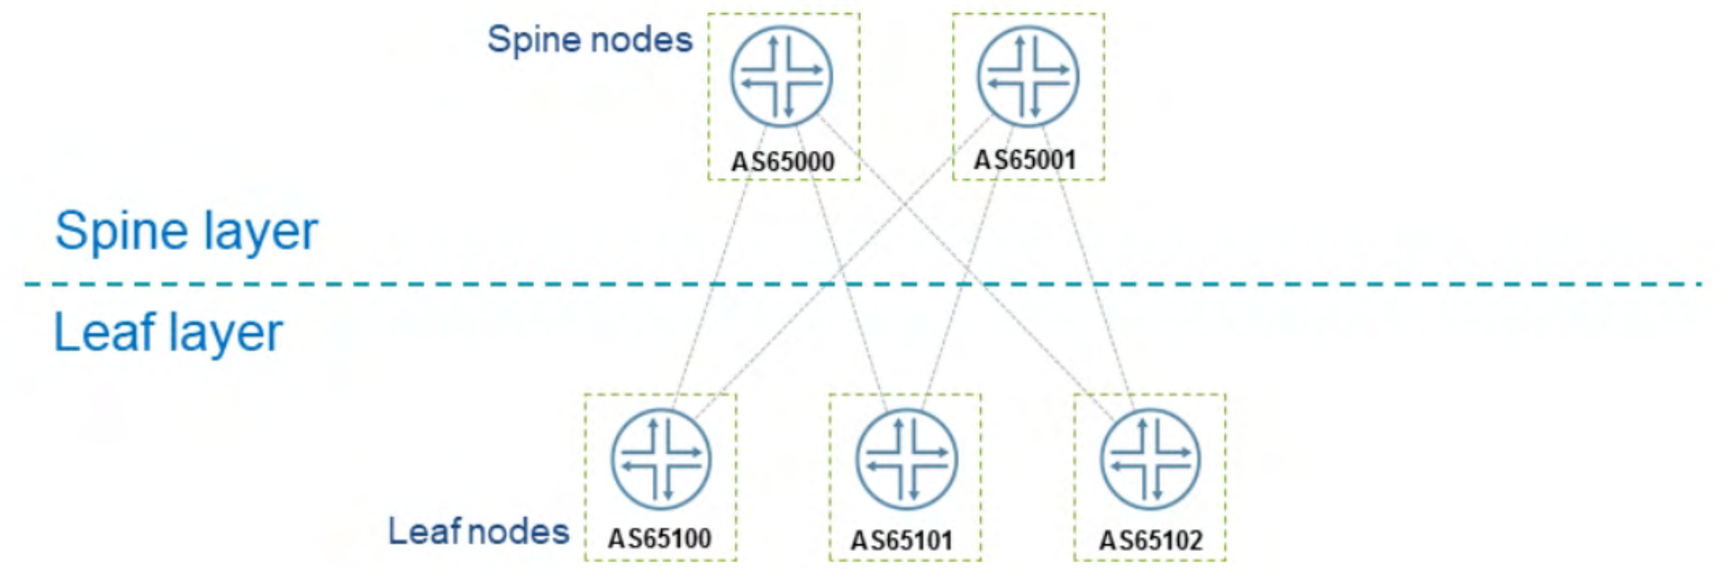

The main function of the underlay is to enable IP-connectivity between the loopback IP addresses of the nodes involved.

The overlay network will use the IP connectivity provided by the underlay for two different purposes: VXLAN encapsulation (data plane) and EVPN signaling (control plane). The following example displays three leaf nodes forming a layer 2 overlay network atop the EBGP-based underlay that is made up out of 5 devices:

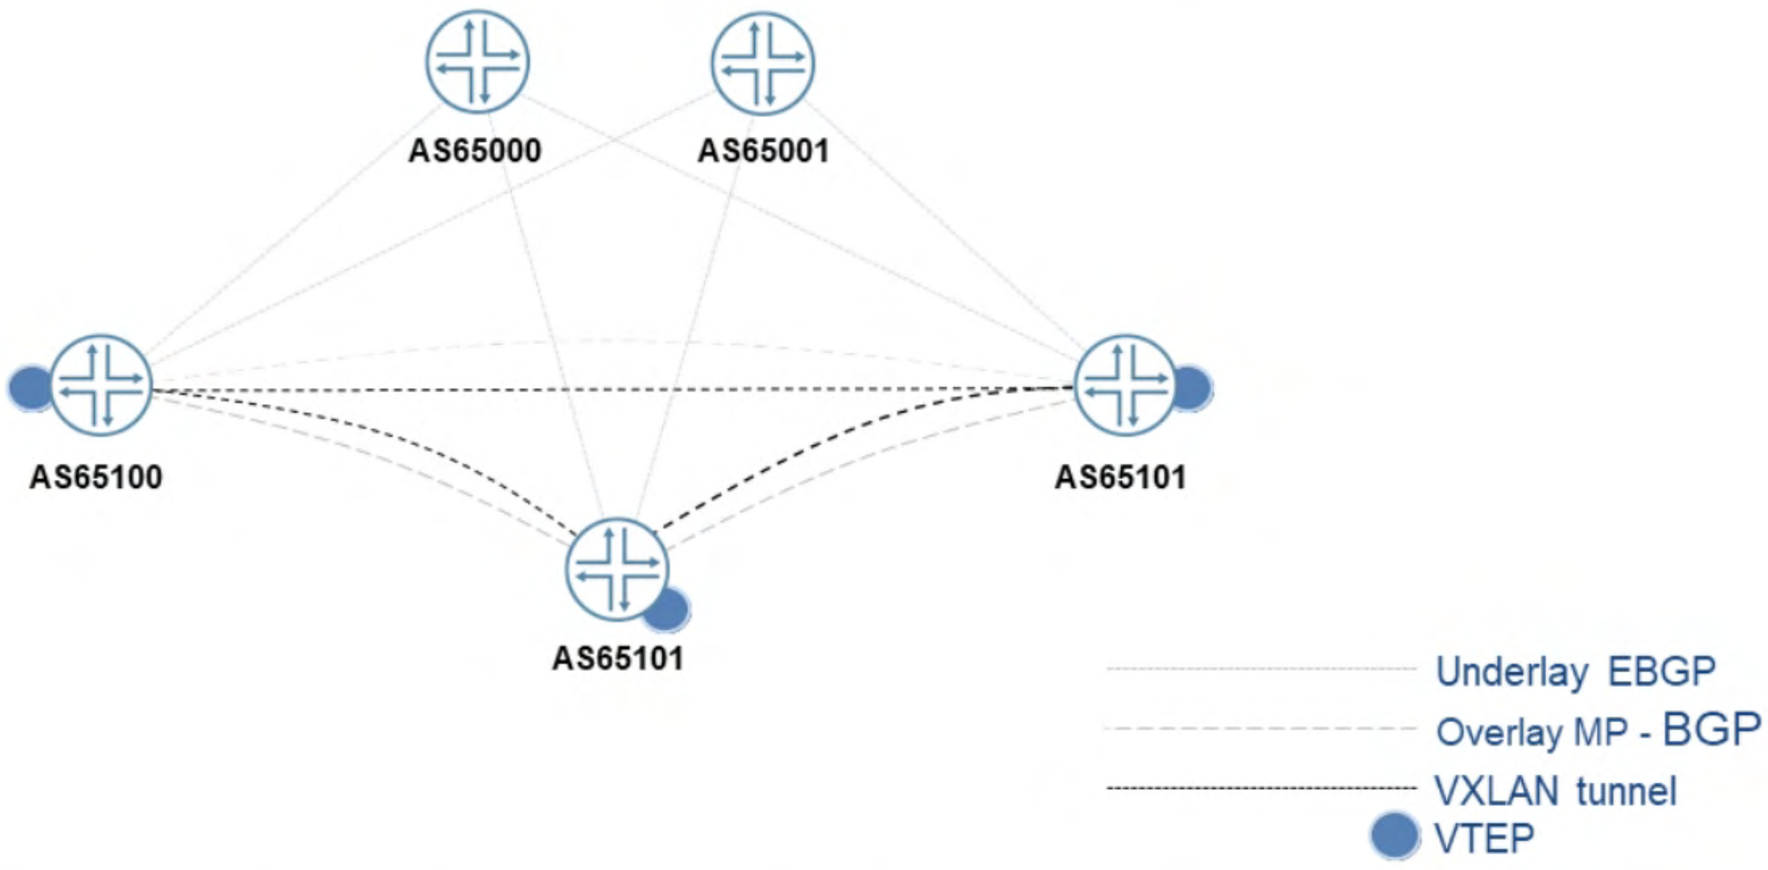

The three leaf nodes run a separate MP-BGP session that is configured for '**evpn signaling**'. This BGP session is sourced from the loopback interface and it is used to signal EVPN address information. EVPN is used to advertise and learn about the location

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

of VXLAN segments and the hosts connected to those segments. The loopback IP address on each of the leaf nodes also serves as a VTEP or VXLAN gateway. The VTEP handles the encapsulation and decapsulation of VXLAN packets.

The overlay network can also be configured with a layer 3 gateway for inter-VRF or inter-VXLAN routing. In the following design, the spine nodes are used in the overlay network:

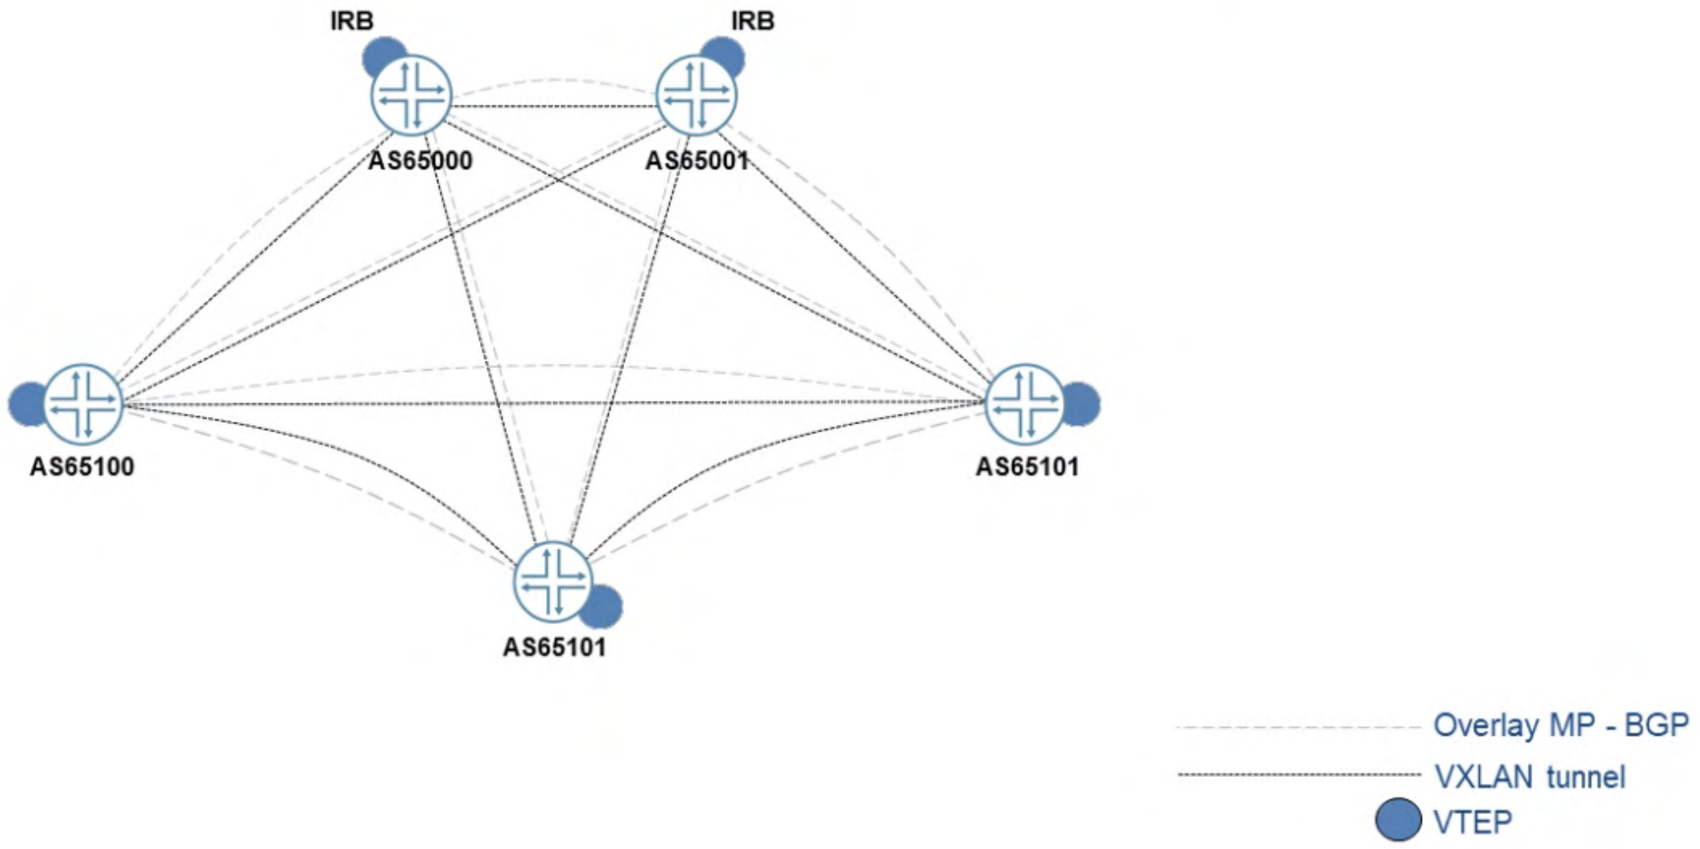

Even though it is possible to configure a single layer 3 gateway for a VXLAN segment on the spines, there are two options in which spine nodes can be used to offer redundant inter-VRF or inter-VXLAN routing.

The first option is to configure both IRB interfaces with a '**virtual-gateway-address**'. In this case, both devices will share a virtual gateway and use the VRRP MAC address as layer 2 address without actually running VRRP. The second option is to configure both IRB interfaces with the same MAC and IP address:

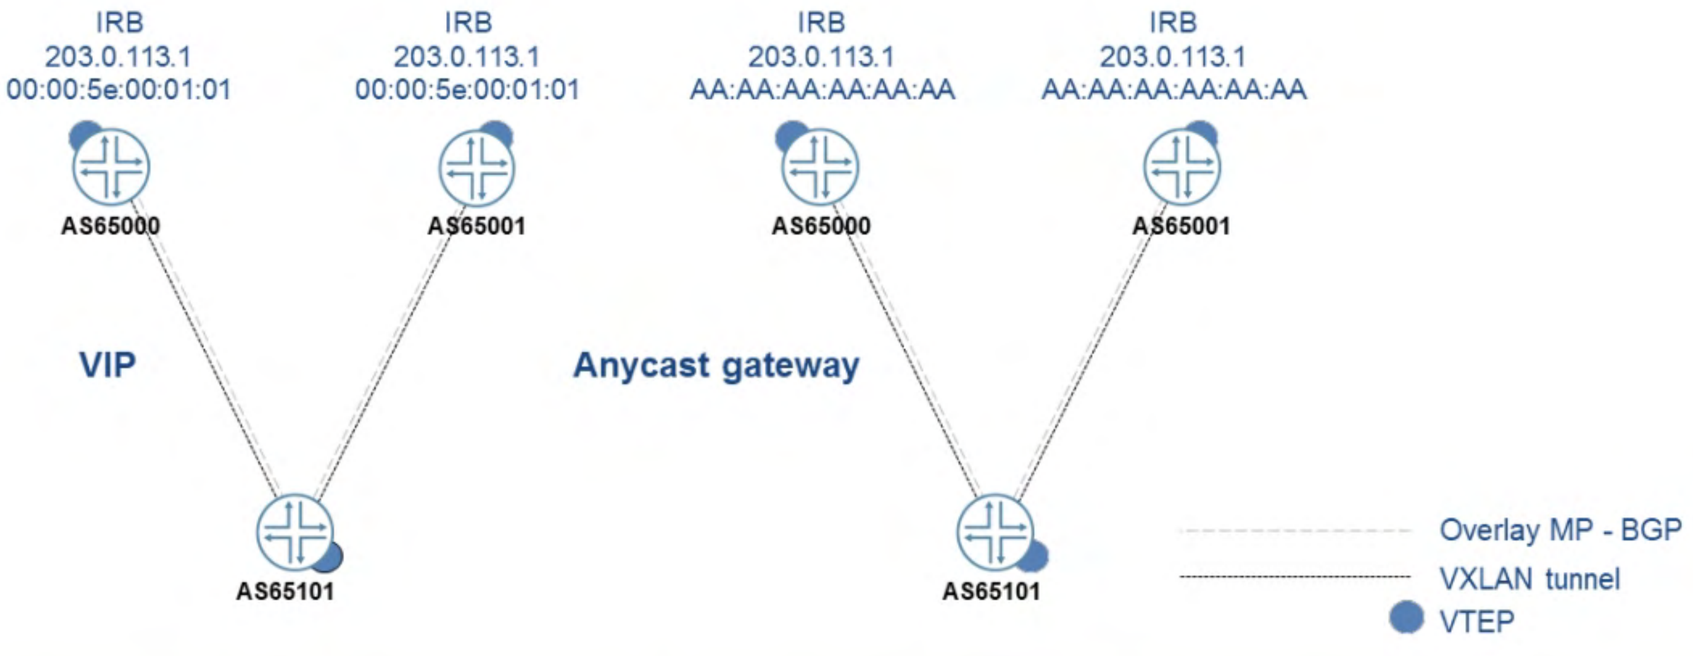

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

In both scenarios, the gateway has an identical MAC and IP address on both nodes. As a consequence, the rest of the nodes in the fabric can load share traffic towards the gateway and the hosts are oblivious to any failover.

Redundancy can also be offered at layer 2. This is done by multihoming servers/networks to multiple leaf devices of the overlay network. Typically, the device that is connected to the overlay will run LACP and see the multihomed connection as a single LAG:

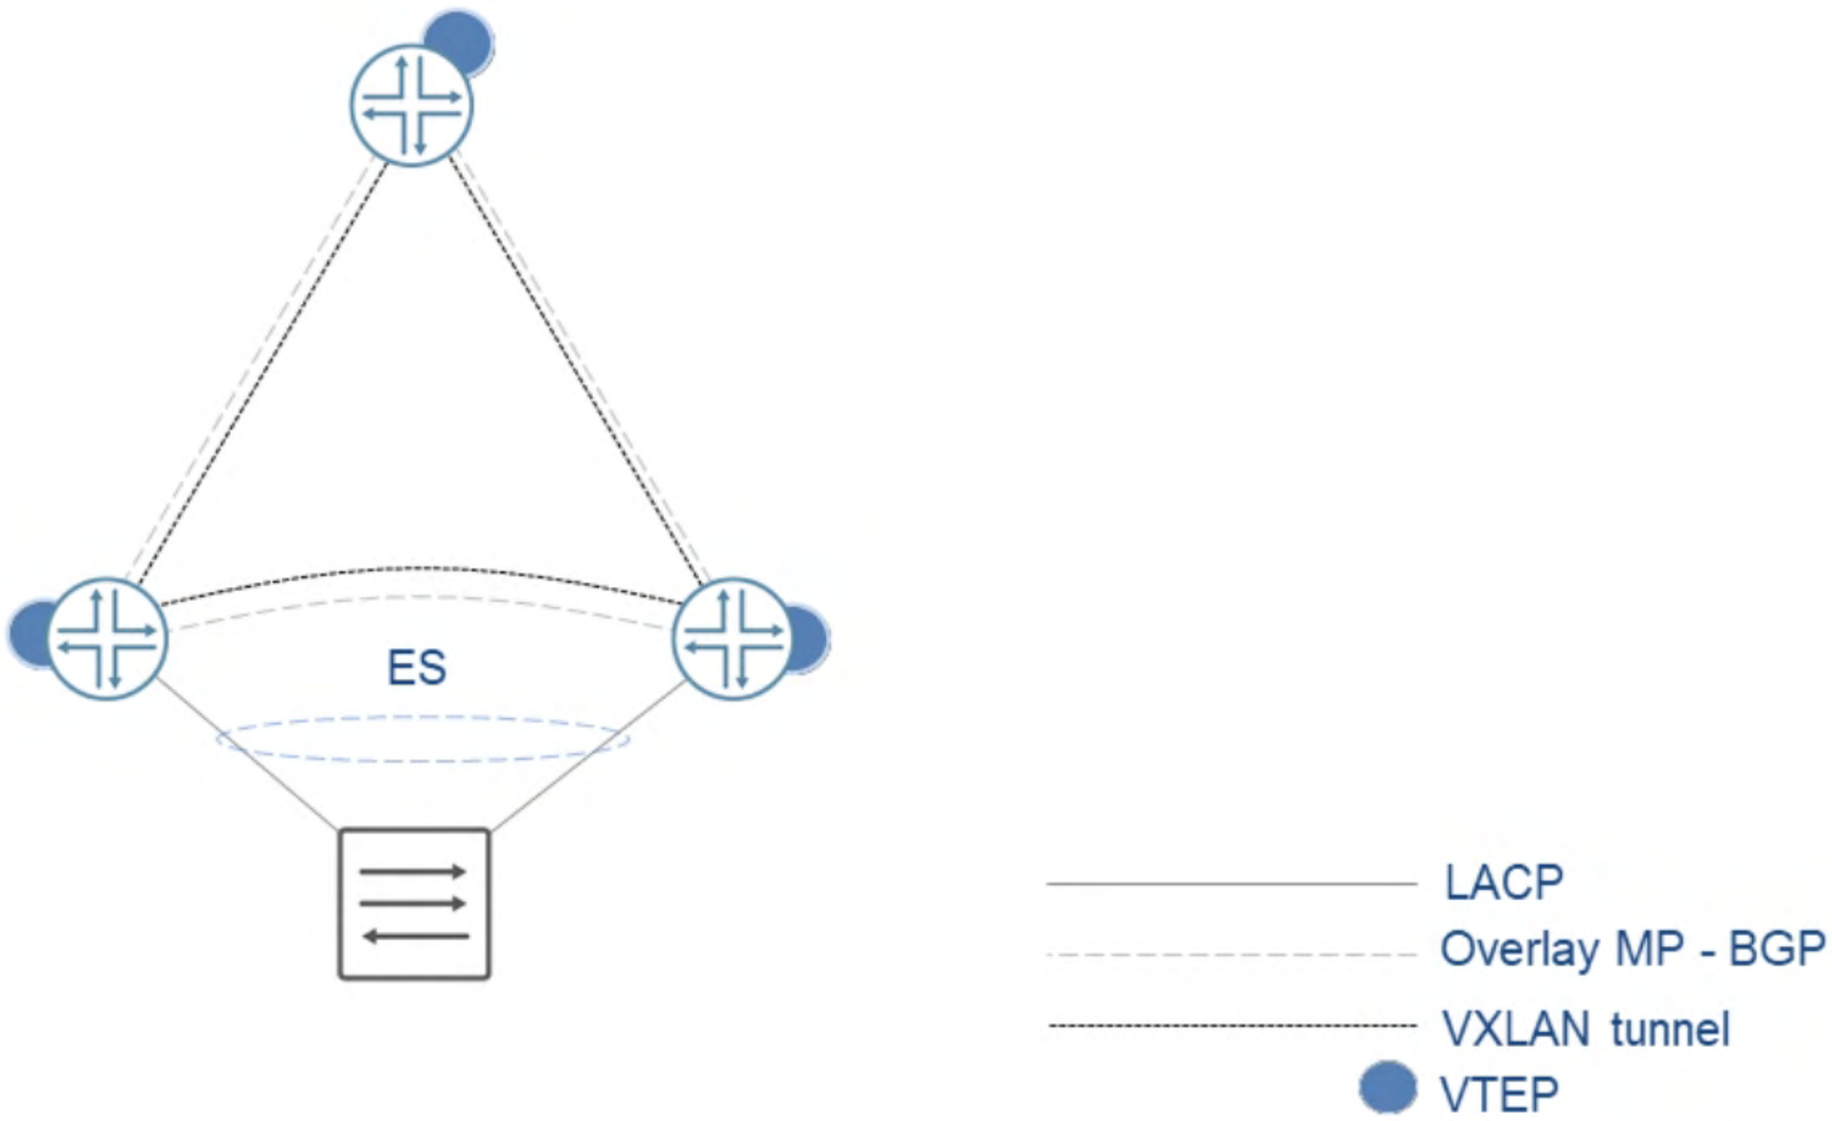

From the overlay network point-of-view, the LAG is spread out across 2 different physical devices. Using EVPN, both devices signal to the rest of the network that they are connected to one and the same Ethernet Segment (ES). When the ES is configured as an active-active ES, both links can be used to forward traffic. To prevent traffic loops, only the PE acting as the Designated Forwarder (DF) will forward BUM traffic. The DF is elected per broadcast domain using the EVPN Ethernet Segment route.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: A Basic EBGP-based Overlay Network

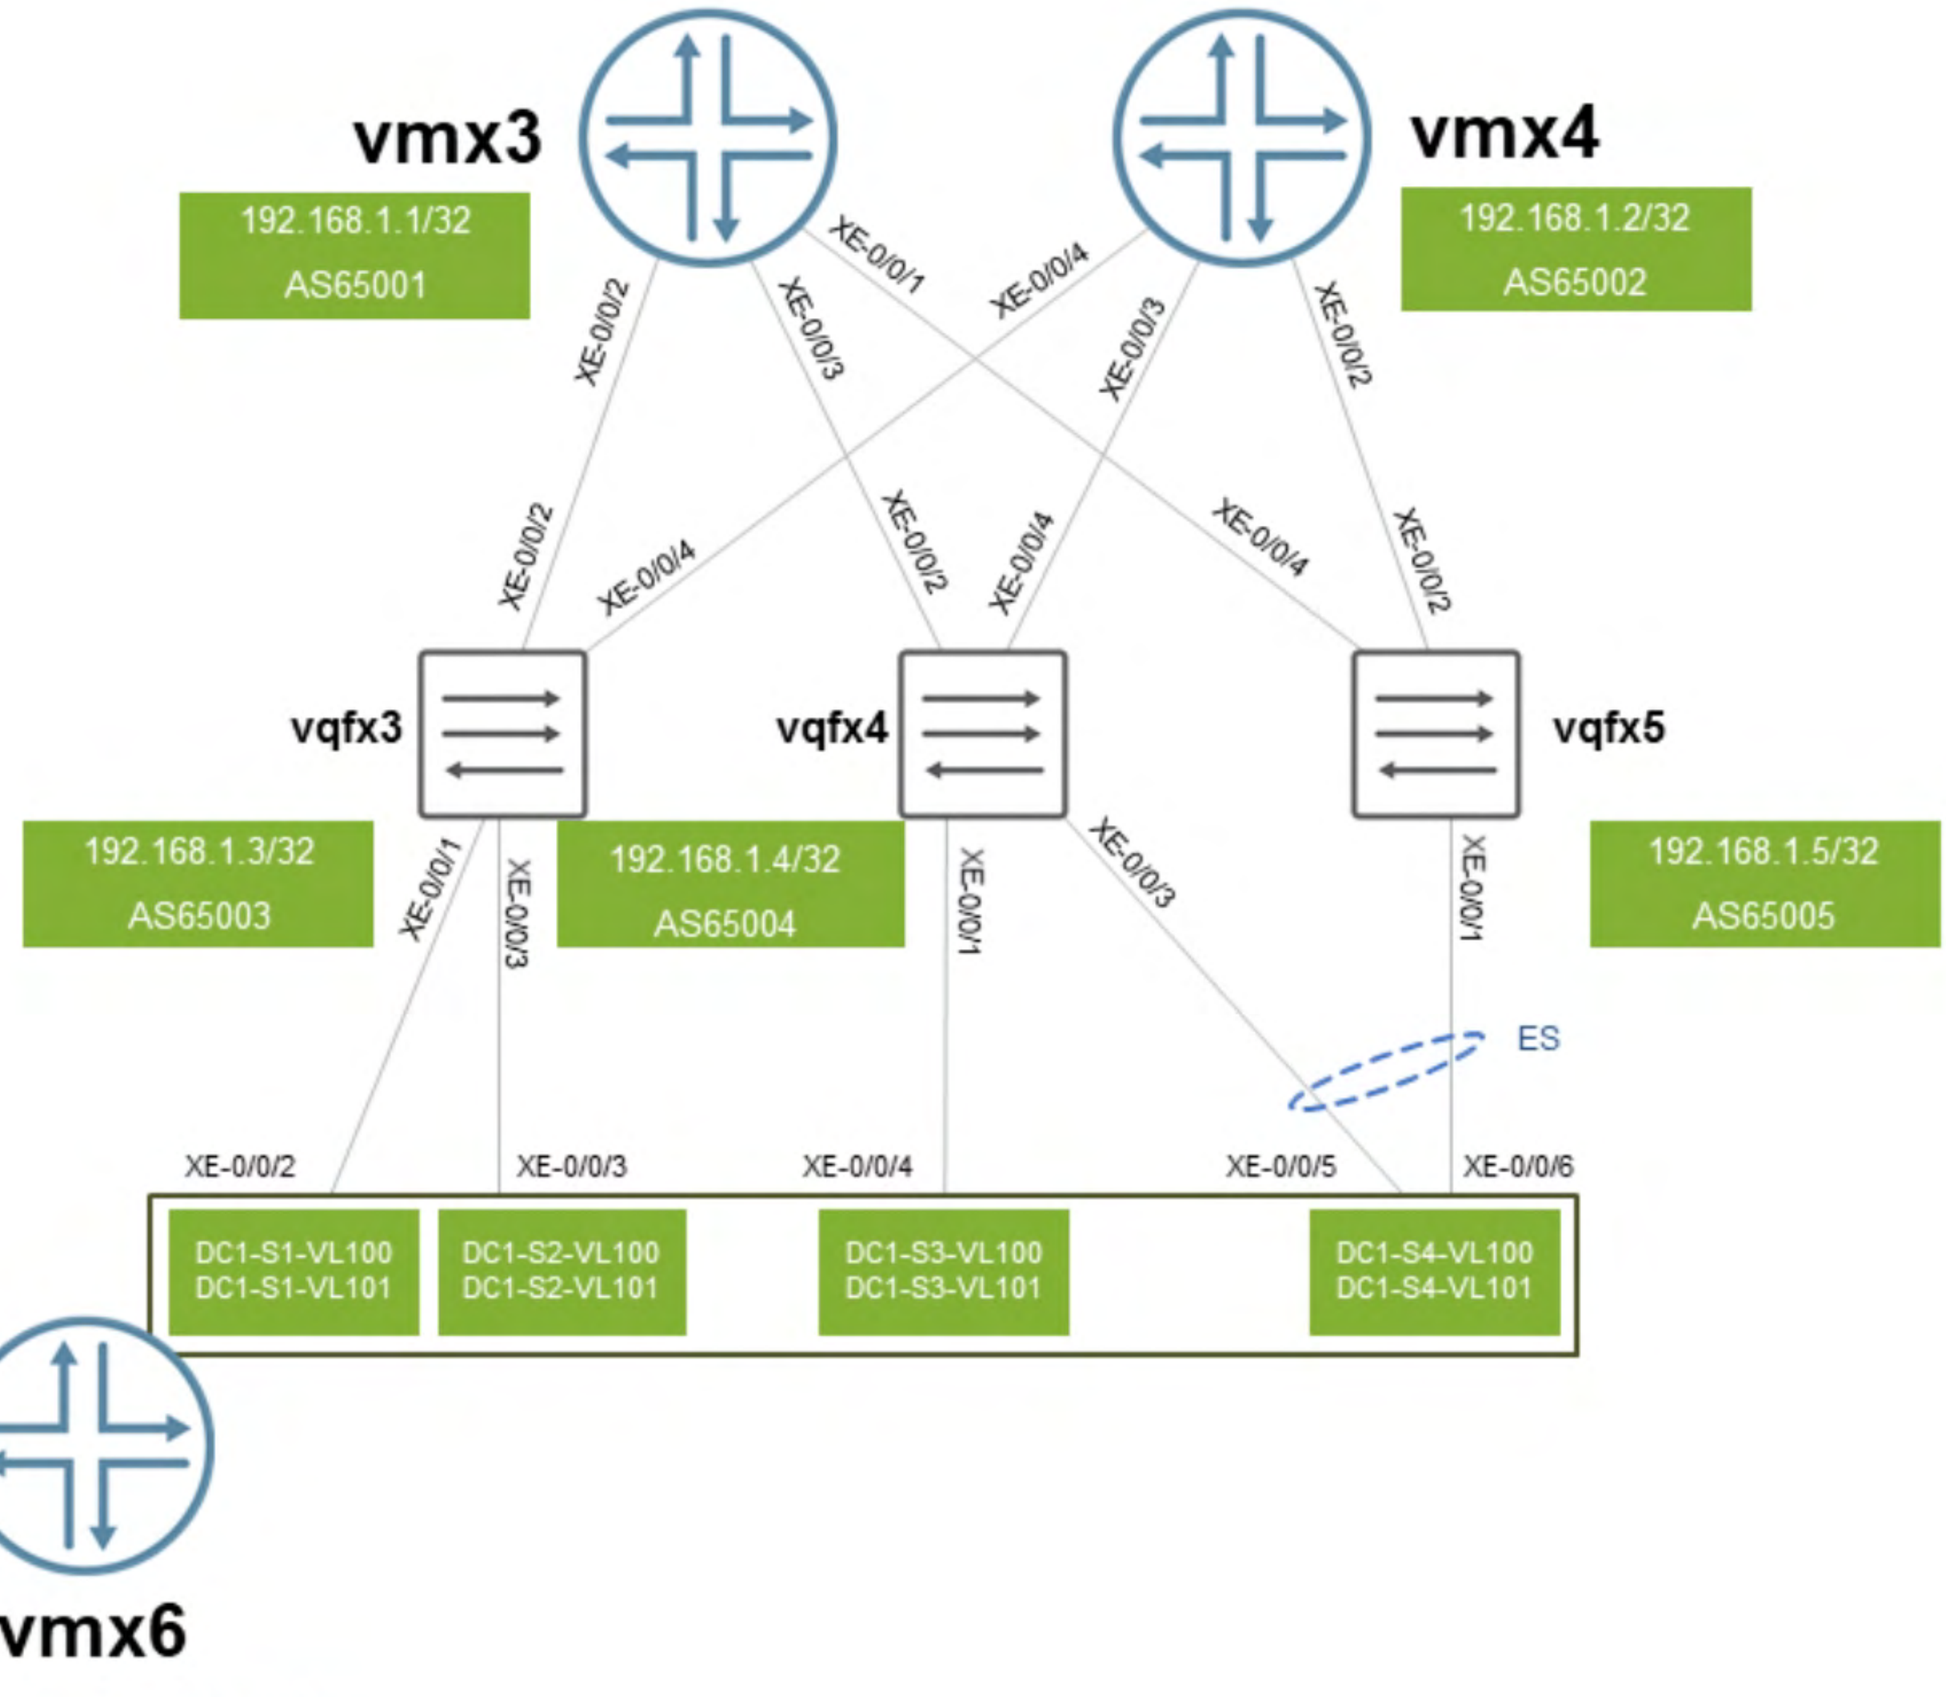

Figure 6 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

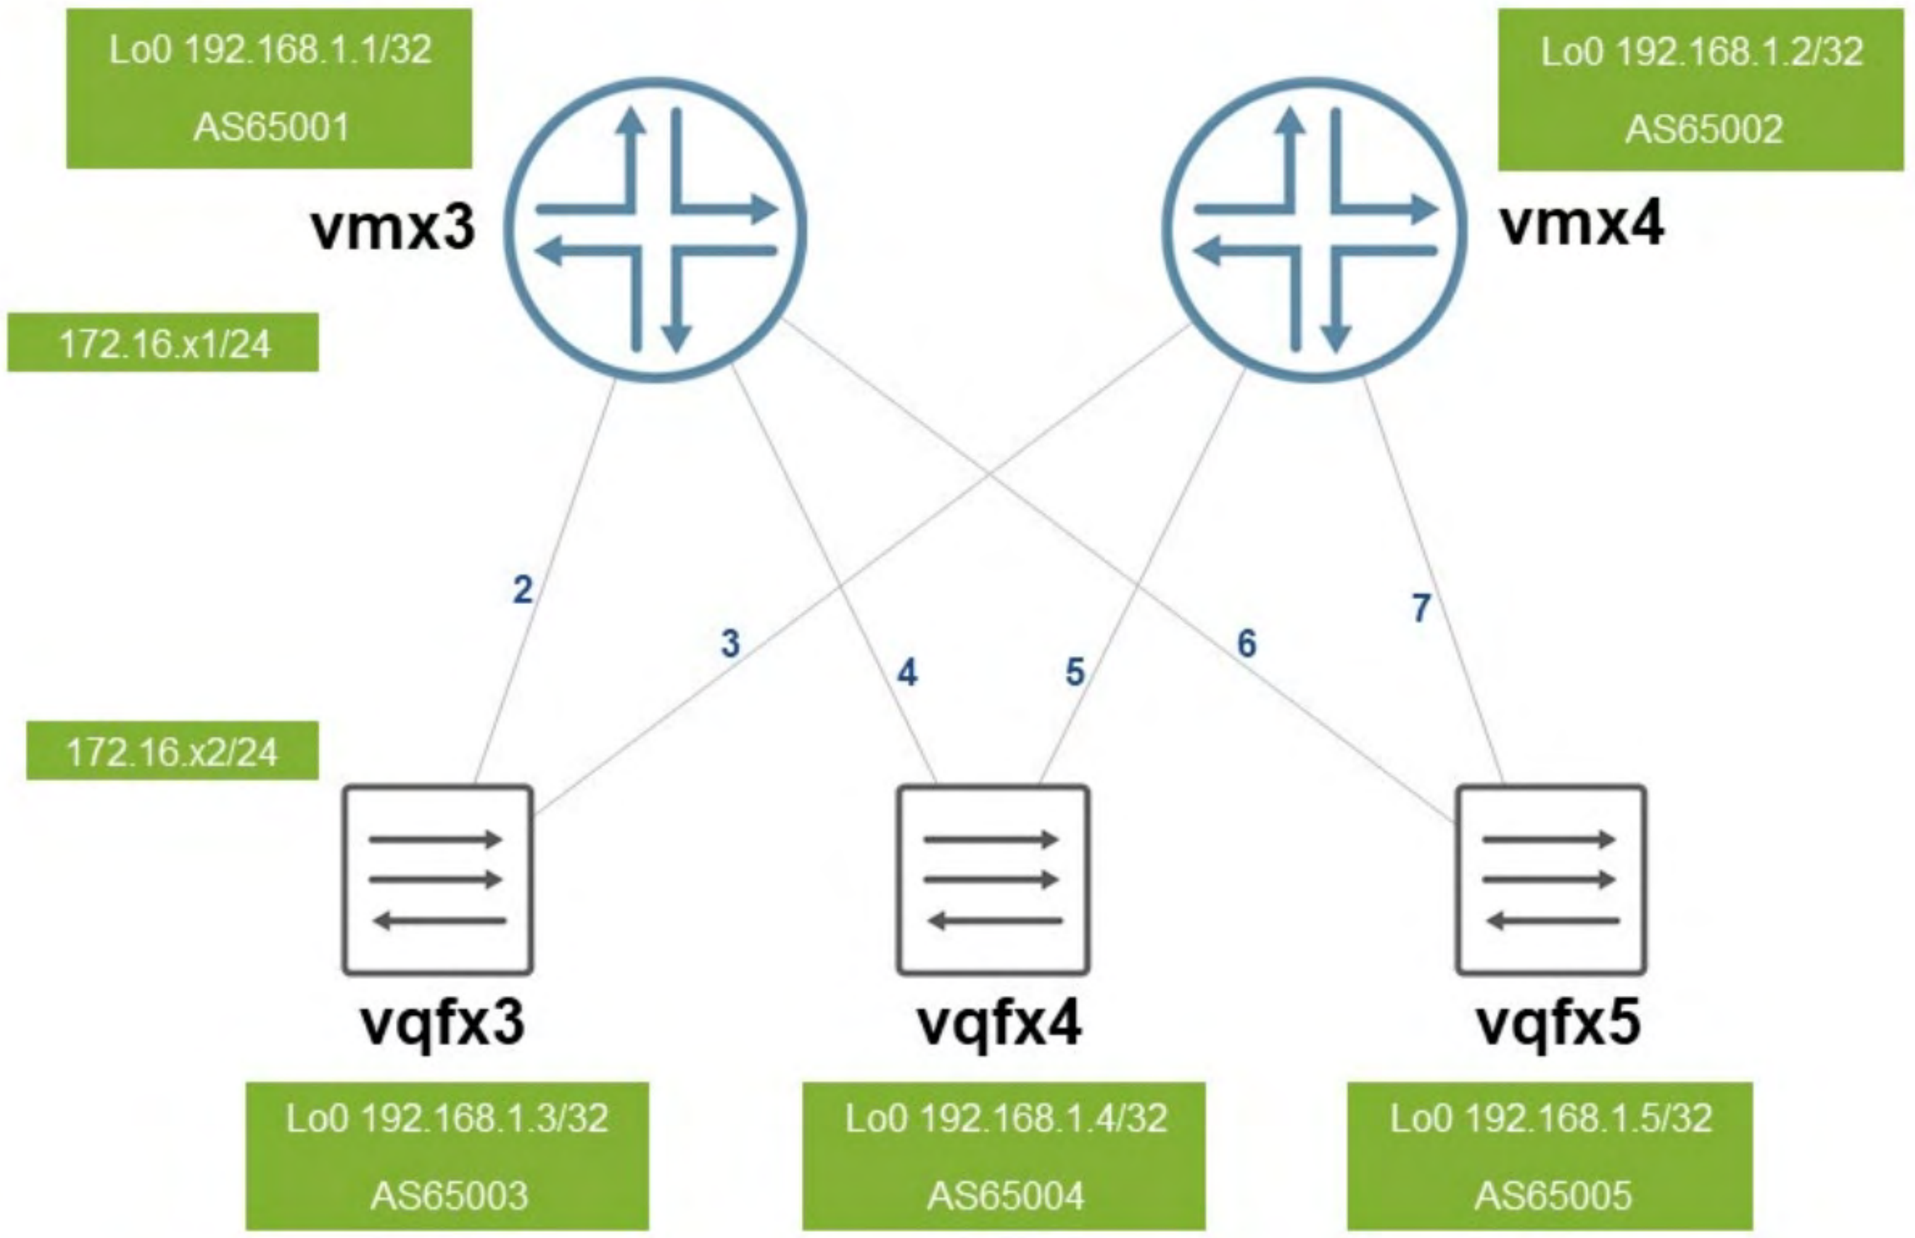

Figure 2 - Logical topology

Note: the starting topology is configured with IP-addressing only.

Task 3.1: Configure the EBGP Underlay

- Configure the underlay network according to the information provided in the logical topology.
- Ensure IP-connectivity between all of the loopback IP addresses.
- Make sure all BGP sessions are 'Established' and that sessions state changes are logged to syslog.

Task 3.2: EBGP Underlay Optimization

- Make sure that the devices load-balance traffic.
- Authenticate BGP sessions with the MD5-hashed password ‘jncie’.
- Make sure that the loss of a BGP neighbor is detected by BFD in exactly 10 seconds. You are not allowed to change the BGP hold-time.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.3: IBGP Overlay Configuration

- Configure a full-mesh of IBGP sessions sourced from the local loopback interface.
- The IBGP sessions should:
    - use AS65500
    - be authenticated using the MD5-hashed password 'jncie'
    - detect the loss of a BGP peer in 6 seconds
    - only support the exchange of EVPN routes
    - be able to load-balance traffic

Task 3.4: Configure QFX Switch-level EVPN

- Use ingress-replication and VXLAN encapsulation
- Use Lo0 as the VTEP source interface
- Use route-target 1:1 for the fabric

Task 3.5: VXLAN Segment Configuration

- Configure a VXLAN segment for S1, S2 and S3 on vQFX3 and vQFX4.
- Traffic to and from the servers should be tagged with VLAN 100.
- Use route-target 100:100 for the VNI.

Task 3.6: Active-Active Multihoming

- Configure a VXLAN segment for S4 on vQFX4 and vQFX5.
- Configure the VXLAN segment as a single Ethernet Segment
- Use the VLAN and route-target that was used in the previous assignment.

Task 3.7: Configure vMX Virtual-Switch-Level EVPN

- Configure a VLAN-aware EVPN service on both vMX3 and vMX4.
- Use ingress-replication and VXLAN encapsulation.
- Use Lo0 as the VTEP source interface.
- Use route-target 1:1 for the fabric.
- Import the VNI that is active on the leaf switches.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.8: Redundant Layer 3 VXLAN Gateway Configuration

- Configure a layer 3 gateway for the VXLAN segment offered to servers located in VLAN 100 on both vMX routers
- Use IP 10.200.100.252/24 on vMX3 and IP 10.200.100.253/24 on vMX4. Configure an IP gateway on the routers for hosts inside VNI 100. Use IP address 10.200.100.254 for this.

Task 3.9: Inter-VXLAN Routing

- Configure an additional VXLAN segment for all servers in VLAN 101 using the following parameters:
    - VLAN 101
    - route-target 101:101
    - Subnet 10.200.101.0/24 (10.200.101.252 on vMX3, 10.200.101.253 on vMX4 and 10.200.100.254 as VIP)
- Add the VXLAN segment on all server-facing interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: IBGP-based Overlay Network

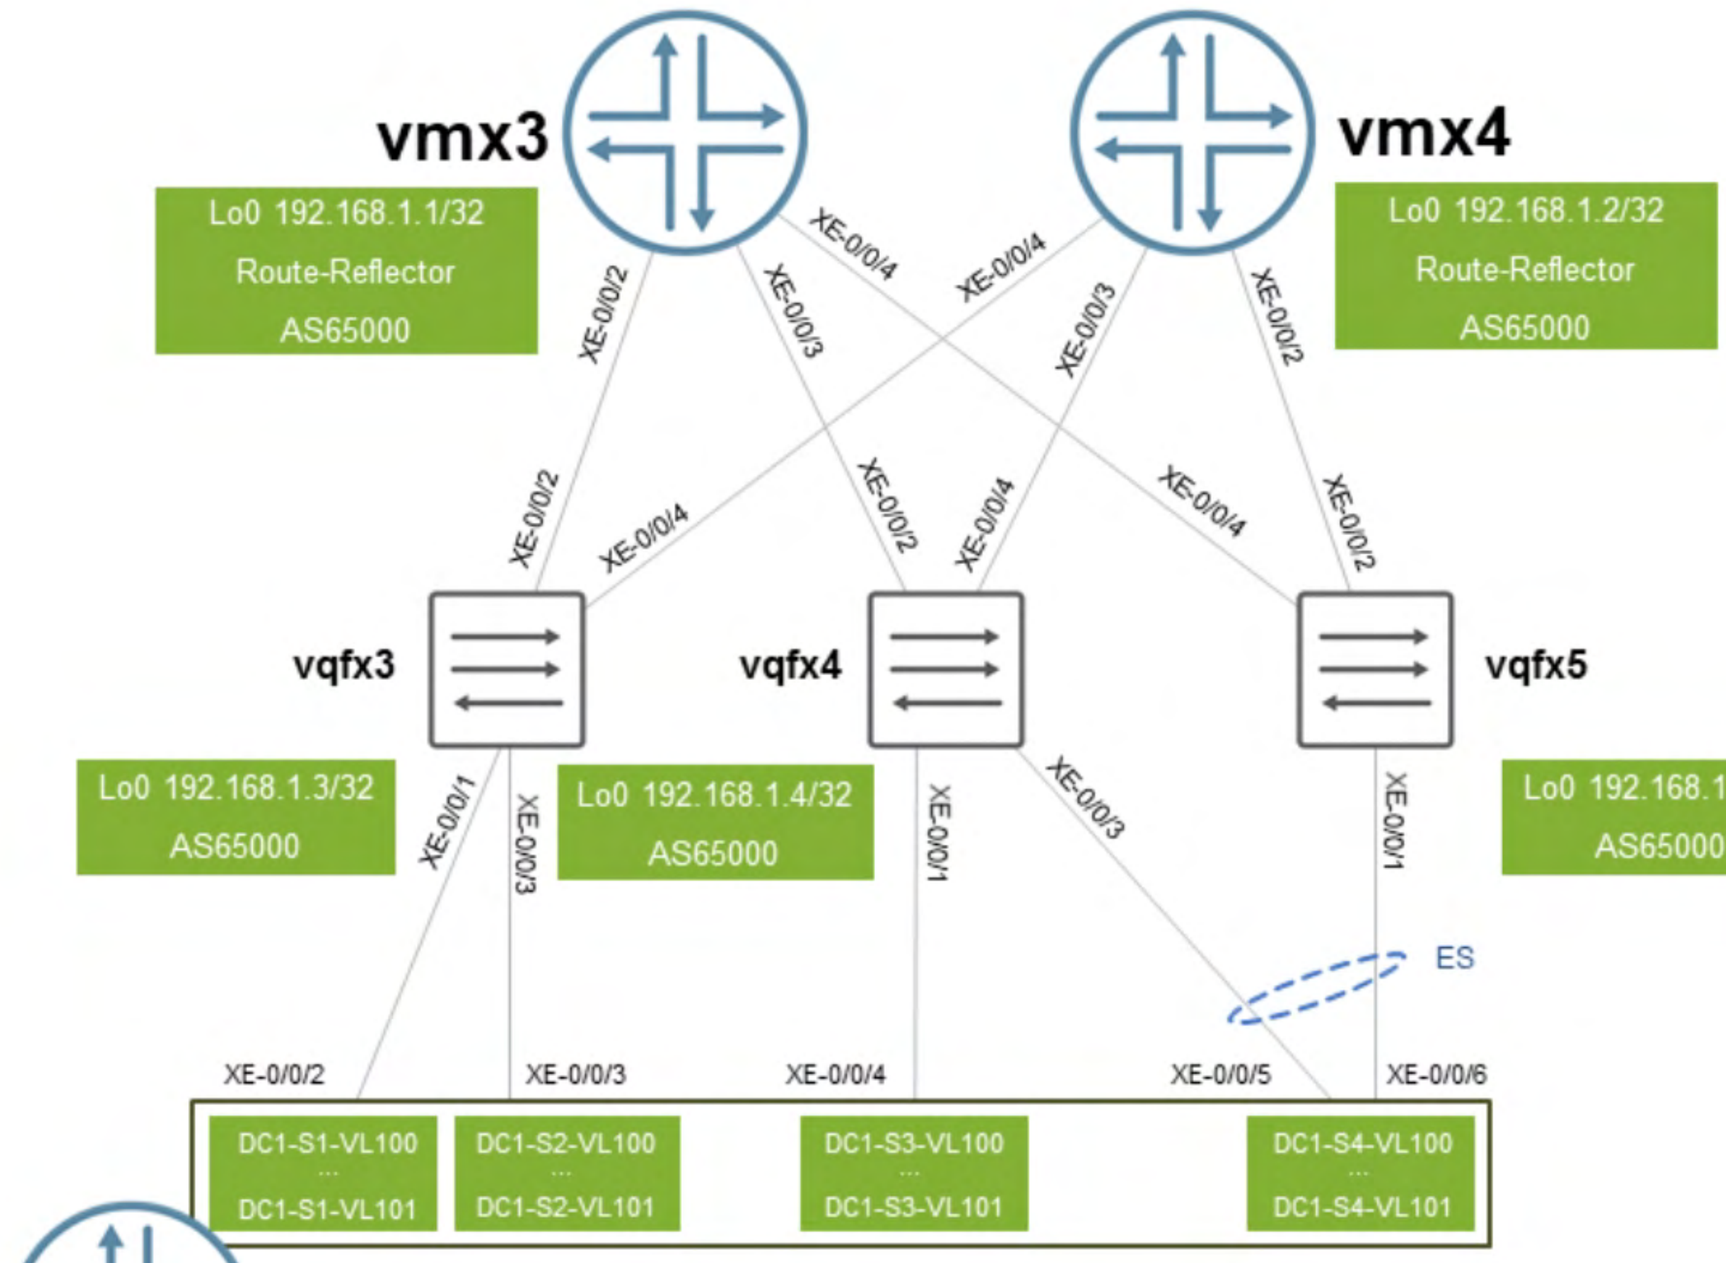

Figure 7 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

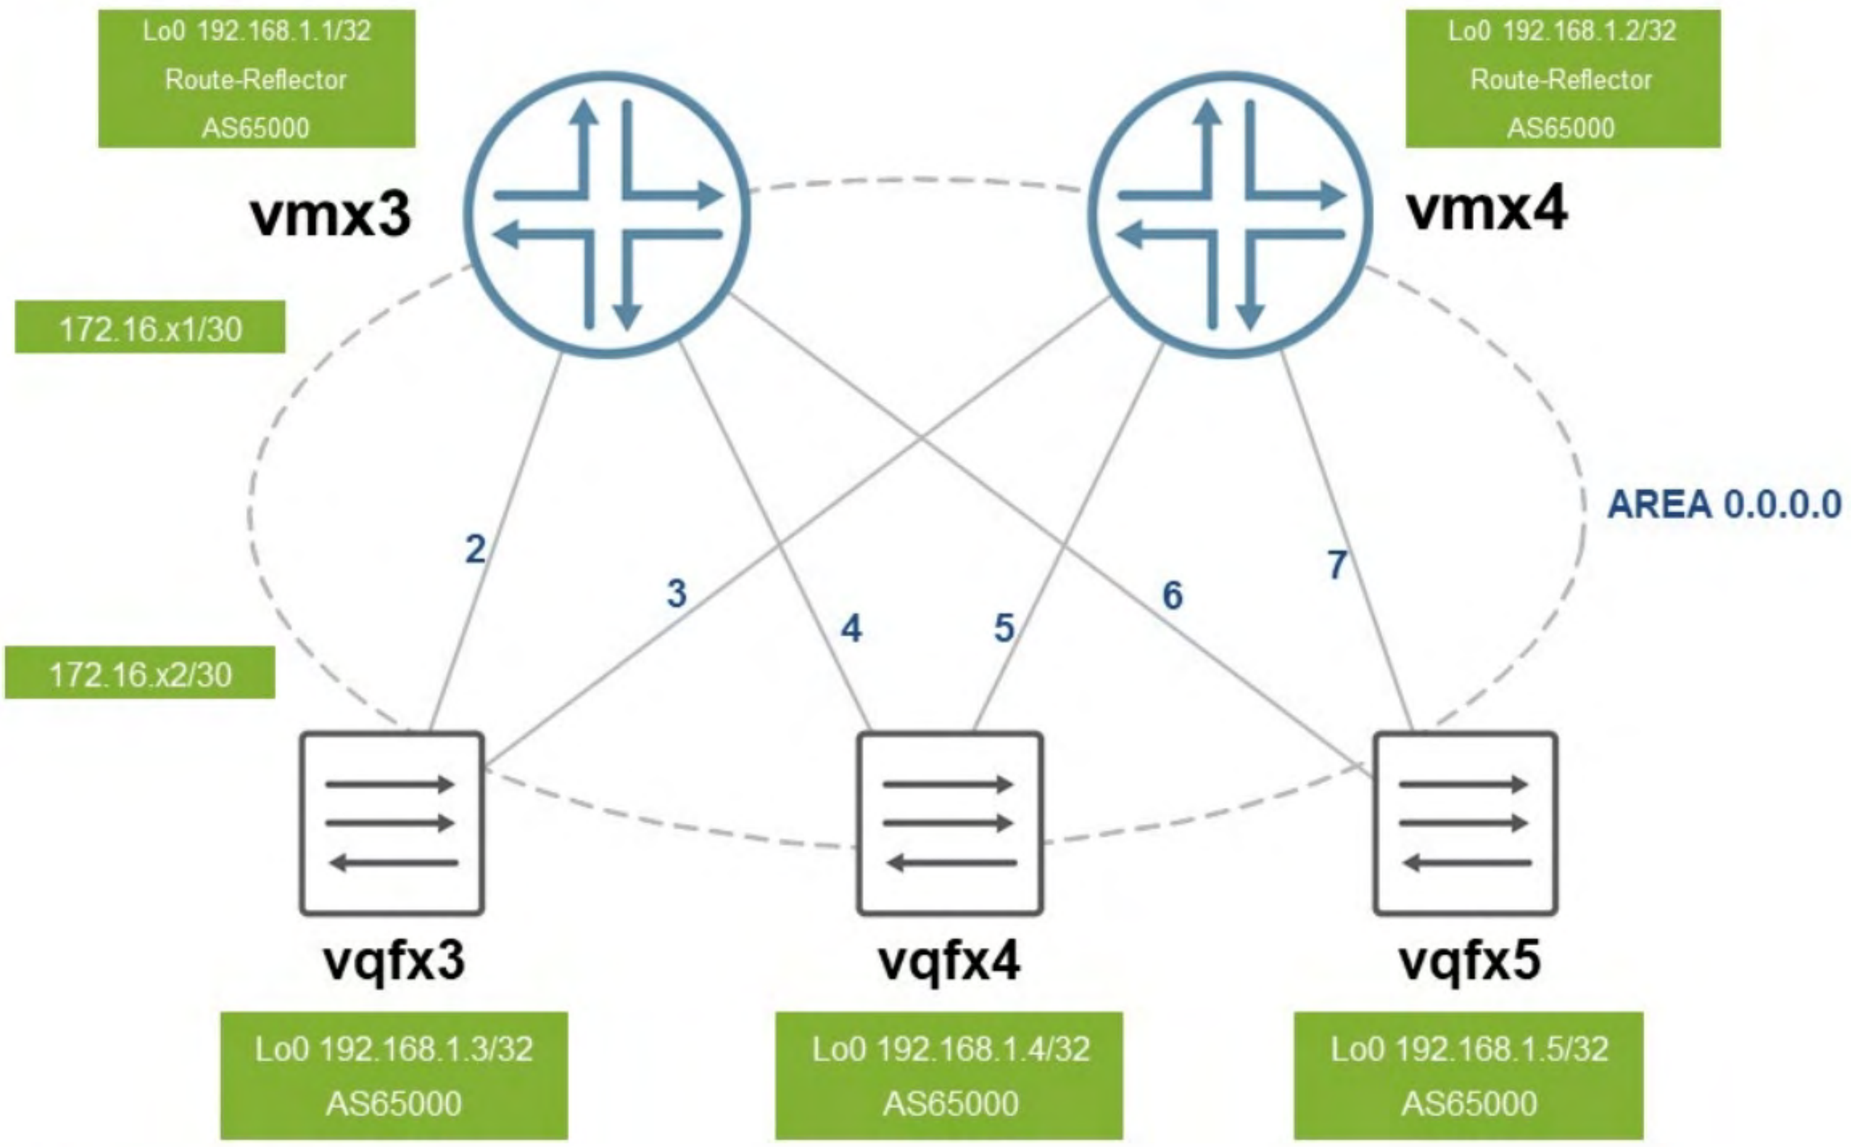

Figure 2 - Logical topology

Note: the starting topology is configured with IP-addressing only.

Task 3.10: IGP Configuration

- Configure a single area OSPF the devices listed in the topology.
- Ensure IP-connectivity between the loopback IP addresses.

Task 3.11: IGP Optimization

- Make sure traffic will be load-balanced across equal cost paths.
- OSPF adjacency loss should be detected within 10 seconds.
- Ensure that no type 2 LSAs are found within Area 0.0.0.0.
- Authenticate OSPF using MD5 authentication. Use the password ‘jncie’.

Task 3.12: IBGP Configuration

- Configure IBGP for AS65000 based on the information provided in the logical topology.
- Ensure that:
    - vMX3 and vMX4 function as route-reflectors

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- the loss of a BGP neighbor is detected within 10 seconds
- BGP sessions only support the exchange of EVPN routes
- BGP sessions are authenticated using the MD5-hased password 'jncie'

Task 3.13: Configure QFX Switch-Level EVPN

- Use ingress-replication, VXLAN encapsulation and Lo0 as the VTEP source interface.
- Configure the QFX to use route-target 1:1 for the overlay.
- VXLAN segments should carry a target community that differs from the fabric community.
- Prepare the overlay to allow the import of VXLAN segments that do not have explicitly defined target communities.

Task 3.14: VXLAN Segment Configuration and Active-Active Multihoming

- Enable layer 2 communication for devices in VLAN 100 using VNI 100.
- Do not explicitly configure the VXLAN segments with a target community.
- Enable active-active multihoming on vQFX4 and vQFX5. Assume that the connected device does not support LACP.

Task 3.15: VXLAN Segment Configuration for VNI 101

- Enable layer 2 communication for devices located in VLAN 101.

Task 3.16: Additional VXLAN Segment Configuration

- Enable layer 2 communication for devices located in the following VLANs:

VXLAN / VNI	**Route Target**
102	
103	
104	
105	target:65000:105
106	target:65000:106
107	target:65000:107
108	target:65000:108
109	target:65000:109
110	target:65000:110

- Only configure route-targets for VXLAN segments when they are mentioned.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.17: MX Virtual-Switch with Selective VXAN Import

- Import the VXLAN segments 100, 101, 102, 103 and 104 into a routing-instance called 'fabric'.
- Make sure that no other address information is placed into the newly created instance.

Task 3.18: Layer 3 VXLAN Gateway Configuration

- Configure the following layer 3 gateways on vMX3 and vMX4:

VXLAN / VNI	Gateway
100	10.200.100.254/24
101	10.200.101.254/24
102	0.200.102.254/24
103	0.200.103.254/24
104	0.200.104.254/24

- Configure vMX3 and vMX4 with the IP addresses that precede the virtual-gateway-address.

Task 3.19: Layer 3 VXLAN Gateway Configuration

- Configure the following layer 3 gateways on vMX3:

VXLAN / VNI	Gateway
105	10.200.105.254
106	10.200.106.254

- Configure the following layer 3 gateways on vMX4:

VXLAN / VNI	Gateway
107	10.200.107.254
108	10.200.108.254

- Only import VXLAN segments into the routing-instance when the vMX has a gateway configured for that VXLAN segment.
- Make sure there is IP connectivity between the different subnets

Task 3.20: Layer 3 VXLAN Gateway Configuration

- Configure the following anycast gateways on vMX3 and vMX4:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

VXLAN / VNI	Gateway
109	10.200.109.254
110	10.200.110.254

Part 3: IP Fabric and Controllerless Overlay

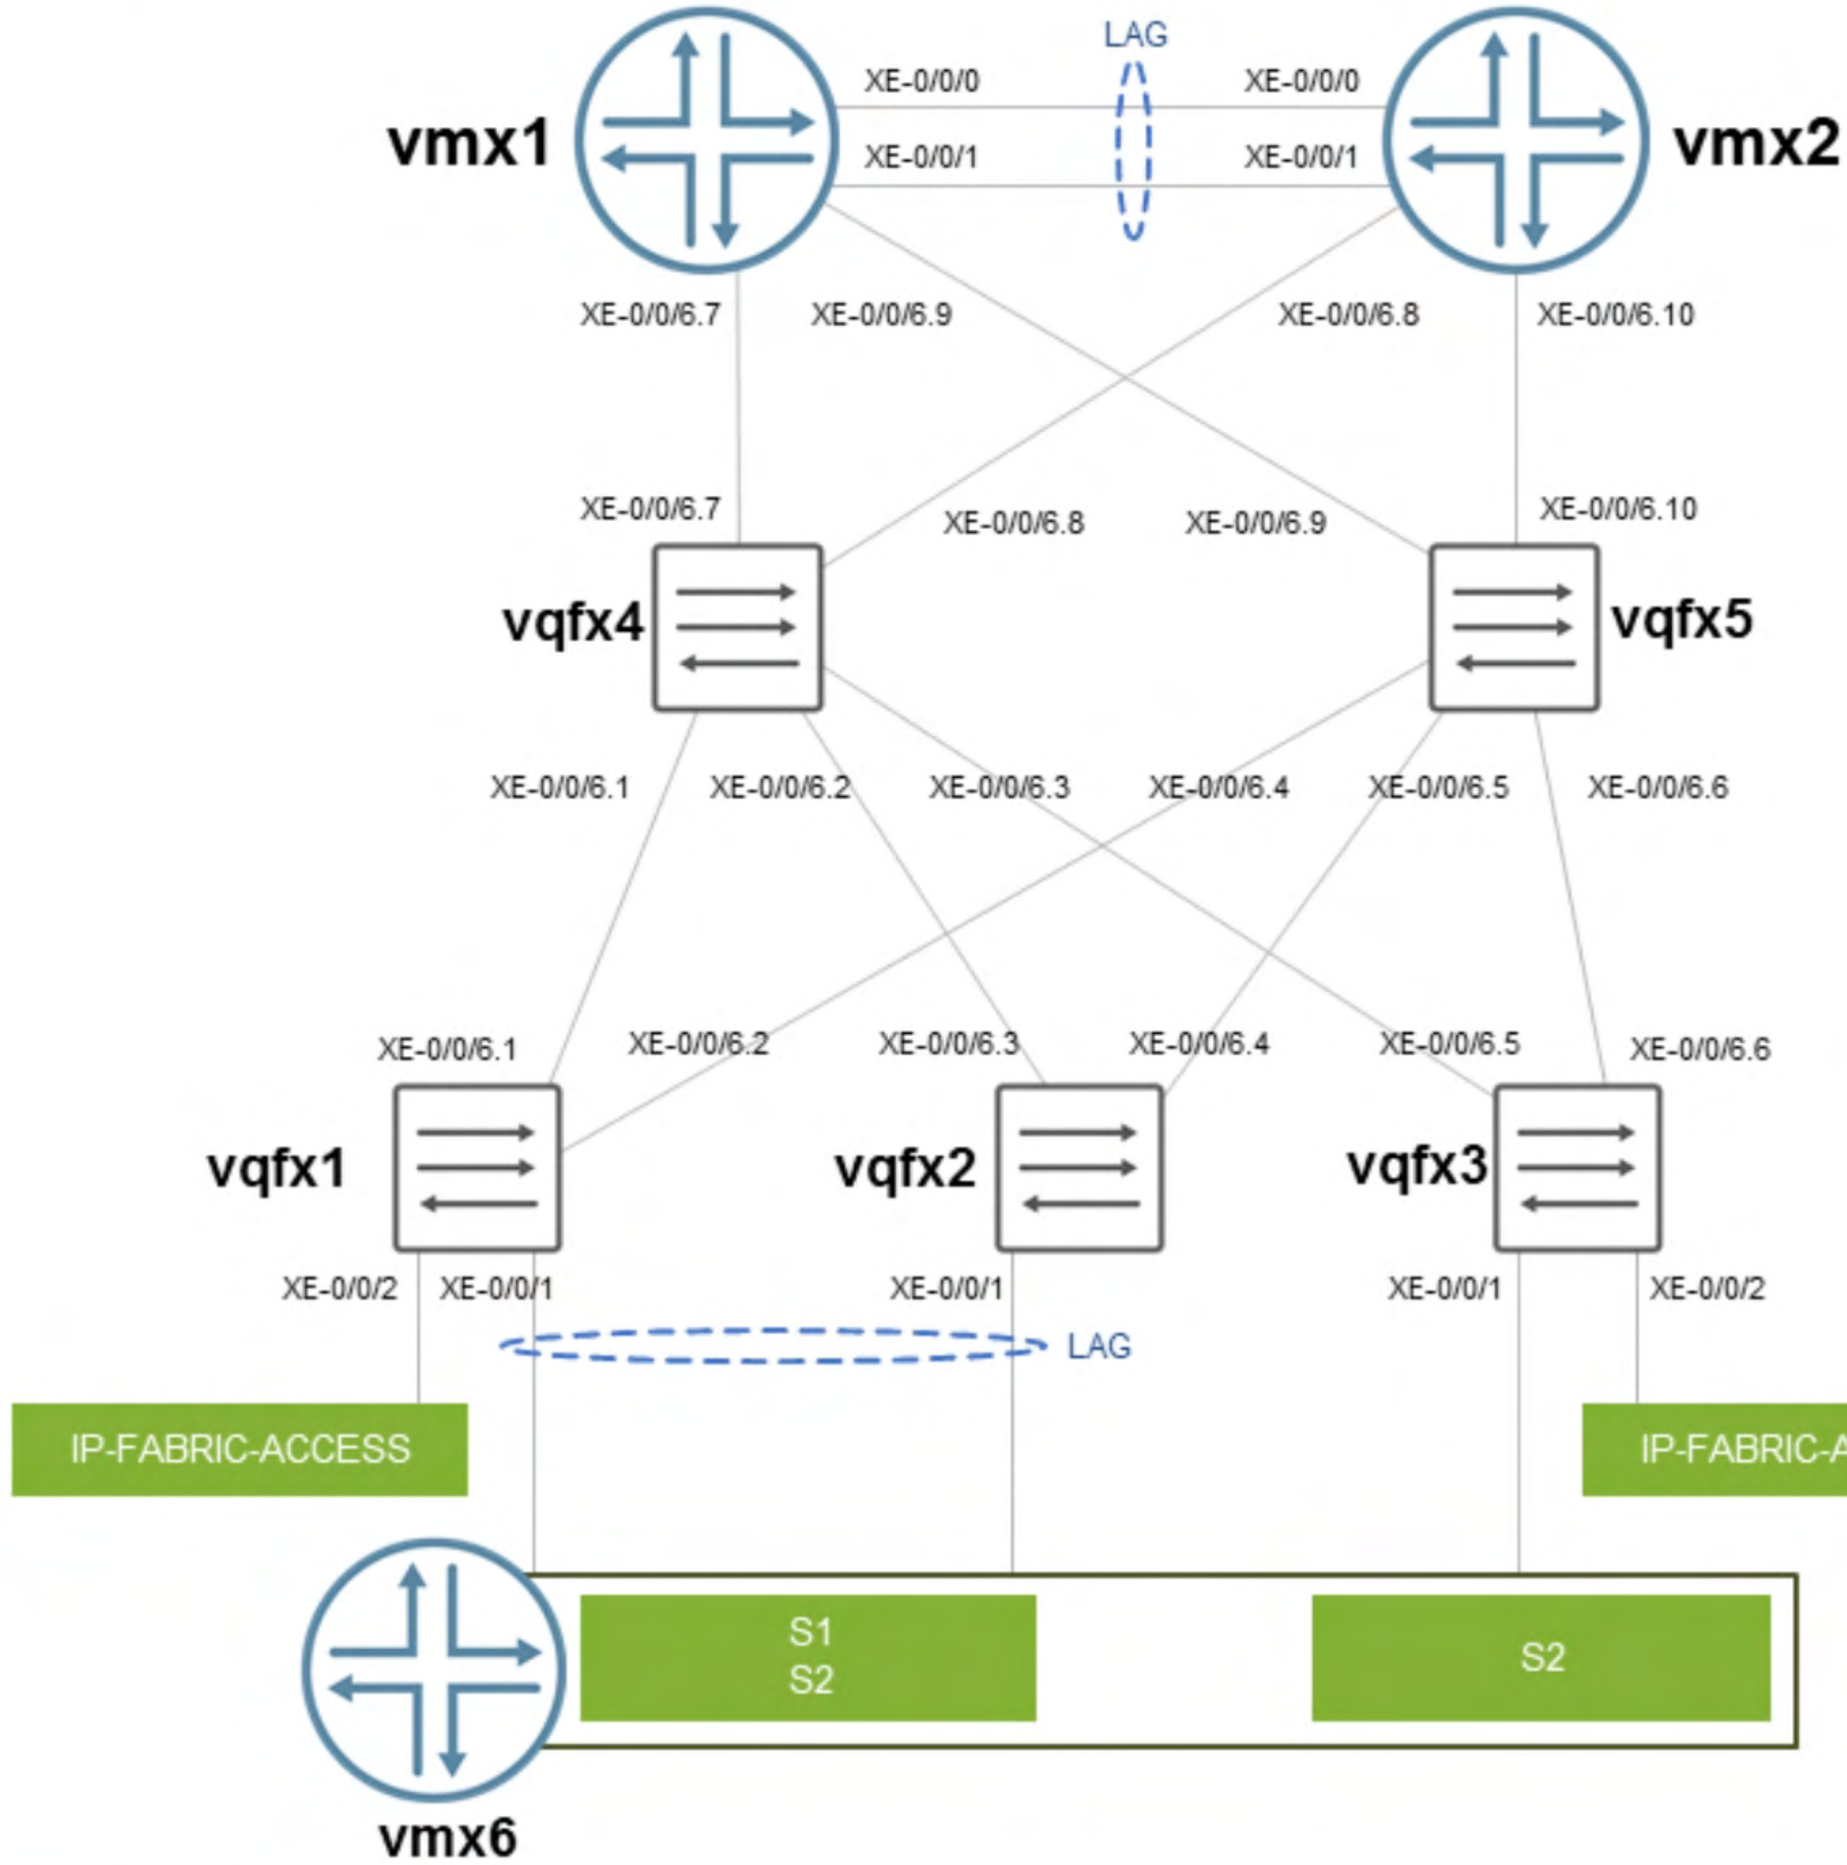

Figure 8 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

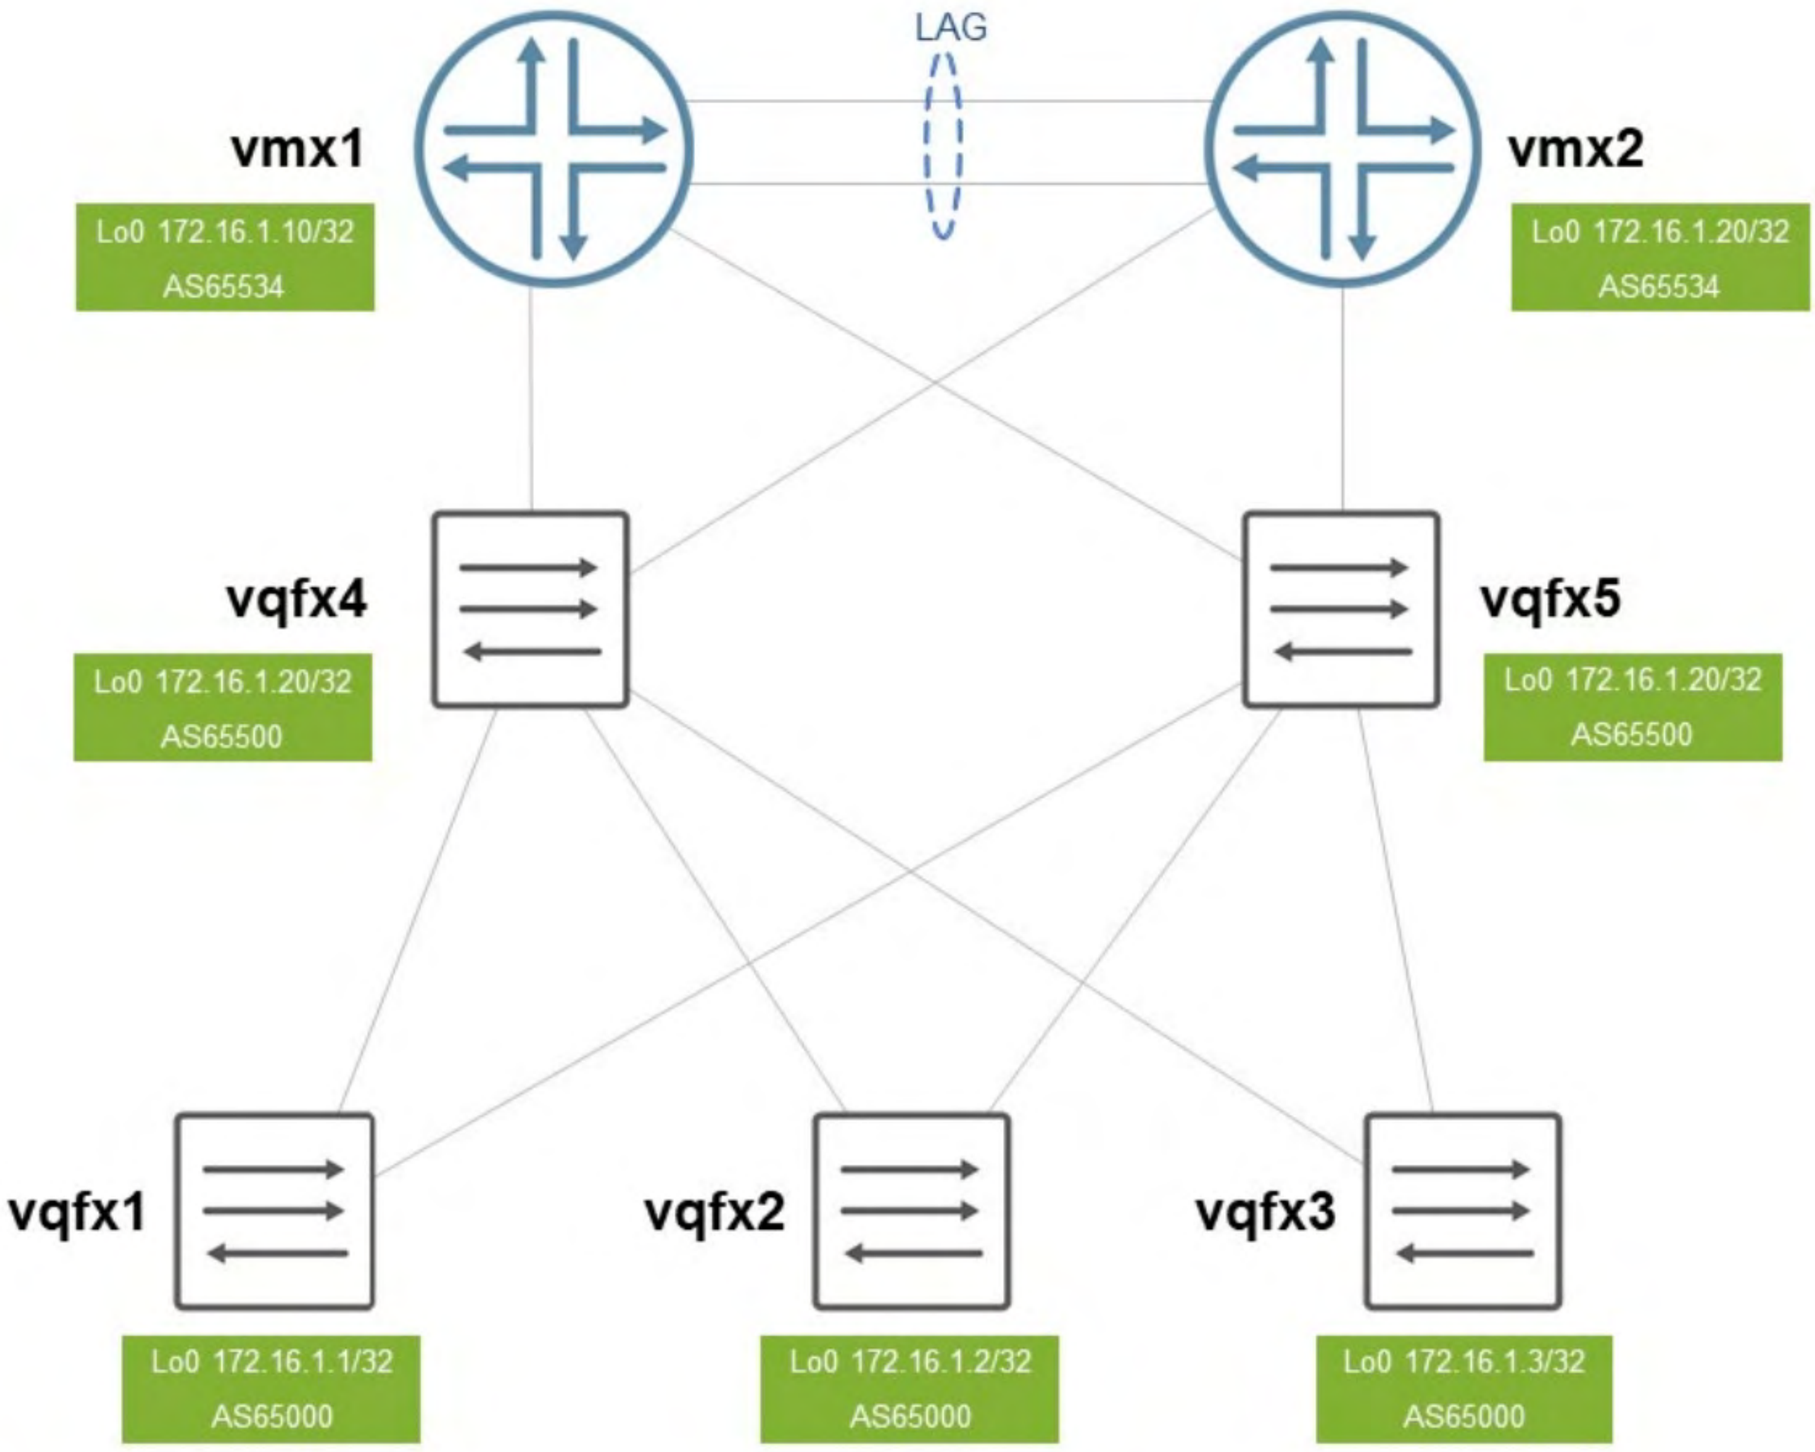

Figure2 - Logical topology

Note: starting topology is configured with IP-addressing only.

Task 3.21: EBGP for the IP-Fabric

- Configure the devices displayed in the topology to form an EBGP-based IP-Fabric.
- Members of the IP-Fabric should advertise their own loopback IP address.
- Configure the same export policy on every device. Do not use the 'interface' statement in any of the match conditions.
- Use the autonomous systems displayed in the logical topology.
- Members of the fabric that have equal-cost paths towards other nodes should utilize all available links.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.22: Completing the IP-Fabric

- BGP speakers should be able to detect peer loss in 9 seconds without the use of BFD.
- All control plane protocol messages exchanged should use the MD5-hased password ‘jncie’.
- Advertise gateway addresses configured on interface xe-0/0/2 on both vQFX1 as well as vQFX3.
- Whenever an IP address is added to interface xe-0/0/2 on either vQFX1 or vQFX3, it should be advertised without any routing-policy change.

Task 3.23: Controllerless Overlay

- Have vQFX1, vQFX2 and vQFX3 form an overlay. Use the IP-Fabric as underlay network.
- The servers S1 and S2 should be able to communicate inside VLAN 200 and VLAN 300. The servers are expecting tagged frames. It should be possible to tunnel VLANs inside VLAN 300.
- Ensure that the connections to S1 and S2 are redundant.
- New VLANs added to the server-facing interfaces should only require configuration in the [edit vlans] and [edit interfaces] stanza.
- Every VNI should have a different RT.

Task 3.24: Expanding the Controllerless Overlay

- Incorporate vMX1 and vMX2 into the controllerless overlay and have the vMX routers peer with every leaf switch.
- Configure the following default-gateway on the vMX

	IPv4 Gateway	**IPv6 Gateway**		**LAN/VNI**	**`Target**
vMX1	192.168.220.254 (VIP)	2001:100:200::1 (VIP)		400	65534:1
vMX2	192.168.220.254 (VIP)	2001:100:200::1 (VIP)		400	65534:1

- Only import the VXLAN segment on the switches connected to S3.
- Ensure that S3 has IP connectivity to the IRB interfaces configured on the vMX devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.25: IP-Fabric Routing Policies

- Create a policy called 'drain' on both vQFX4 as well as vQFX5. Do not apply the policy but ensure that it can be applied using only a single configuration command. The policy should ensure that, when applied, the vQFX will become less likely to receive traffic compared to other vQFX.
- vQFX3 should send traffic towards vMX2 via vQFX5 and use vQFX4 as backup. The vQFX3 leaf should continue to load share towards vMX1, vQFX1 and vQFX2.
- On vQFX3, you are only allowed to apply a policy at the BGP group level.
- Ensure that the vMX routers can use standard communities to differentiate between the routes that were received from the vQFX leaf switches when using show commands.
- Ensure IP connectivity between the servers with a layer 3 gateway on the vQFX and servers with a layer 3 gateway on the vMX.

Chapter 4: Data Center Interconnect

To interconnect the compute, storage and networking resources located in different data centers, we can use a variety of layer 2 and layer 3 solutions.

At first, layer 2 solutions only included L2-circuits and BGP signaled layer 2 VPNs. These solutions offer a point-to-point layer 2 connection without MAC learning. To be able to offer point-to-multipoint connections between data centers, VPLS was developed. A VPLS can be either BGP signaled or LDP signaled. MAC learning for VPLS is performed in the data plane.

Layer 3 solutions to interconnect data centers have been very consistent the past few years and included various tunneling mechanisms such as IP in IP, GRE, IPsec and MPLS L3VPNs.

The latest addition to the DCI options is EVPN. EVPN offers a control plane based protocol that can work in tandem operation with various encapsulation protocols (GRE, VXLAN, MPLS) to offer layer 2 services. An EVPN can also be combined with an MPLS L3VPN, offering full-mesh VPN connectivity between data centers. This way, data centers are interconnected on both layer 2 as well as layer 3.

Having covered VXLAN in chapter 3 already, we will discuss MPLS and label distribution protocols, MPLS L3VPN and finally EVPN using MPLS as an encapsulation strategy. We will conclude this chapter by examining the different datacenter interconnect strategies that EVPN has to offer.

MPLS and Label distribution protocols

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

MPLS was originally limited to ISP networks but has been making an entry into both datacenter as well as enterprise networks for quite some time now.

In an MPLS network, Label Switched Paths (LSPs) are signaled between MPLS routers, called Label Switched Routers (LSRs). The LSPs are unidirectional and can be used to switch traffic across the network. LSPs are signaled using a label distribution protocol. The two dominant protocols in use today are the Label Distribution Protocol (LDP) and the resource reservation Protocol (RSVP).

An LSR performs a label lookup to determine which interface is to be used as the outgoing interface. At the edge of the network, the ingress LSR pushes a label onto a packet between its layer 2 and layer 3 header. This is called a 'push' operation. Intermediate LSRs along the LSP will use this label to do a lookup to find the proper outgoing interface and 'swap' the label. On egress, the MPLS label is removed from the packet in an operation referred to as a 'pop' operation. In case of Penultimate Hop Popping (PHP), this operation is performed by the second to last LSR of the LSP. PHP can optimize label operations by avoiding a double lookup on the egress PE. With PHP, the egress PE only has to do an IP lookup (when the payload is IP). Without PHP, the egress PE would have to pop the label and then perform an IP lookup.

Label distribution protocols

LDP was developed to offload the forwarding plane and to build a full mesh of LSPs. LDP signaled LSPs do not offer any traffic-engineering capabilities and always follow the IGP's best path. All LDP routers exchange multicast messages to form a neighbor relationship. As soon as this relationship is formed, the routers will try to set up an LDP session over a TCP connection. This session is then used to exchange labels. These labels are exchanged on a hop-by-hop basis. By default, Junos advertises a label for its lo0 address and for every label-mapping that was learned from other LDP neighbors.

LDP routers will use label mappings for prefixes only when the mapping is received on the interface that corresponds to the IGP best path. Through policy configuration, it is possible to make an LDP router inject label mappings for prefixes in addition to its lo0 interface.

Compared to LDP, RSVP offers additional complexity, faster traffic restoration and more control over traffic flows. Routers configured for RSVP will form neighbor relationships with each other. As soon as LSPs are configured, RSVP routers will start to signal LSPs. By default, RSVP signaled LSPs will only install the /32 that corresponds to the egress point of the LSP into the inet.3 table.

The RSVP signaling protocol offers a plethora of ways for administrators to define what path an LSP should take. RSVP can use a database called the TED (Traffic Engineering Database). This database, filled by OSPF or IS-IS, can contain additional information about the links in the network. CSPF, an SPF similar to that of other link-state protocols,

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

can use the information from the TED (available bandwidth, link colors, separate TE-metrics) offering network administrators additional knobs to control traffic in a dynamic way.

Additionally, an explicit route object (ERO) list can be defined by configuring each hop from the ingress LSR to the egress LSR. Two types of hops can be defined in these ordered sequences: strict hops and loose hops. A strict hop has to be an IP address of a directly connected neighbor interface per hop along the path that the LSP should take. A loose hop can be anywhere between the ingress and egress point of the LSP. With loose hops, the actual path between the local node and the loose hop will be calculated by the IGP.

MPLS L3VPN

With MPLS L3VPN, the MPLS infrastructure can be used to offer a VPN to tenants or services. MPLS L3VPN traffic can be switched across a common MPLS infrastructure by using two labels, a bottom label and a top label. The top label is the label that has been distributed by a label distribution protocol. It is used by LSRs to switch traffic towards the correct egress LSR. The bottom label (aka VRF label) defines the VPN membership of a packet. The ingress LSR pushes this label onto a packet so that the egress LSR knows what VPN to associate the traffic with. The intermediate LSRs that switch MPLS L3VPN traffic do not need to be VPN-aware, they only need to look at the top label when forwarding traffic.

The exchange of VPN routing information across the MPLS infrastructure can be handled by the edge routers of the network. The edge routers exchange VPN related address information (and the VRF label) through the use of MP-BGP. This is possible after the BGP sessions have been enabled for the proper address families ('inet-vpn unicast' for IPv4 and 'inet6-vpn unicast' for IPv6).

The two most important concepts to understand about the VPN routing exchange through the use of MP-BGP are the Route-Distinguisher (RD) and the Route-Target (RT). Every L3VPN can be configured with its own address space. The subnets in use can even overlap. This is made possible through the use of the RD. Every VPN prefix has an RD prepended to it to in order to make the address unique.

The RT is an extended BGP community that defines VPN membership. When routers exchange VPN routing information, routes will have an RT attached to them. On import, the value of the RT is used to determine into what VPN the route should be installed.

The edge router of the MPLS L3VPN can be used to offer a gateway to an endpoint, but it can also be used to extend the VPN to include non-MPLS routers. This can be done using static or dynamic (OSPF, RIP or BGP) routing.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Enabling the infrastructure for MPLS L3VPNs requires the following steps:

1. Configuring an IGP on the LSRs to support both a label distribution protocol and to provide for IP-connectivity for BGP sessions.
2. Establishing a full-mesh of LSPs between the edge LSRs through the use of LDP or RSVP.
3. Enabling MP-BGP, with the proper address families, between the PE routers.
4. Configuring a routing-instance of the type 'vrf' on the PE node for every VPN.
5. (Optional) Configure static or dynamic routing between the edge LSR and the devices connected to the MPLS L3VPN.

MPLS EVPN

In chapter 3, we used EVPN signaling together with VXLAN encapsulation. In this chapter, we will delve further into EVPN using MPLS as the encapsulation mechanism:



MPLS EVPN opens the possibility to route MAC addresses across an MPLS core. Within one and the same Ethernet VPN, MP-BGP can use the EVPN address family to signal both layer 2 as well as layer 3 information.

When MPLS is chosen as the encapsulation mechanism, an EVPN can enjoy the benefits that MPLS has to offer. When RSVP is used as a signaling mechanism for labels, the EVPN traffic can reap the benefits of traffic-engineering and traffic protection that RSVP has to offer. And for both LDP and RSVP signaled MPLS networks, it means that the other MPLS applications and benefits can be used on the same MPLS core network.

The majority of the benefits we saw when covering the controllerless overlay still apply when MPLS is used in the data plane. These benefits include the following:

- clean separation of the control plane and the data plane.
- active-active multihoming for both layer 2 as well as layer 3 traffic.
- integrated L2 and L3 solution.
- MAC-mobility across different data centers.
- optimal routing between the data centers for both layer 2 as well as layer 3.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Another important note worth touching is that we can use EVPN for both VLAN-aware and VLAN-bundle services.

VLAN-aware bundle.

Introduced prior to the VLAN-bundle service, the VLAN-aware bundle offers exactly what the name implies. The devices configured with a VLAN-aware bundle service are aware of the different customer VLANs. VLAN-aware bundle services offer a bridge-domain per customer and when using this service, the customer VLANs are bundled into a single instance of the type 'virtual-switch'. Data plane separation between the VLANs is provided through the creation of a dedicated bridge-domain for each VLAN. It is possible to use different VLAN IDs on different interfaces within a single bridge-domain.

VLAN-bundle/VLAN-based EVPN.

With VLAN-bundle services, multiple VLANs are mapped to a single EVI. Because multiple VLANs are mapped to a single EVI and share the bridge table in the MAC VRF, the control plane overhead is significantly lowered. The drawback is that the entire VLAN range has to be permitted and that VLAN translation is no longer an option. The routing-instance has to be configured with 'instance-type evpn'.

Data Center Interconnect types.

We can distinguish different types of Datacenter Interconnects using EVPN and/or MPLS L3VPN. In this chapter, we will discuss the following types of DCI:

- Directly connected using EVPN-VXLAN.
- Internet or MPLS L3VPN connected using EVPN-VXLAN.
- EVPN-MPLS.
- EVPN in tandem operation with MPLS L3VPN.
- EVPN over MPLS with L3VPN and local stitching.

Directly connected using EVPN-VXLAN

This type of DCI offers the easiest implementation. It involves two devices connected together using a dark fiber or an L2VPN:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

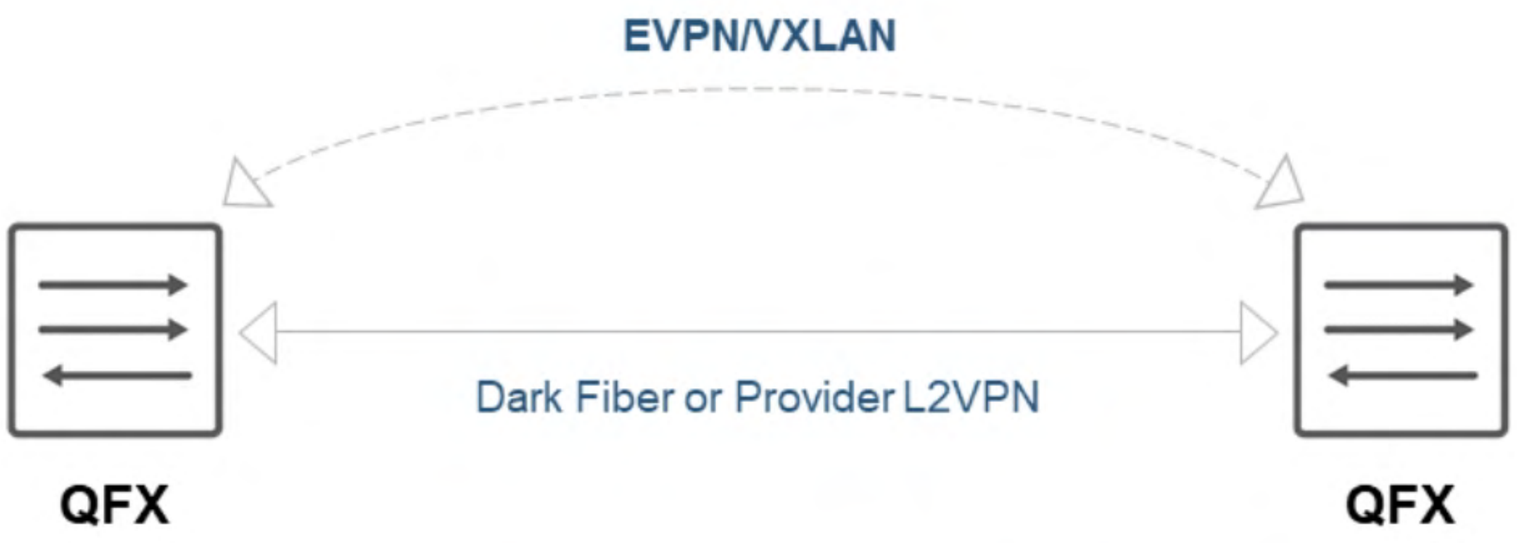

The switches are directly connected to each other through the use of a dark fiber or L2VPN. They run either EBGP or IBGP, together with an IGP, in order to provide for IP connectivity between the loopback interface. This IP-connectivity is then used to enable EVPN signaling and VXLAN-tunneling between the devices.

Internet or MPLS L3VPN connected using EVPN-VXLAN

Instead of connecting the devices that run EVPN-VXLAN, you can also use the Internet or an MPLS L3VPN to provide for IP connectivity between the DCs:

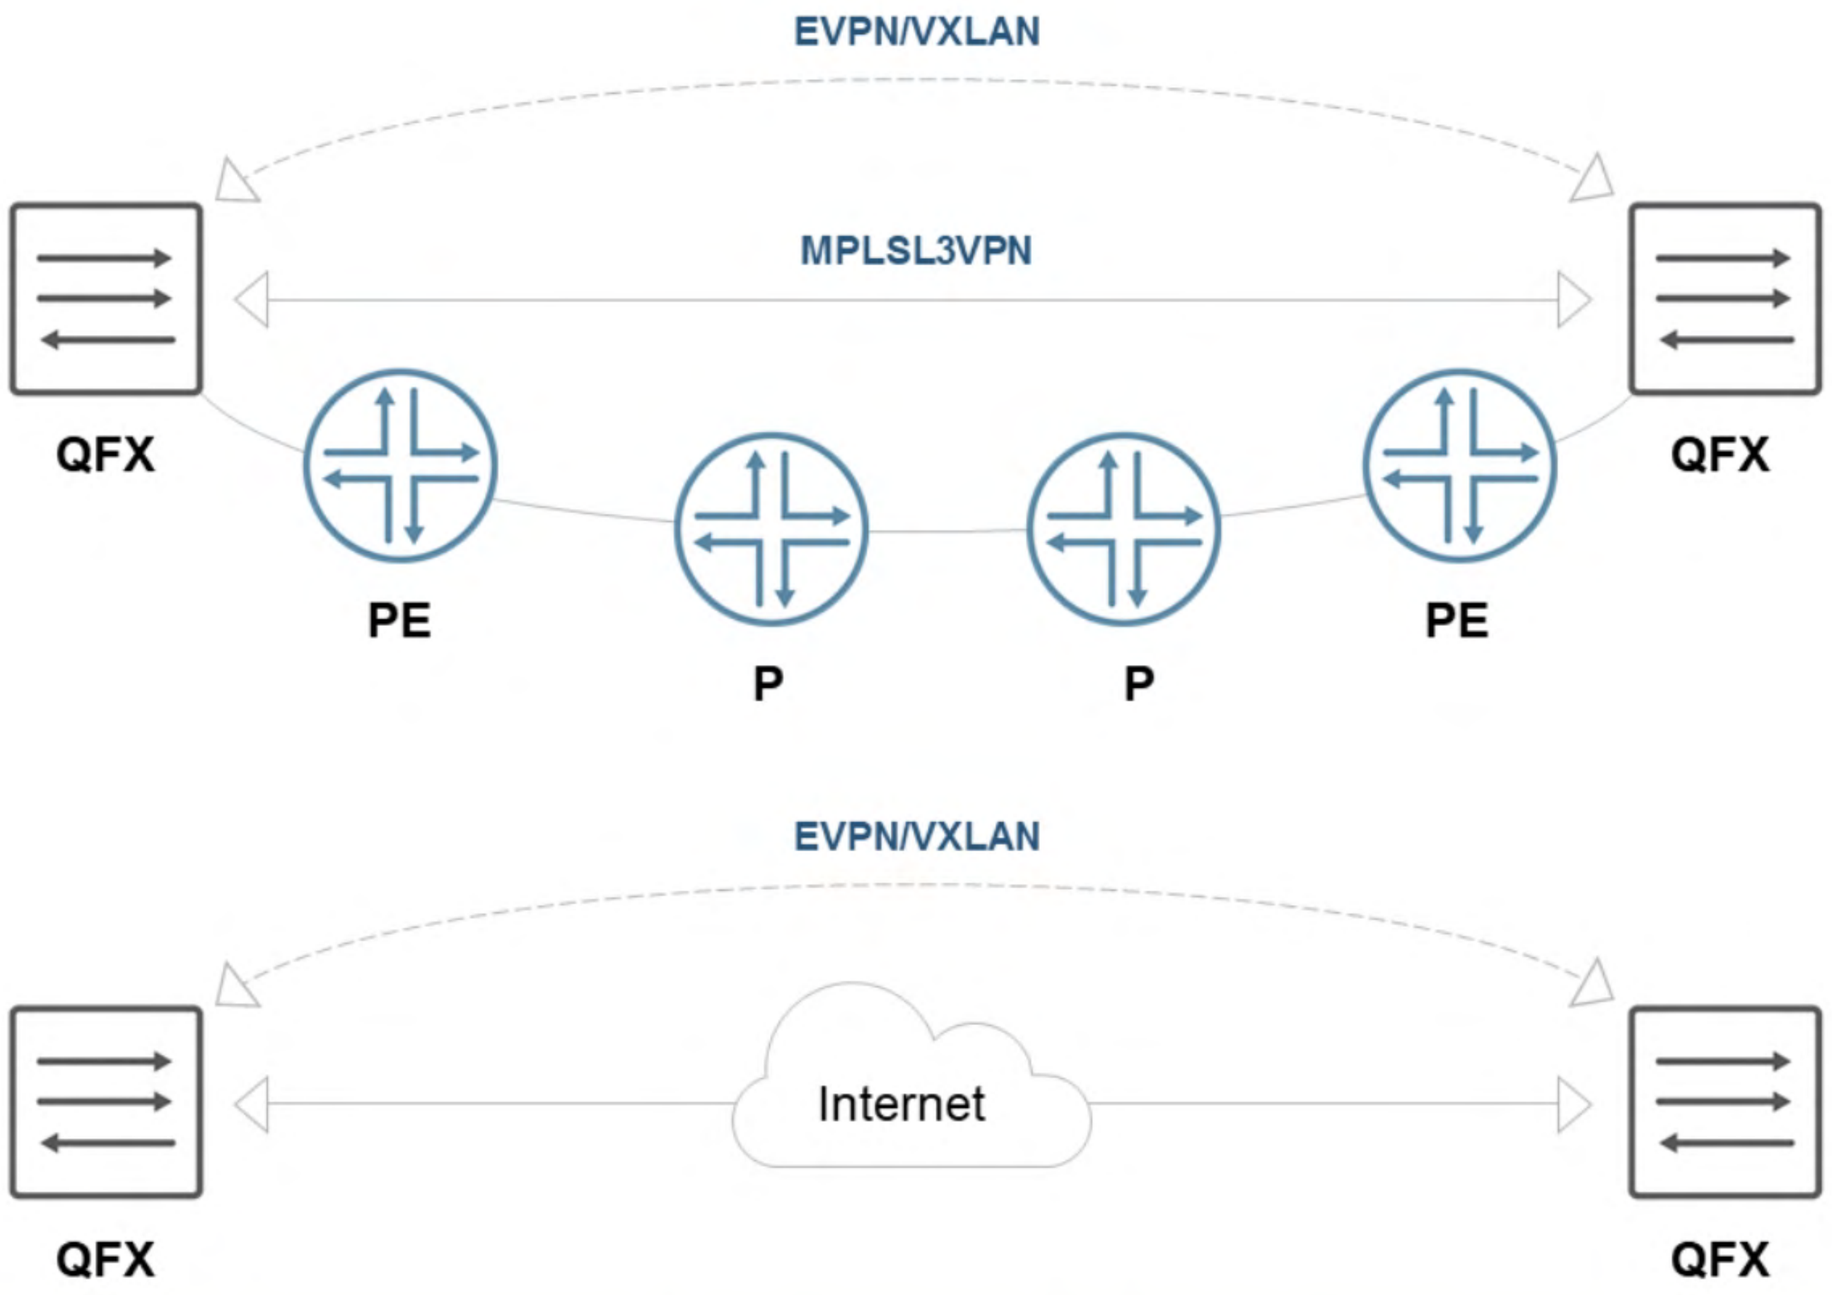

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Using the MPLS L3VPN or the Internet, the devices are provided with IP connectivity. This IP connectivity can then be used for EVPN-VXLAN.

EVPN-MPLS

When MX routers are deployed at the datacenter edges, they can be used to offer this type of DCI:

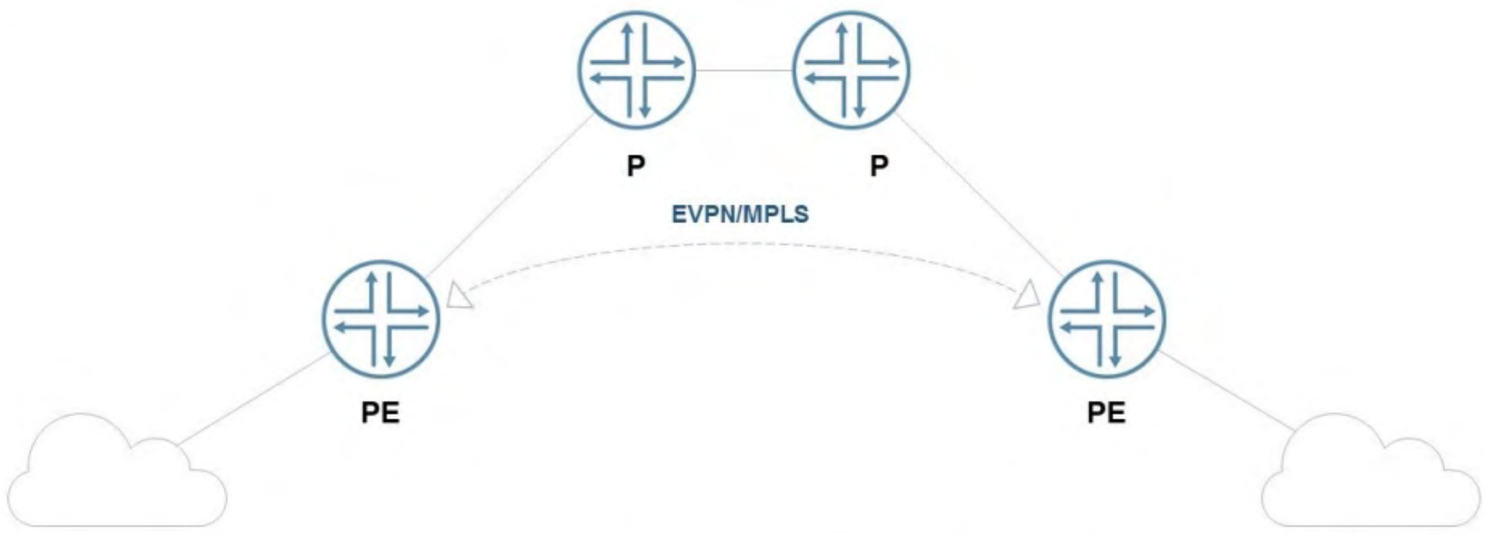

Here, the routers will use EVPN signaling and MPLS as the encapsulation mechanism. Optionally, you can configure the PE with an IRB interface to offer L3 connectivity.

EVPN in tandem operation with MPLS L3VPN

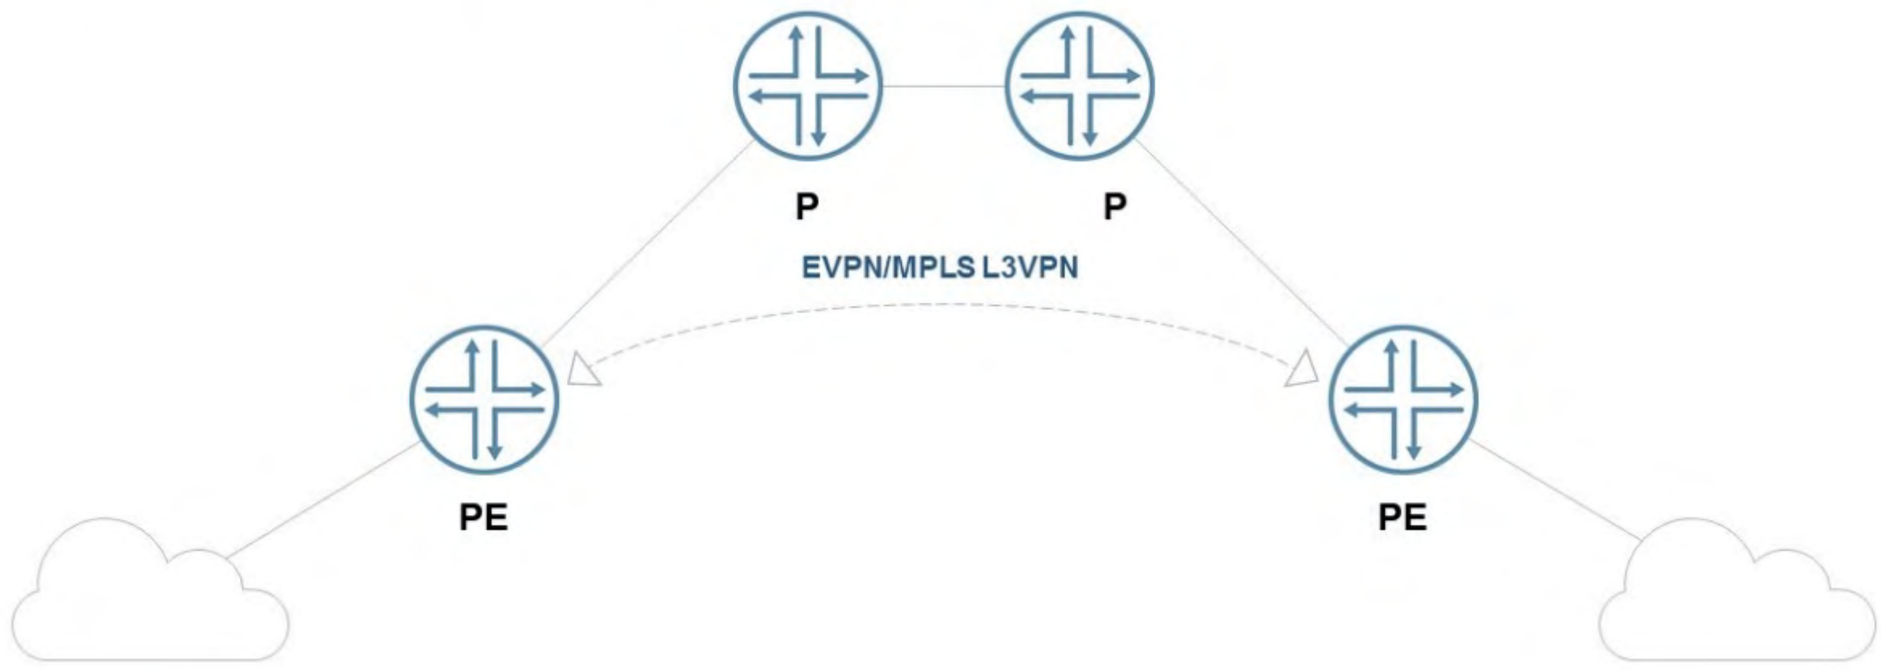

With this option, the routers will use EVPN signaling and MPLS as the encapsulation mechanism. Additionally, the routers are configured with an IRB interface that is placed into an MPLS L3VPN. This way, the DCI offers a full-mesh of L2 and L3 connectivity.

EVPN over MPLS and local stitching

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

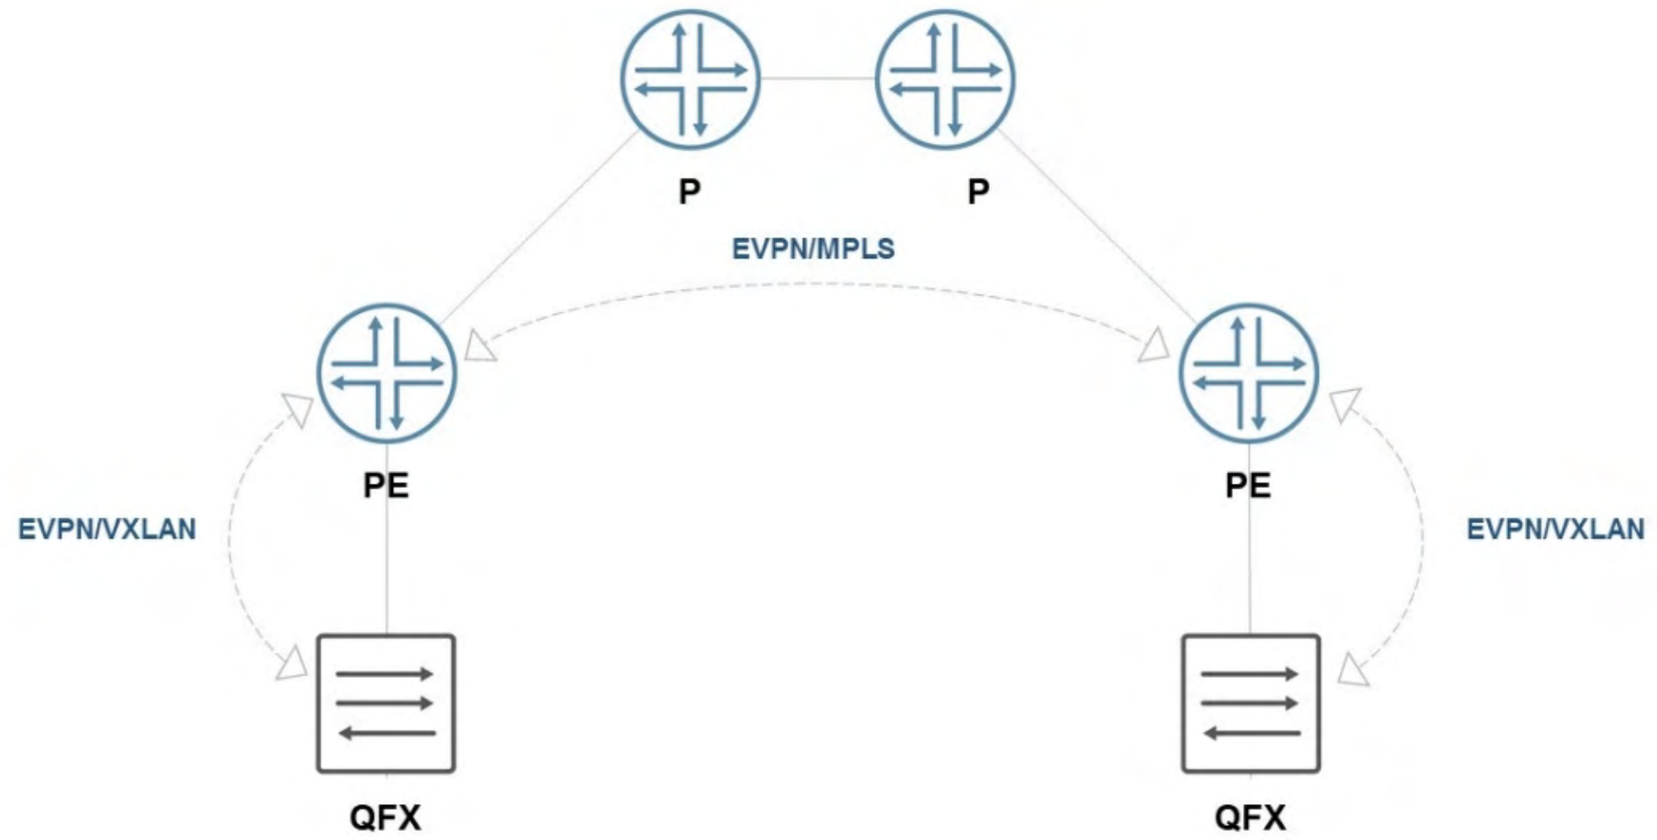

This is the most exotic scenario, mainly targeted at large scale deployments. In this case, an EVPN using MPLS as an encapsulation mechanism is realized between the routers the data centers. Additionally, in each datacenter, the routers and switches are configured to perform EVPN-VXLAN. This EVPN-VXLAN domain is stitched into the EVPN-MPLS domain through the use of a logical tunnel interface on the MX-router. The end-result is a DCI with considerable scaling. Ultimately, the same setup could be used as a DCI that includes a separate controllerless underlay per datacenter:

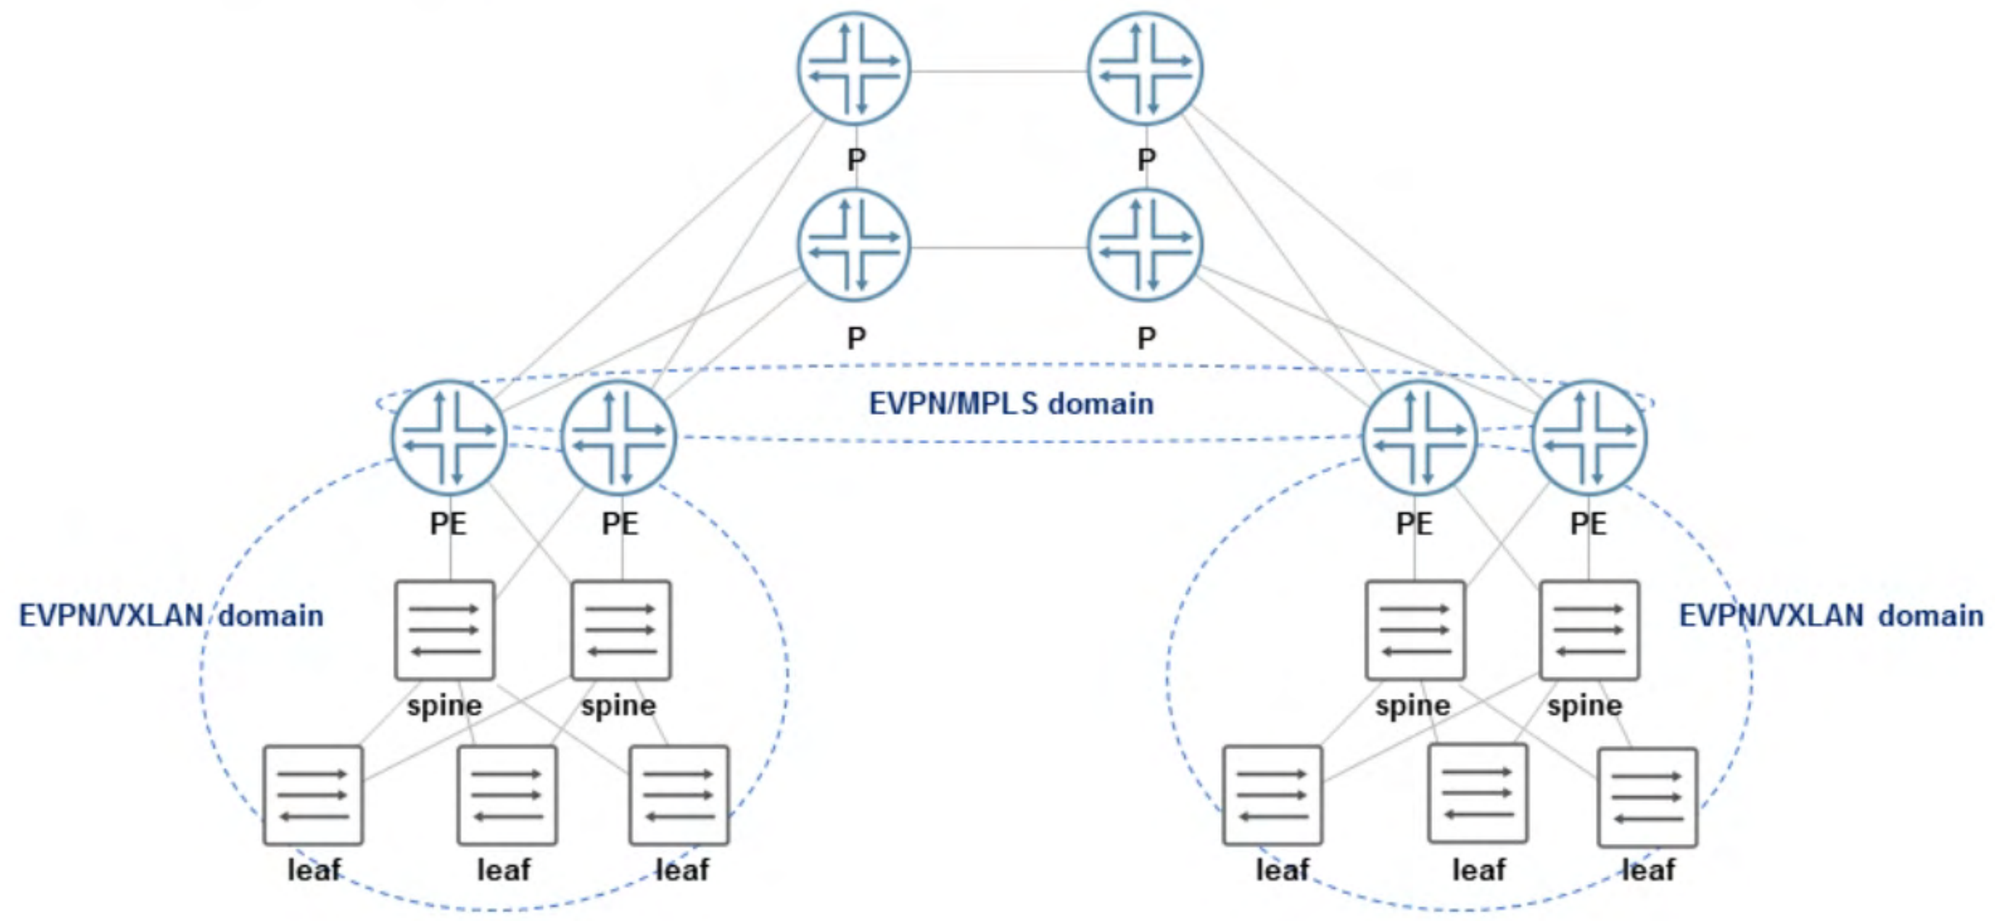

This scenario also offers the option to configure an IRB interface to provide for L3 connectivity. This IRB interface, located on the MX, could be placed into an MPLS L3VPN.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

This way, the DCI connects together two controllerless underlay networks, offering a full-mesh of L2 and L3 connectivity between the data centers.

Part 1: MPLS and MPLS L3VPN

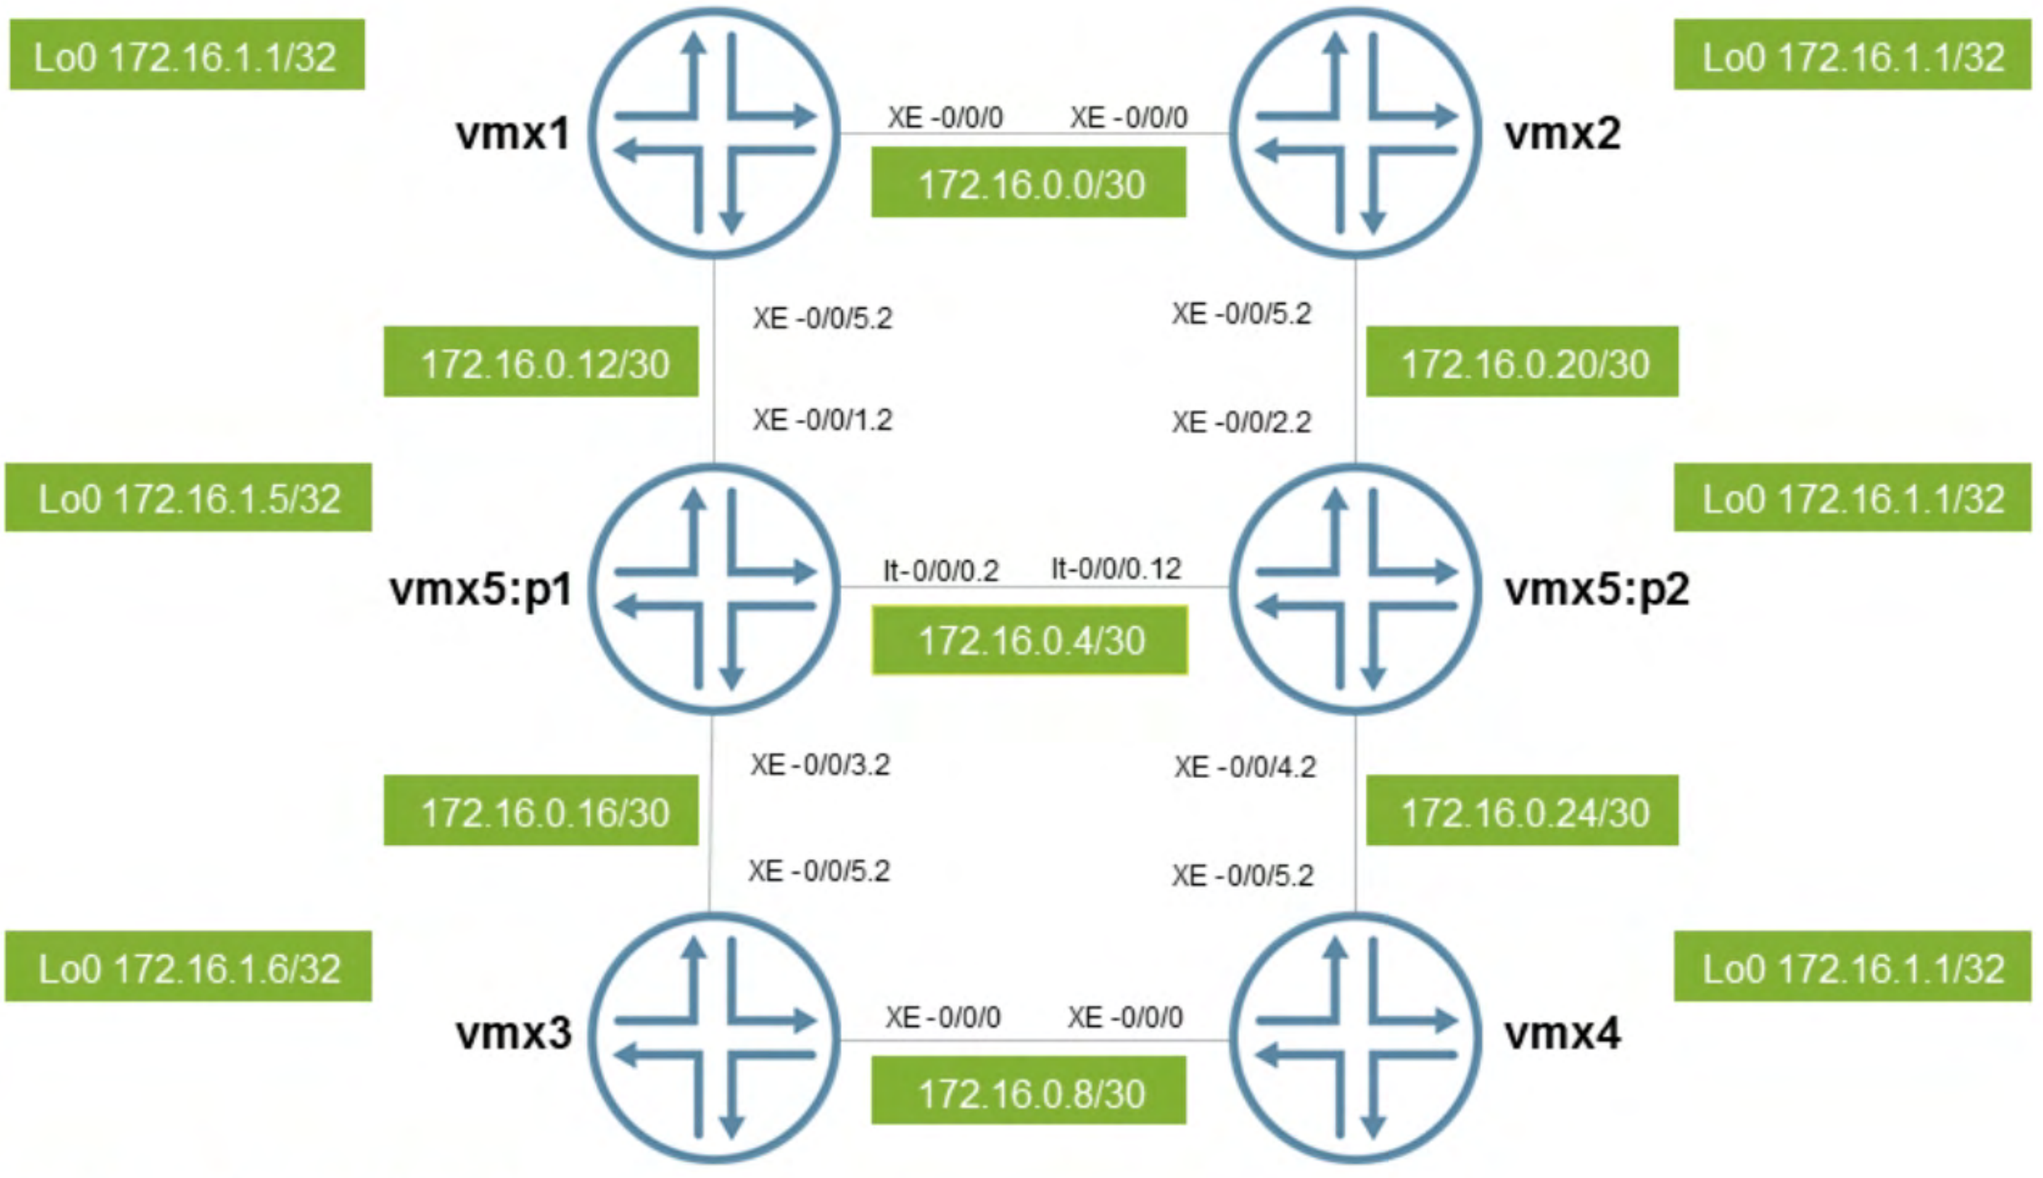

Figure 9 – MPLS network topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

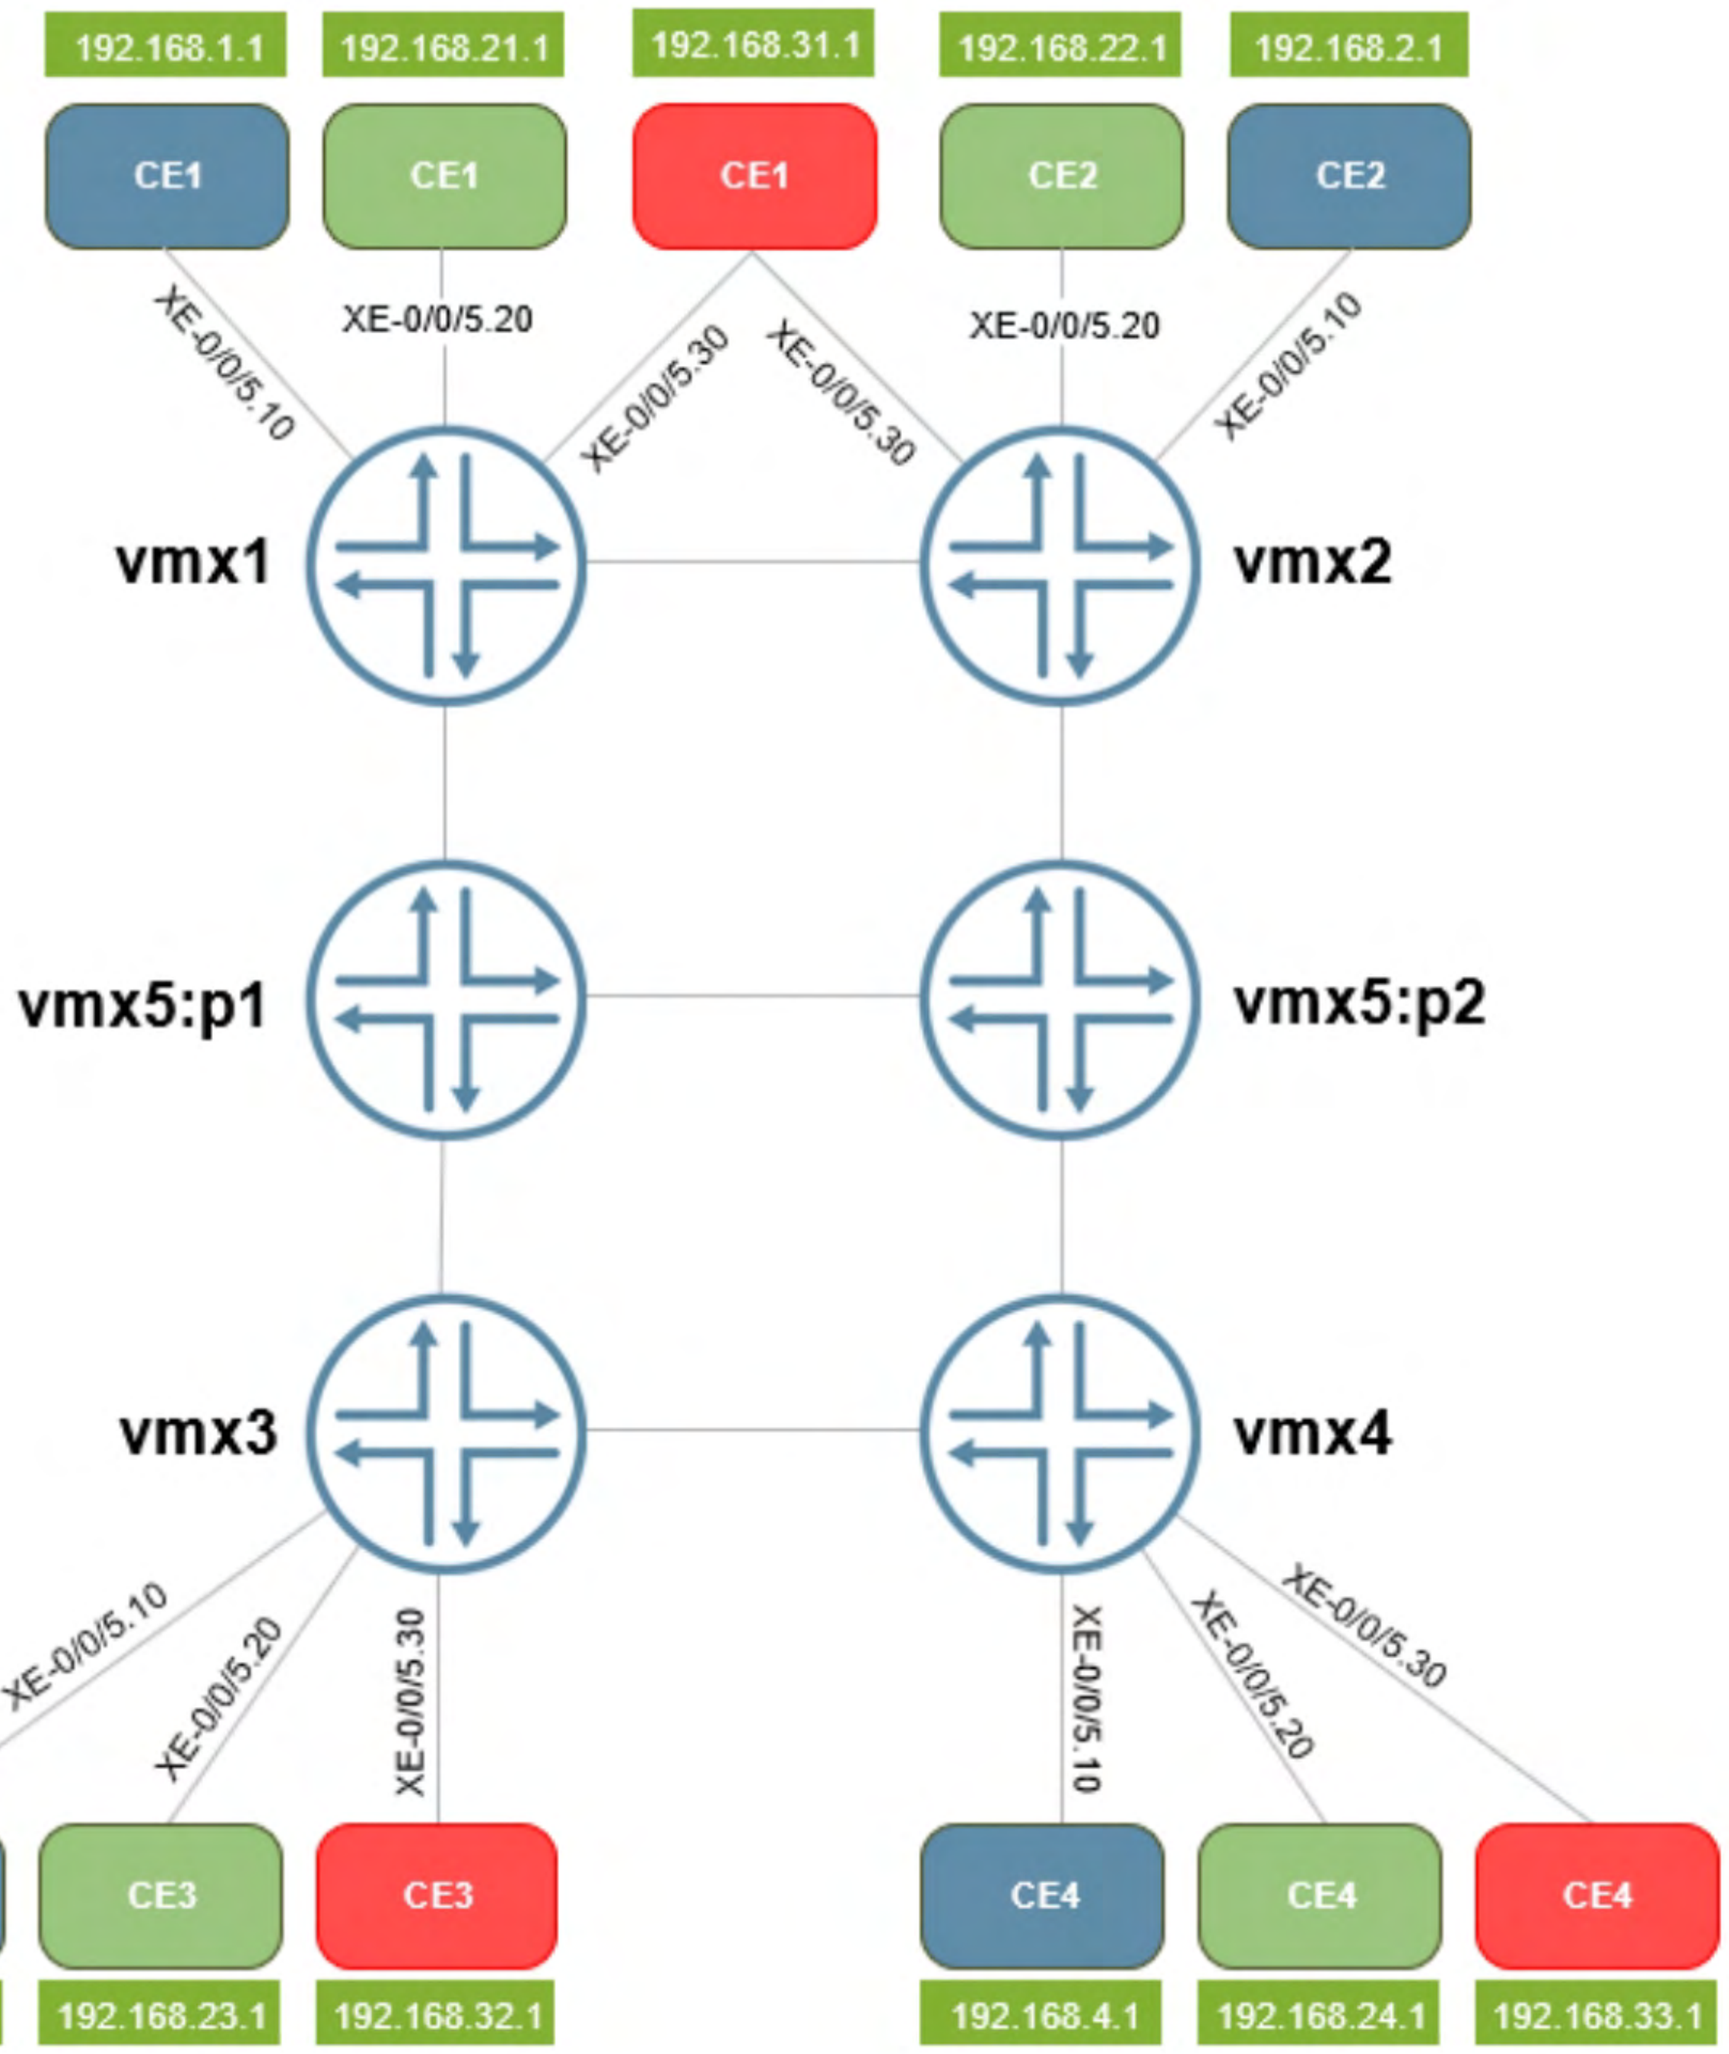

Figure 2 - VPN topology

Notes:

- On vMX5, three logical systems are configured. These are P1, P2 and CE.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.1: IGP Configuration

- Configure all vMX routers to be part of a single area OSPF network.
- Authenticate the OSPF neighbor relationship within the MD5 password 'jncie'
- Make sure OSPF will calculate path cost using 100Gbit as a reference.
- Advertise the loopback IP address into OSPF on every vMX router.
- Make sure the OSPF database only contains router LSAs.

Task 4.2: LDP Configuration

- Enable LDP neighbor relationships between the vMX routers.
- Authenticate LDP sessions using the MD5 hash 'jncie'.
- LDP should use the same metric as the IGP.
- The IGP should prefer links with a working LDP session.

Task 4.3: MPLS LSP Options Configuration

- Enable all protocols to use LDP signaled LSPs for forwarding.
- Tunnel ICMP messages inside MPLS.
- The MPLS network should be represented as a single hop to future VPN customers.

Task 4.4: BGP Configuration

- Enable a BGP full-mesh between vMX1, vMX2, vMX3 and vMX4
- Authenticate all BGP sessions using the password 'jncie'.
- The BGP session should allow for the exchange of MPLS L3 VPN addresses only.
- Use autonomous system 65000.

Task 4.5: MPLS L3VPN JNCIE-BLUE

- PE-CE interfaces configured with VLAN-ID 10 should be placed into a VPN called 'JNCIE-BLUE'.
- Enable the use of the following subnets on CEs located in the MPLS L3VPN:

CE device	**Subnet**
CE1	192.168.1.0/24
CE2	192.168.2.0/24
CE3	192.168.3.0/24
CE4	192.168.4.0/24

- Do not use any dynamic routing protocols between the PE and the CE device. You are allowed to use static routes.
- The PE interface routes should be reachable for remote sites.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- This VPN, as well as future VPNs, should be active on vMX1, vMX2, vMX3 and vMX4.
- Choose any route-target you want.

Task 4.6: MPLS L3VPN JNCIE-GREEN

- PE-CE interfaces configured with VLAN-ID 20 should be placed into a VPN called 'JNCIE-GREEN'.
- Configure OSPF as PE-CE routing protocol.
- Choose any route-target you want.
- The PE interface routes should be reachable for remote sites.

Task 4.7: MPLS L3VPN JNCIE-RED

- PE-CE interfaces configured with VLAN-ID 30 should be placed into a VPN called 'JNCIE-RED'.
- Configure BGP as PE-CE routing protocol. CE devices in this VPN are using AS65500.
- The site connected to vMX1 and vMX2 is multihomed. Make sure traffic to and from that site traverses vMX1.
- Choose any route-target you want.
- The PE interface routes should be reachable for remote sites.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: MPLS, MPLS L3VPN and EVPN.

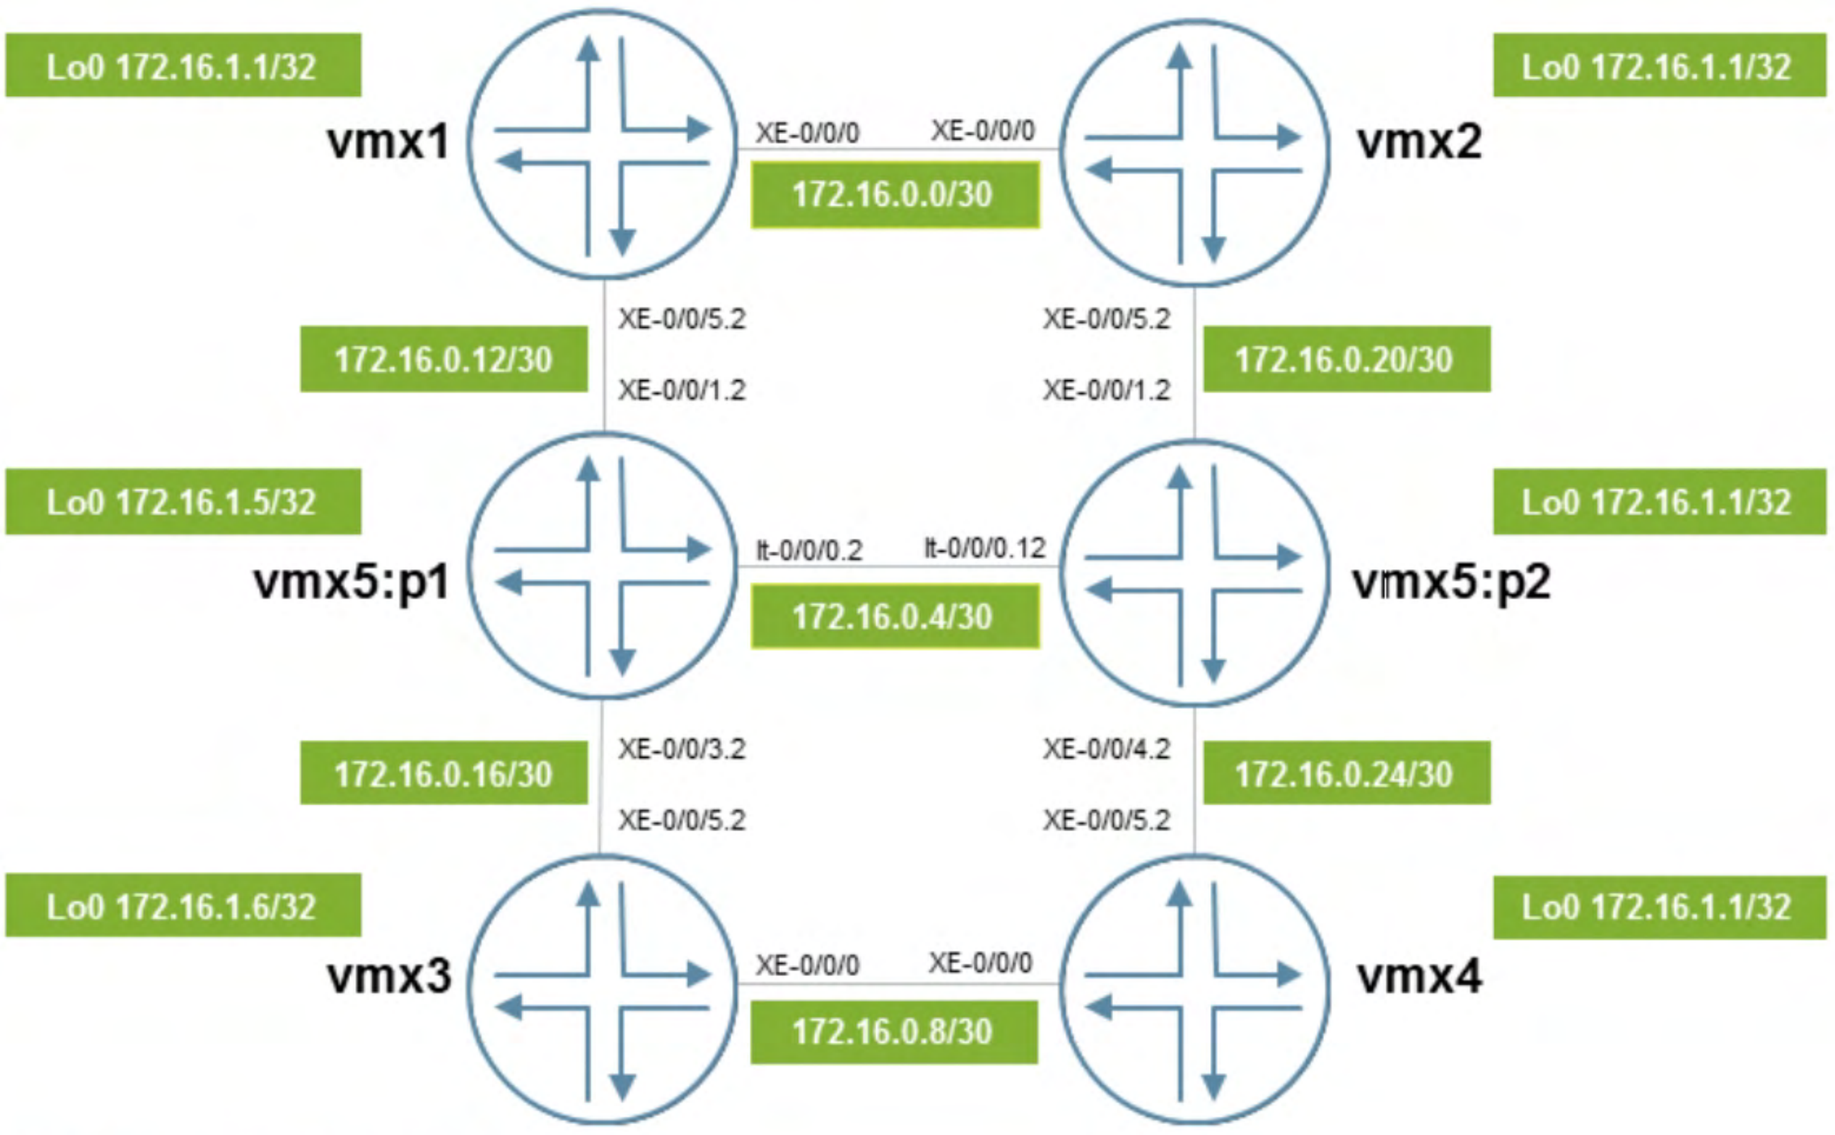

Figure 10 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

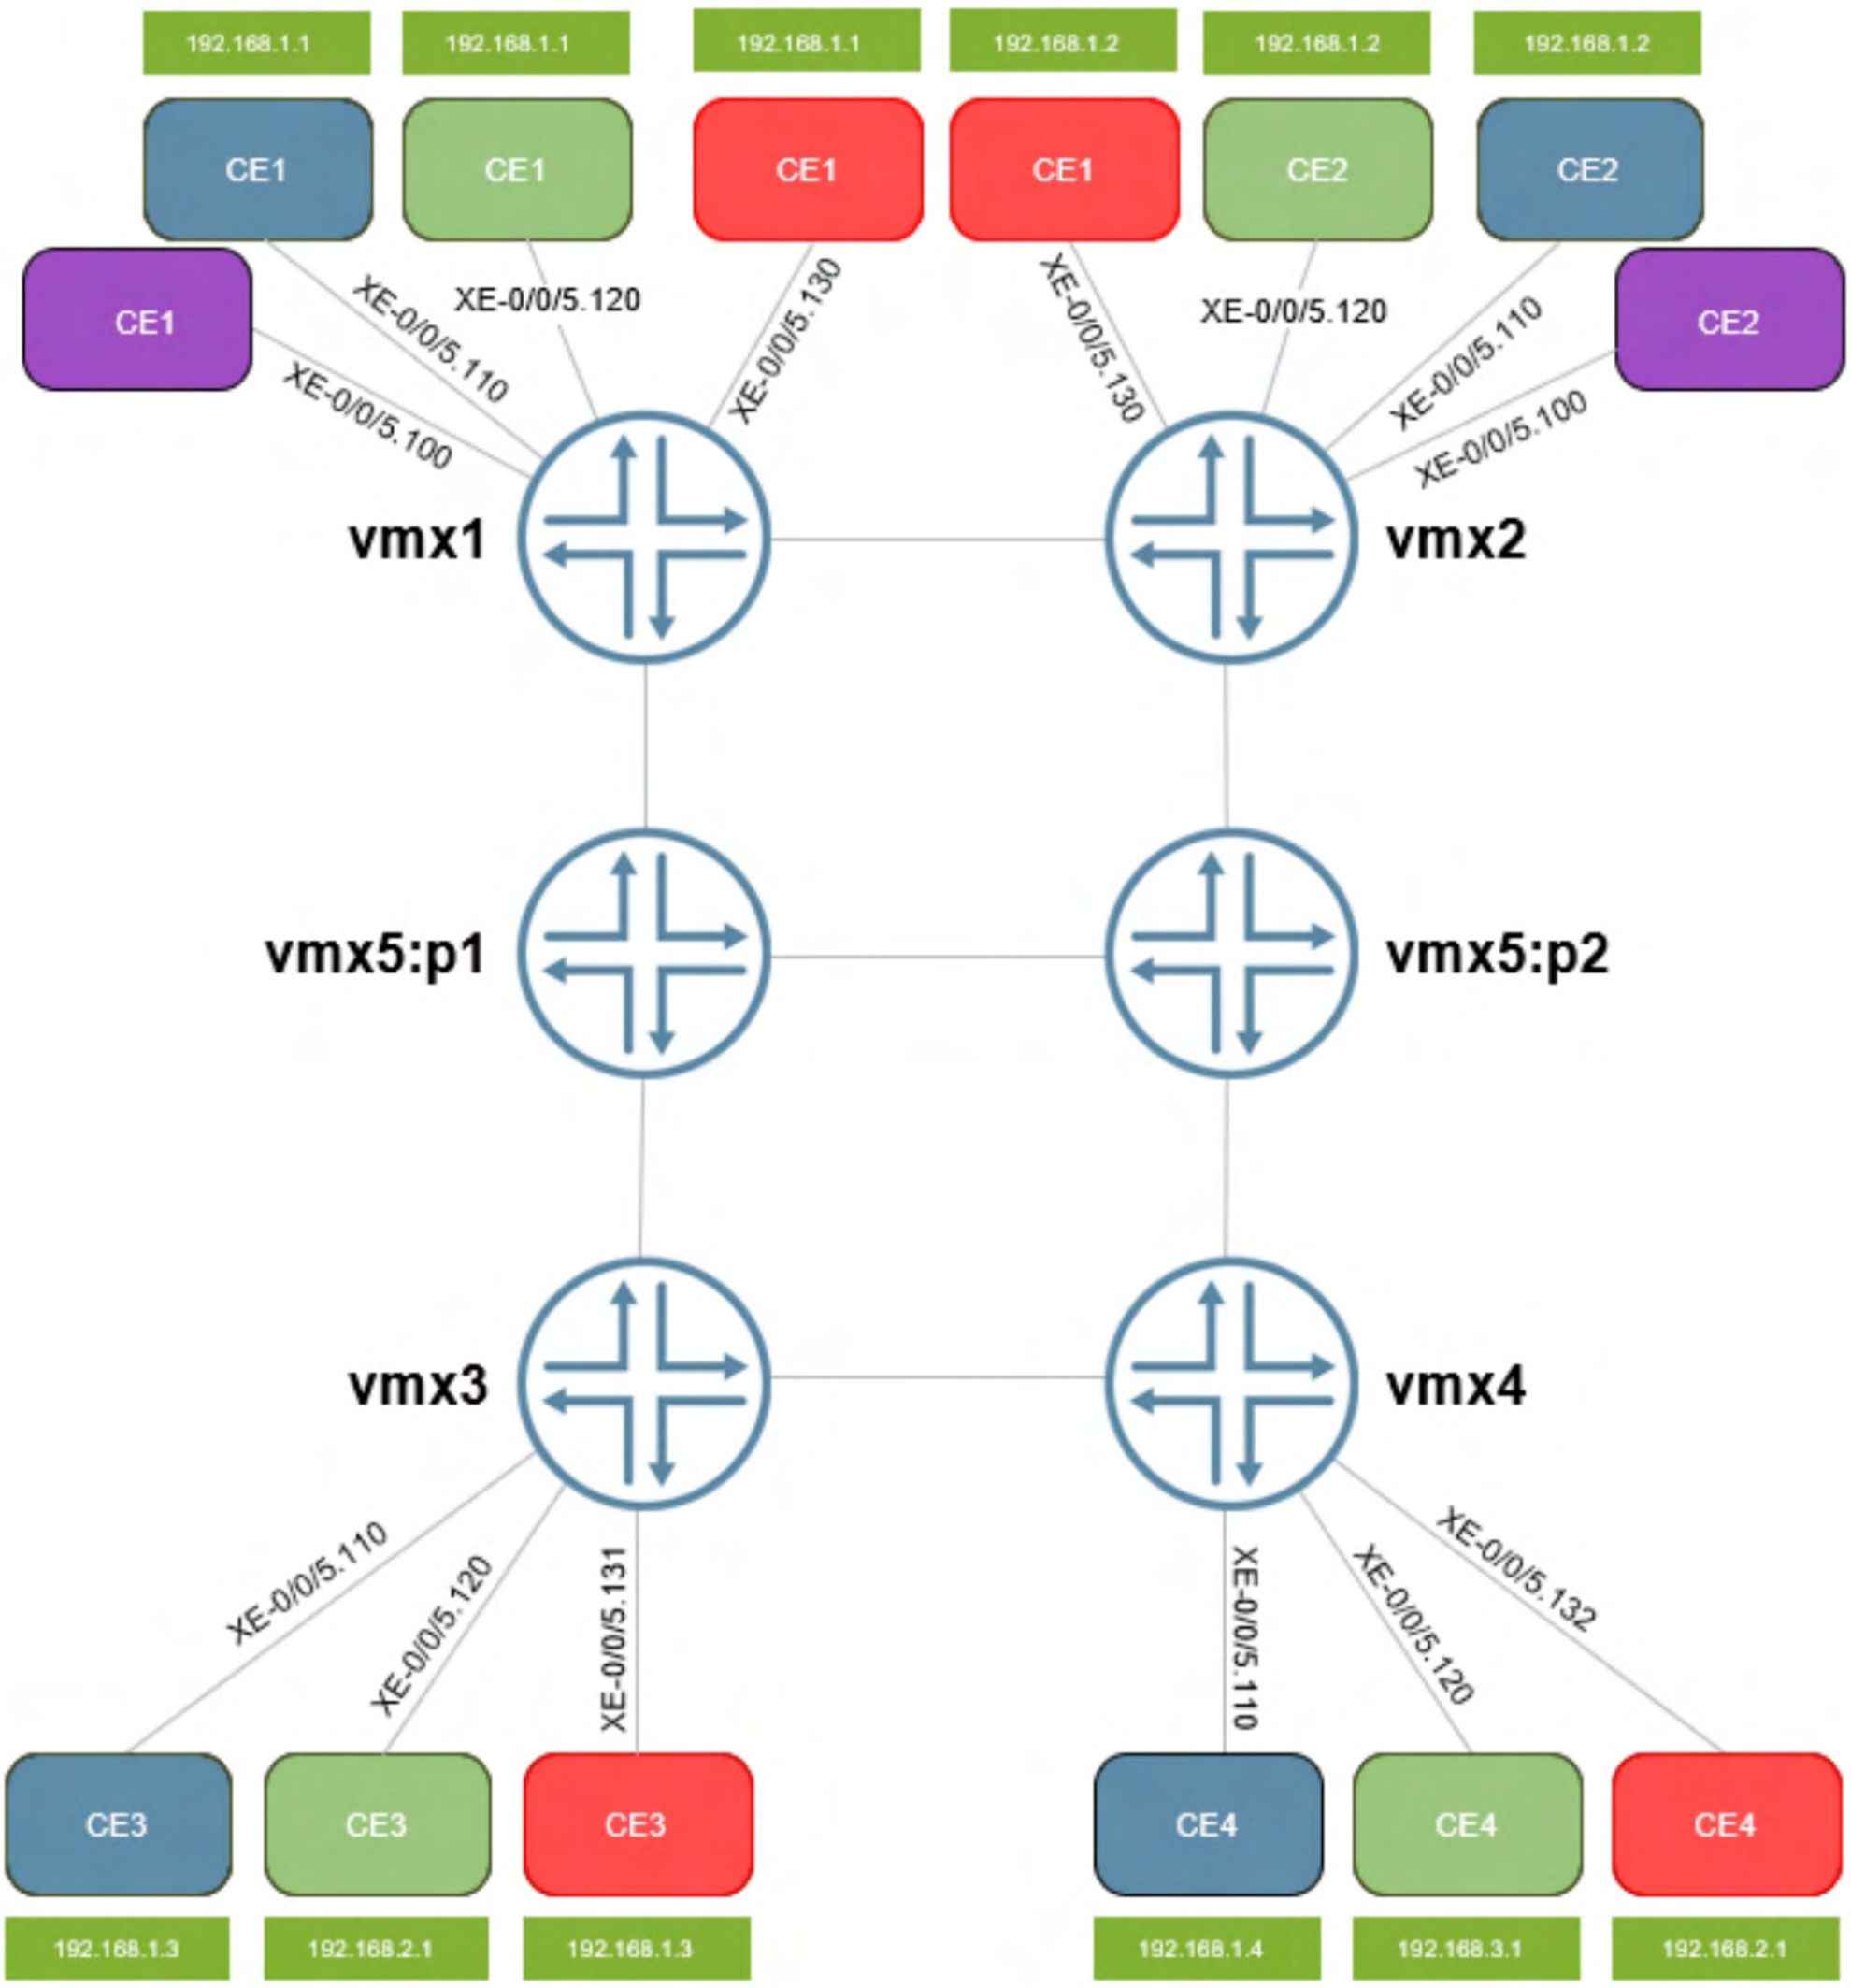

Figure 2 – VPN topology

Notes:

- On vMX5, three logical systems are configured. These are P1, P2 and CE.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.8: RSVP Configuration

- Configure OSPF to support traffic engineering.
- Enable RSVP on all core interfaces.

Task 4.9: RSVP LSP Configuration

- Enable a full-mesh of RSVP signaled LSPs between vMX1, vMX2, vMX3 and vMX4.
- Enable a backup path for all the LSPs that is ready to take over traffic in case of any failures.

Task 4.10: BGP Configuration

- Configure a full mesh of IBGP sessions between vMX1, vMX2, vMX3 and vMX4.
- Use autonomous system 65000.
- Authenticate BGP sessions using the MD5 hash ‘jncie’.
- Enable the sessions to support EVPN as well as L3VPN.

Task 4.11: EVPN JNCIE-PURPLE

- Purple CEs should be placed into an EVPN called ‘JNCIE-PURPLE-EVPN’.
- The PE-CE interfaces must be configured for VLAN 100.
- The EVPN is used for layer 2 services only.

Task 4.12: EVPN JNCIE-BLUE

- Blue CEs should be placed into an EVPN called ‘JNCIE-BLUE-EVPN’.
- The EVPN is used for layer 2 and layer 3 services.
- Configure a layer 3 gateway on every vMX.
- Use IP address 192.168.1.254/24 as the default gateway address.
- Make sure P1 and P2 can exchange traffic with hosts inside the EPVN.

Task 4.13: EVPN JNCIE-GREEN

- The green CEs connected to vMX1 and vMX2 should be placed into an EVPN called ‘JNCIE-GREEN-EVPN’.
- Configure an MPLS L3VPN called ‘JNCIE-GREEN’ for this customer.
- Configure the following gateways on the PEs:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

PE	Gateway
vMX1	192.168.1.254/24
vMX2	192.168.1.254/24
vMX3	192.168.2.254/24
vMX4	192.168.3.254/24

Task 4.14: EVPN JNCIE-RED

- The red CEs connected to vMX1, vMX2 and vMX3 should be placed into an EVPN called 'JNCIE-RED-EVPN'.
- Configure a distributed layer 3 gateway for this EVPN on vMX1, vMX2 and vMX3:

Gateway	Real IP	Virtual IP
vMX1	192.168.1.251/24	192.168.1.254
vMX2	192.168.1.252/24	192.168.1.254
vMX3	192.168.1.253/24	192.168.1.254

- Configure a layer 3 gateway for this customer on vMX4. Use address 192.168.2.254/24 and VLAN-ID 132.
- Configure an MPLS L3VPN called 'JNCIE-RED' for this customer.
- For the sites where distributed gateways are used, do not advertise the EVPN gateway addresses.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 3: MPLS L3VPN and VLAN-aware EVPN.

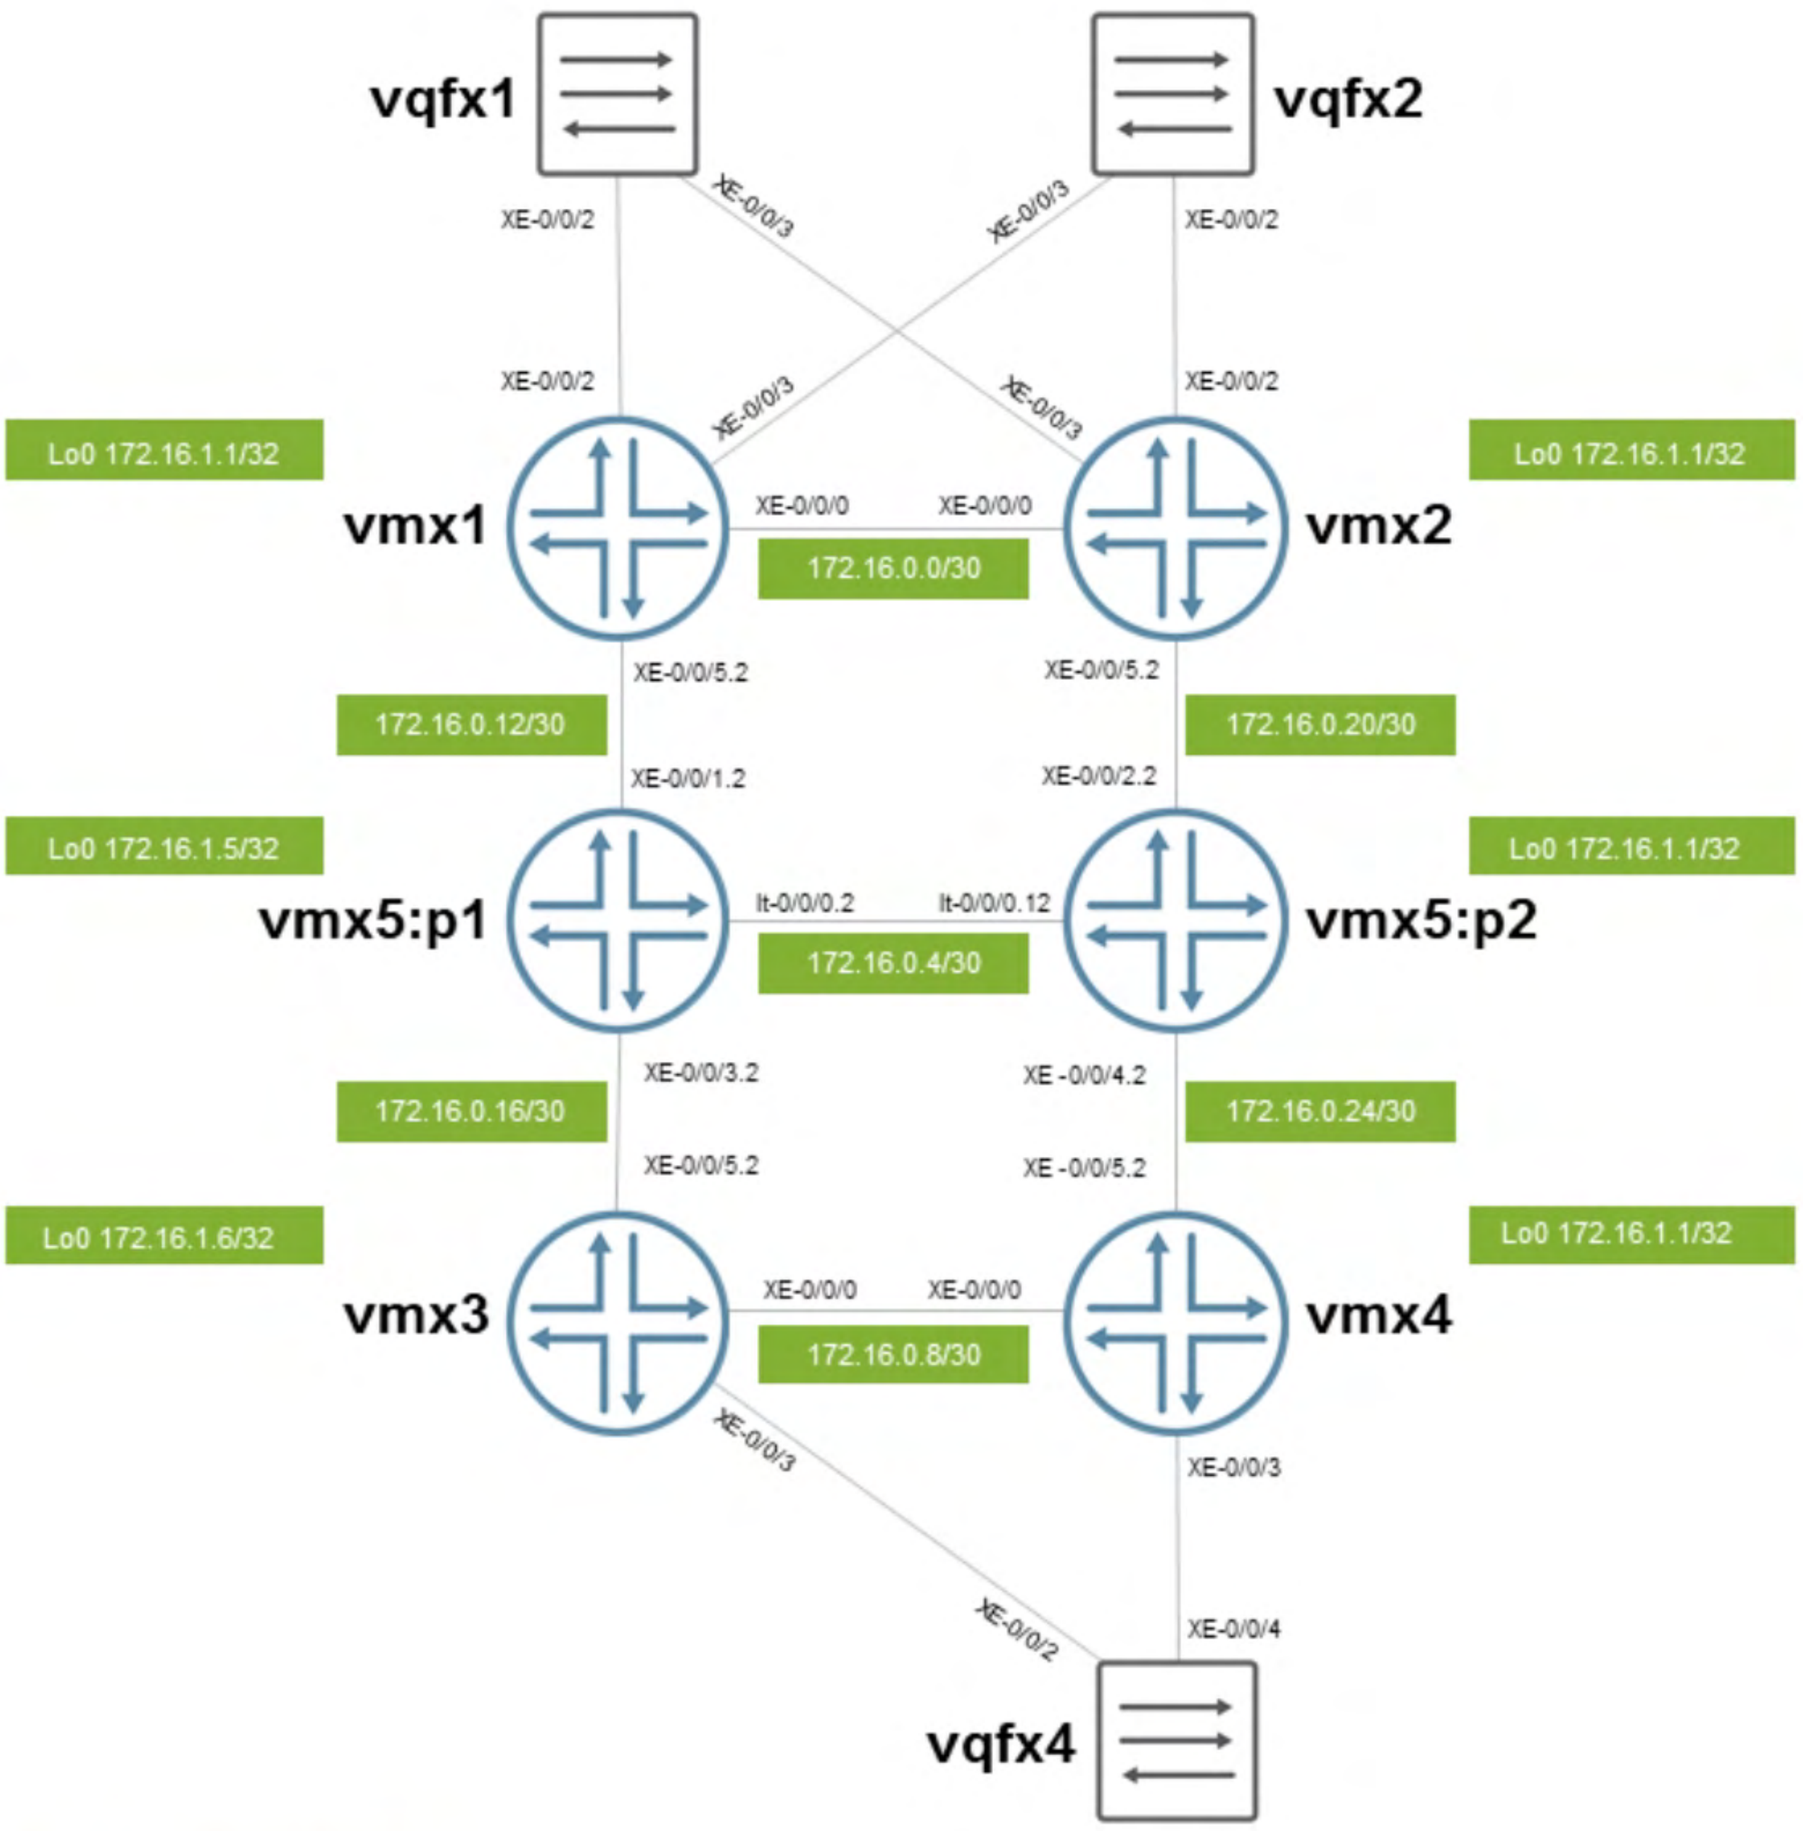

Figure 11 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

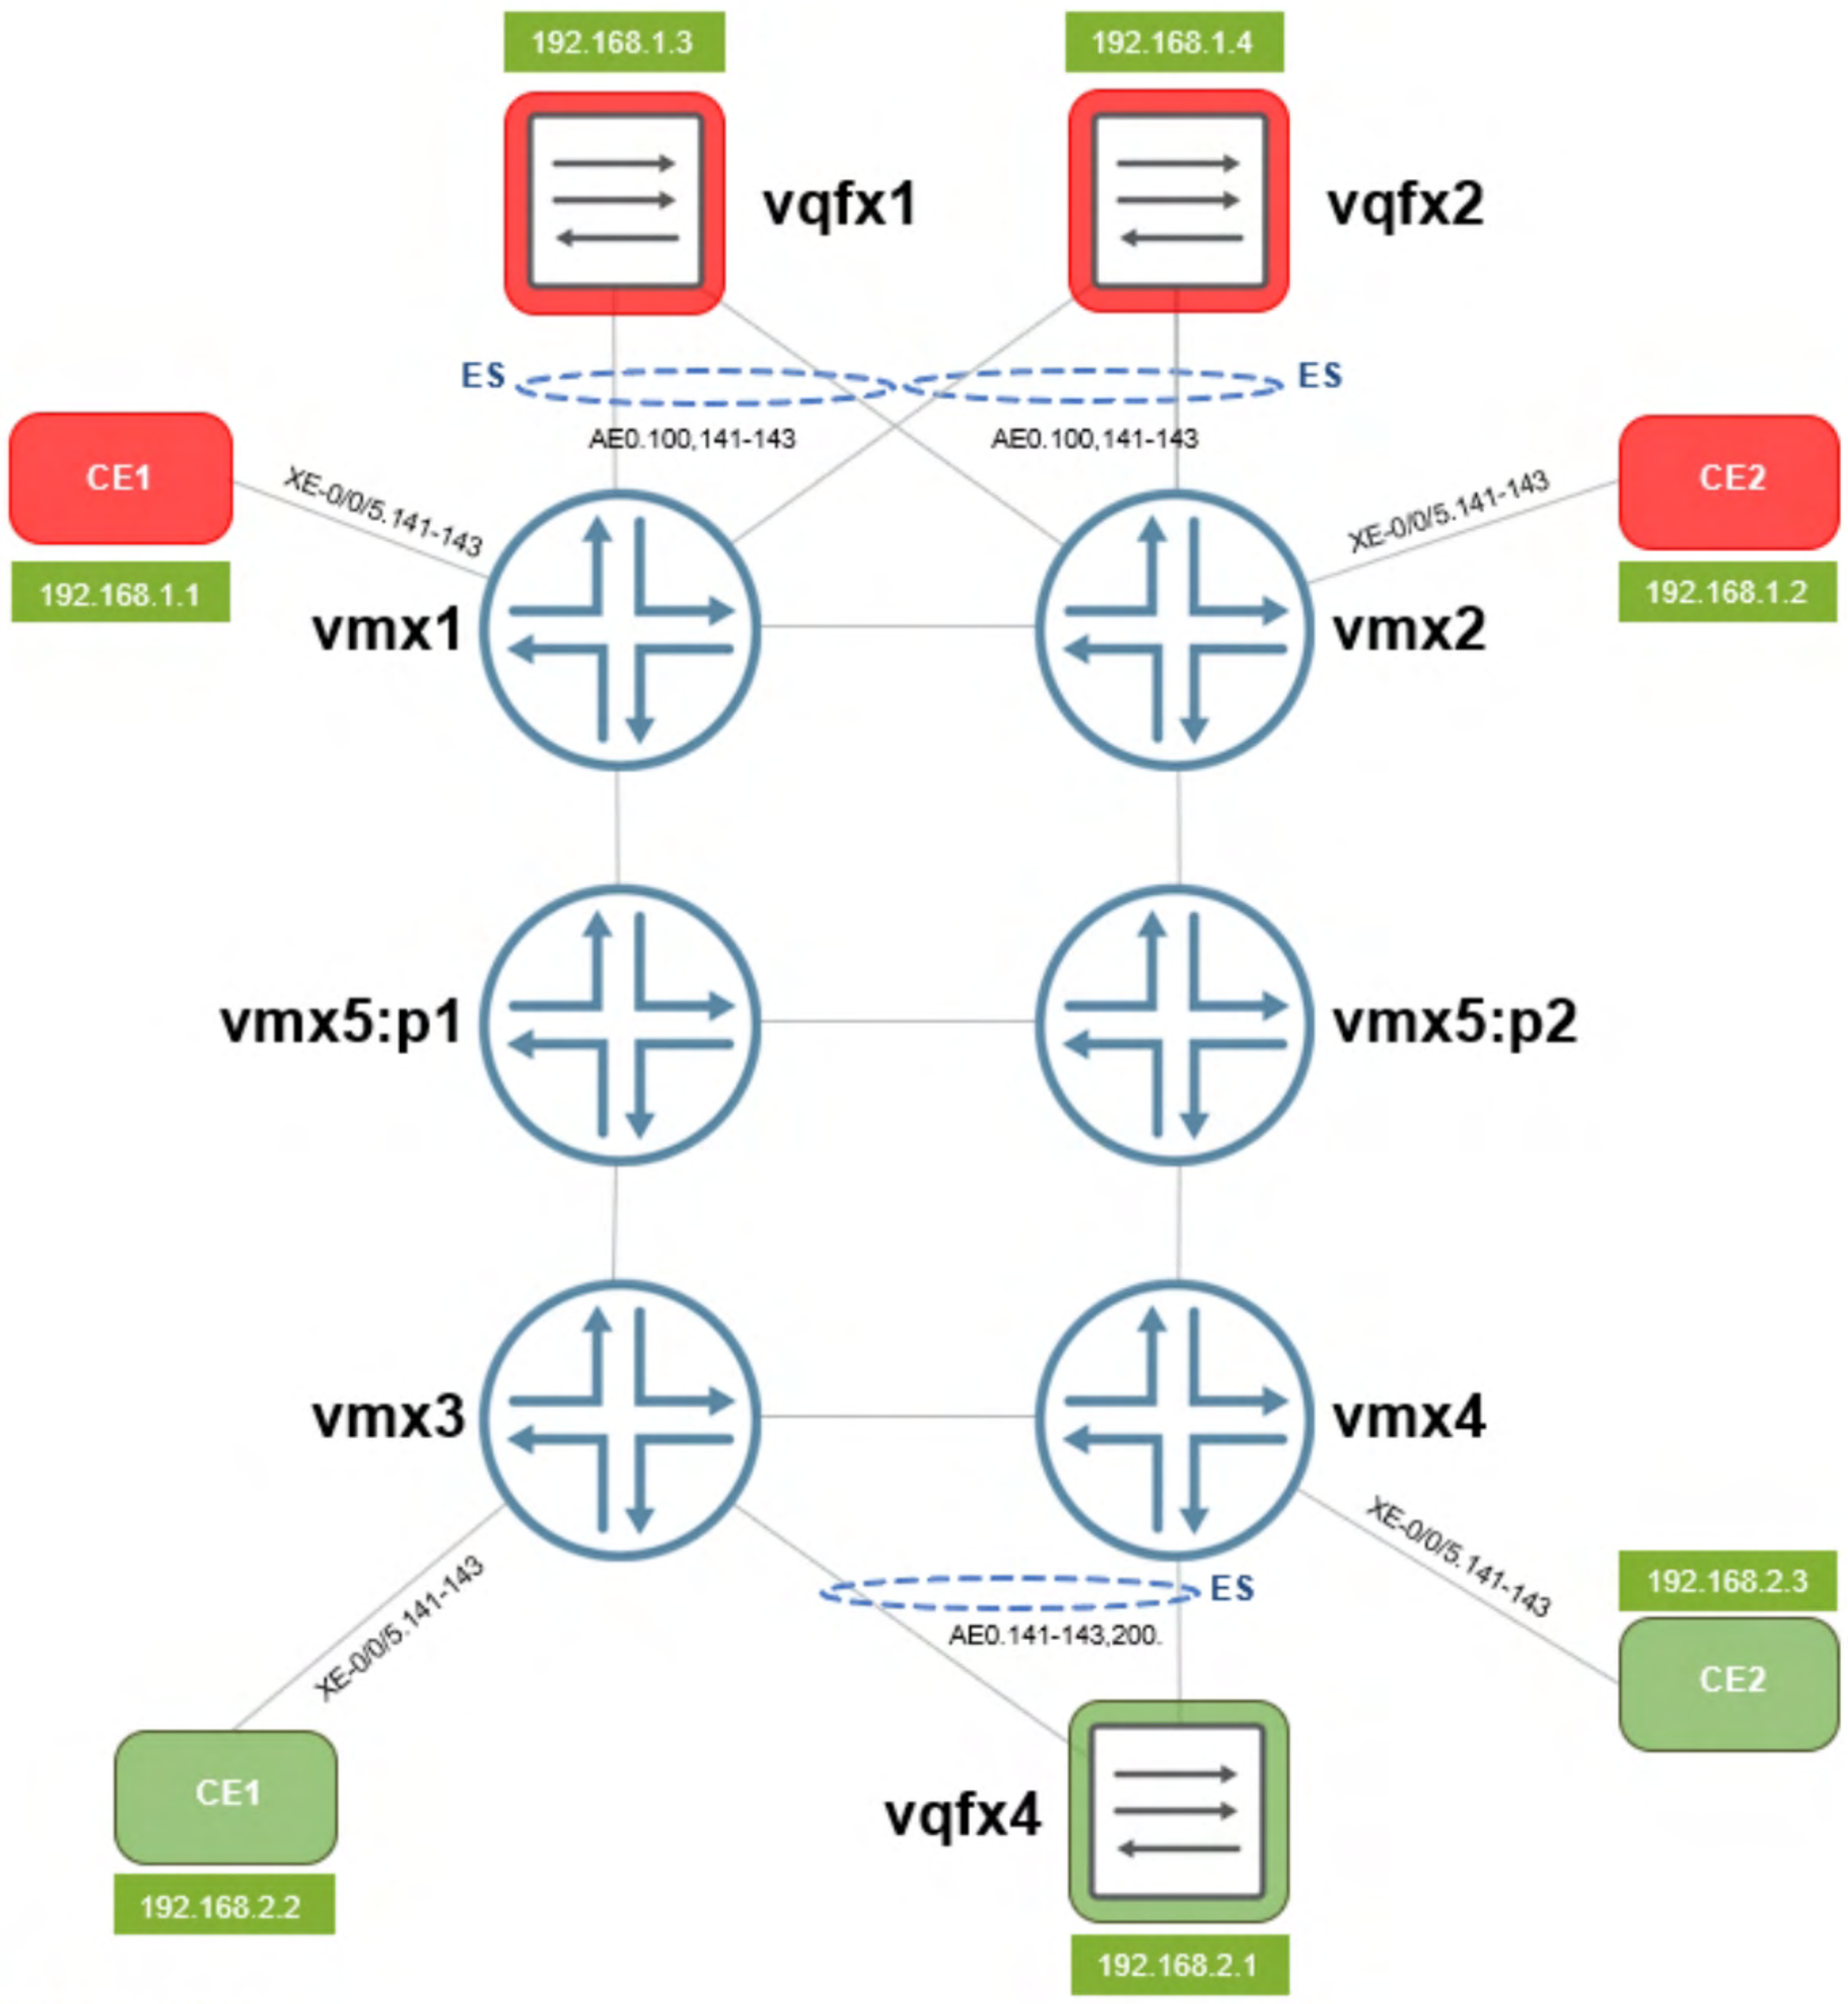

Figure 2 - VPNs

Note: starting topology is configured IP-addressing and LSPs between PE routers.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.15: EVPN North Configuration

- Configure an EVPN called 'jncie-north' on vMX1 and vMX2 for CE1 and CE2.
- Configure a VLAN-aware EVPN using the following parameters:

Gateway	IP address vMX1	IP address vMX2	Gateway address
VLAN 141	192.168.1.251/24	192.168.1.252/24	192.168.1.254
VLAN 142	192.168.1.251/24	192.168.1.252/24	192.168.1.254
VLAN 143			192.168.1.254

- Ensure IP connectivity between hosts inside VLAN 141 and all of the vMX routers.
- Separate layer 3 traffic for VLAN 142 and 143 through the use of MPLS L3VPN.

Task 4.16: EVPN South Configuration

- Configure an EVPN called jncie-south on vMX3 and vMX4 for CE1 and CE2.
- Configure the VLAN-aware EVPN using the following parameters:

Gateway	IP address vMX3	IP address vMX4	Gateway address
VLAN 141	192.168.2.251/24	192.168.2.252/24	192.168.2.254
VLAN 142	192.168.2.251/24	192.168.2.252/24	192.168.2.254
VLAN 143			192.168.2.254

- Ensure IP connectivity between hosts inside VLAN 141 and all of the vMX routers.
- Separate layer 3 traffic for VLAN 142 and 143 through the use of MPLS L3VPN.
- Ensure that hosts inside VLAN 142 have IP-connectivity with hosts inside VLAN 142 that was previously configured for the 'jncie-north' EVPN.
- Ensure that hosts inside VLAN 143 have IP-connectivity with hosts inside VLAN 143 that was previously configured for the 'jncie-north' EVPN.

Task 4.17: EVPN Multihoming

- Configure the following Ethernet segments:

	ES 1	ES 2	ES 3
vMX1	xe-0/0/2	xe-0/0/3	
vMX2	xe-0/0/3	xe-0/0/2	
vMX3			xe-0/0/3
vMX4			xe-0/0/3

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- On vMX1 and vMX2, place the interfaces into the 'jncie-north' EVPN. On vMX3 and vMX4, place the interfaces into the 'jncie-south' EVPN.
- Enable the newly created interfaces for any VLAN that is configured for the EVPN.

Task 4.18: VLAN-based EVPN Configuration

- Configure a VLAN-based EVPN called 'ie-north' on vMX1 and vMX2.
- Configure the following IP addresses on vMX1 and vMX2:

PE device	IP address	Virtual Gateway address	VLAN
vMX1	192.168.100.251/24	192.168.100.254/24	100
vMX2	192.168.100.252/24	192.168.100.254/24	100

- Configure a VLAN-based EVPN called 'ie-south' on vMX3 and vMX4.
- Configure the following IP addresses on vMX3 and vMX4:

PE device	IP address	Virtual Gateway address	VLAN
vMX3	192.168.200.251/24	192.168.200.254/24	200
vMX4	192.168.200.252/24	192.168.200.254/24	200

- Configure an MPLS L3VPN named ‘jncie-vpn’ for all newly configured layer 3 interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 4: Controllerless overlay and EVPN DCI

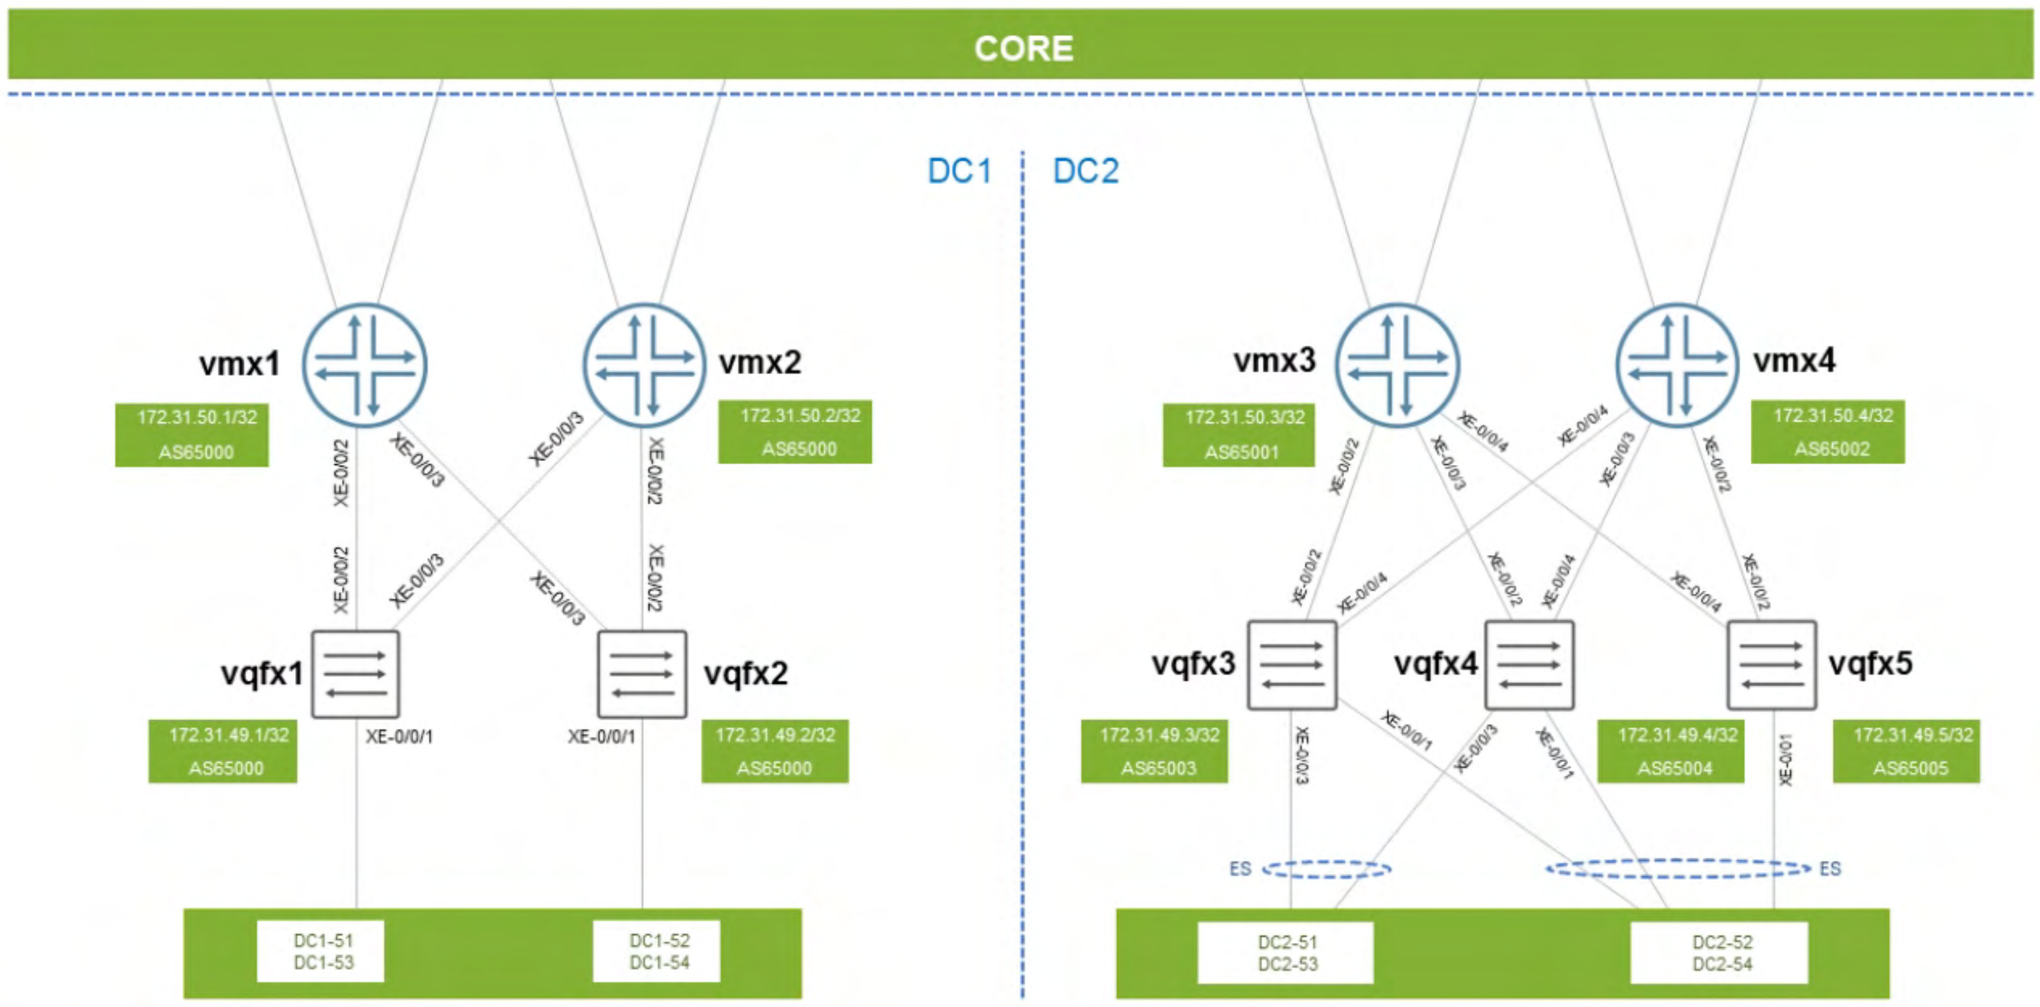

Figure1 –Network topology

Task 4.19: DC1 Controllerless Overlay Part 1: Routing and Signaling

- Ensure IP-connectivity between the loopback IPs of the devices in DC1 by configuring OSPF on the links between the vQFX and the vMX in DC1.
- Source BGP sessions from the loopback interface and enable BGP for EVPN signaling.
- The vQFX devices in DC1 are only allowed to peer with the vMX routers and should receive each other's EVPN routes.
- Use the AS numbers shown in the topology.
- Enable the devices to load-share traffic across links.

Task 4.20: DC1 Controllerless Overlay Part 2: the Overlay

- Configure a VXLAN layer 2 gateway on the vQFX switches for VLAN 200.
- Server-facing interfaces should support tagged traffic. Ensure that the servers can reach the core network.
- The devices should automatically assign RTs for the VXLAN segments.
- Configure the following layer 3 gateway on vMX1 and vMX2:

VLAN	IPv4 gateway	MAC address
200	192.168.200.254	00:00:00:00:02:00

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.21: DC2 Controllerless Overlay Part 1: The Underlay

- Ensure IP-connectivity between the loopback IPs. Use EBGP to accomplish your goal.
- Enable EVPN signaling between the devices. Source the BGP sessions used for EVPN signaling from the loopback interface.
- Use the AS numbers shown in the topology.
- Enable the devices to load-share traffic across links.

Task 4.22: DC2 Controllerless Overlay part 2: The Overlay

- Configure a VXLAN layer 2 gateway on the vQFX switches for VLAN 100.
- Server-facing interfaces should support tagged traffic. Ensure that the servers can reach the core network.
- Configure the following layer 3 gateway on vMX3 and vMX4:

VLAN	IPv4 gateway	MAC address
100	192.168.100.254	00:00:00:00:01:00

- You should be able to hide VXLAN segments on a per-device basis.
- Configure the Ether Segments shown in the topology.

Task 4.23: EVPN Stitching

- Configure a layer 2 VXLAN gateway on all server facing interfaces for VLAN 250. This VXLAN gateway should be active on every vQFX.
- Configure a layer 2 gateway for VLAN 250 on both vMX2 as well as vMX3.
- Any inter-datacenter traffic inside VLAN 250 should be transported across the MPLS LSPs.

Chapter 5: Security

When you are thinking about securing the network, you have three areas of focus and two types of traffic to consider. The areas of focus are the control plane, the data plane and transit traffic. The different types of traffic you need to take into consideration are control plane traffic and transit traffic.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

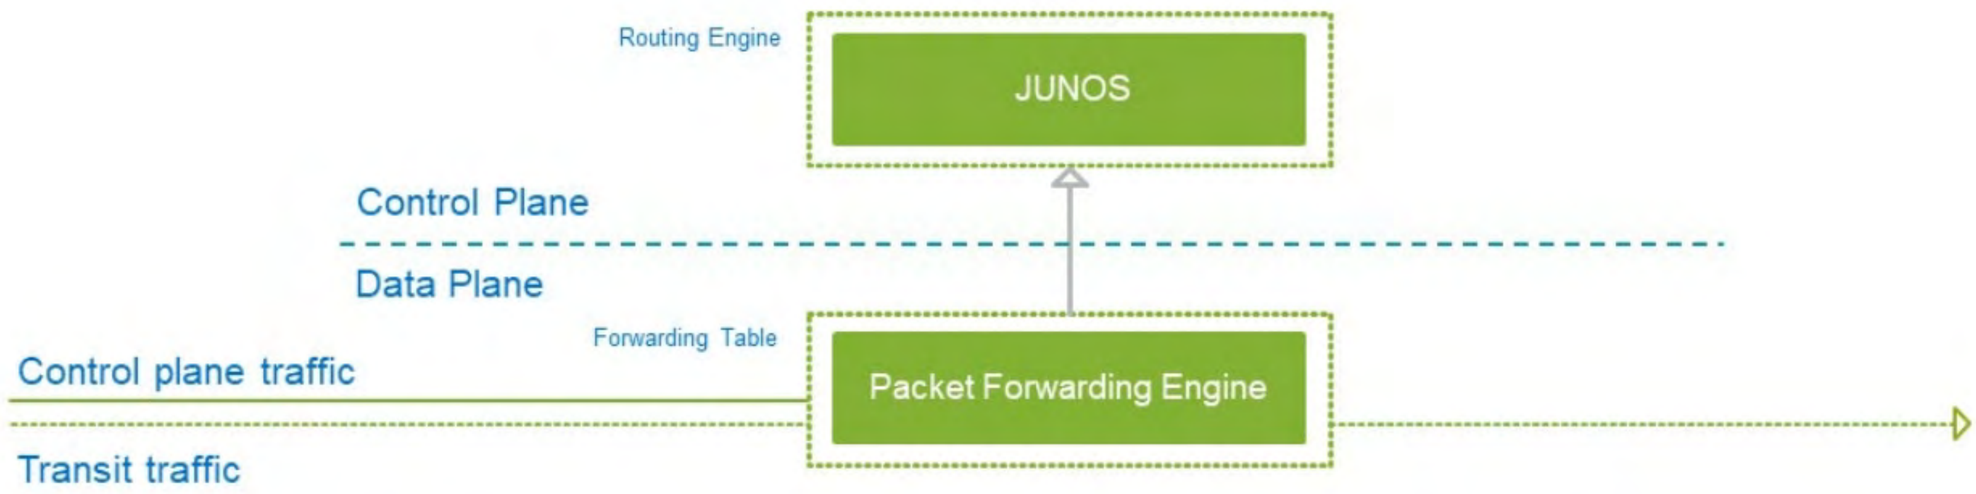

Within Junos, there is a separation between the control plane and the data plane. Inside the control plane, we find the routing-engine (RE), which takes care of system management and the routing process. The RE has an internal link to the data plane where we find the Packet Forwarding Engine (PFE). The RE continuously programs the ASICs in the data plane to make sure traffic is forwarded according to the most up-to-date routing information.

Traffic that is send destined to a device is called control plane traffic. The PFE forwards this traffic to the RE for evaluation. Control plane traffic examples are routing-updates, ARP requests and device management. The other type of traffic is called transit traffic. This traffic is handled completely by the PFE and it is destined to another part of the network.

Control plane protection

The RE inside the control plane needs to be protected at all cost. The RE needs to be protected from unwanted access, protocols misbehaving and denial-of-service. RE protection from denial-of-service should focus on attacks as well as abnormal conditions in the network. It is best practice to permit only the protocols that run on the RE from IP subnets that the RE needs to communicate with and rate-limit this traffic to an amount that suits your networks need. In Junos, tight control over this control plane traffic is achieved through stateless firewall filters that are applied to the lo0 interface.

Firewall filters are a collection of rules that can be used to match traffic and perform an action. The filter is evaluated sequentially, from top to bottom. As soon as a match with a terminating action is encountered, the filter stops the evaluation and applies the action.

Firewall filter actions include the following:

- **accept**: accept the packet
- **discard**: discard the packet without sending an ICMP
- **reject**: reject the packet and send and ICMP
- **policer**: subject the traffic that is matched to a rate-limiter
- **log**: log packet header to file

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- **syslog**: syslog information about the packet
- **count**: increment a named counter whenever there is a match

When you are configuring firewall filters, remember that it is possible to reference the name of a protocol instead of the port number. When you need to configure a lengthy firewall filter, this can be quite a substantial time-saver. Another consideration is to use named terms to improve the readability of the firewall filter during troubleshooting. When firewall filters become very lengthy, consider an input-list of different firewall filters to further increase the readability and manageability.

When you conclude the firewall filter, consider using the 'log' option. This way, when using the '*show firewall log*' command, you can have a look at all the traffic that the firewall filter drops.

Firewall filters can also be configured in such a way that configuration changes will automatically update the match condition. For example, to allow all configured BGP neighbors to be allowed by a firewall filter, you can create a prefix-list using '*apply-path*':

set policy-options prefix-list **bgp-neighbor** apply-path **"protocols bgp group <*> neighbor <*>"**

When we reference this prefix list in a firewall term that allows BGP, we will automatically allow all of the BGP peers configured under the [**protocols bgp**] stanza. To see what the prefix-list is currently matching, issue the ' show configuration policy-options prefix-list **bgp-neighbor** |display inheritance ' command.

Data plane protection

The data plane may have been decoupled from the forwarding plane since a long time, this does not mean you do not have to protect it. The data plane is still susceptible to resource depletion because of malicious attacks or because of abnormal traffic conditions. The following mechanisms can be used to protect the data plane:

Firewall filters and **policers** are the most well-known method of limiting, denying and allowing traffic. They serve to protect the control plane when applied to the lo0 interface, but they can also be applied to all other interfaces of the device as well. They can be used to permit or deny traffic that is send from A to B, rate-limit traffic to prevent resource starvation and control the type of traffic that is allowed to transit through the device. The firewall filter options vary greatly between QFX and MX. The MX is the most advanced in terms of capabilities when compared to a QFX5100.

MAC-limit: to prevent the MAC-table from being overrun, you can define a maximum number of MAC addresses per port, per VLAN/bridge-domain and for the switch as a

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

whole. On MX, you can also configure a MAC-limit for the virtual-switches that are active on the router.

MAC-movement limit: every MAC movement must be processed by a networking device and constant rapid MAC movements can deplete the resource. This knob puts a cap on the number of times a MAC address is allowed to be learned on a different interface per second.

Storm-control: large volumes Broadcast, Unknown unicast and Multicast (BUM) traffic and consume the bandwidth of an entire network. Storm control can be leveraged as a mechanism to put a hard limit to the maximum allowed levels of BUM-traffic. Note that storm control on QFX will not police link-local multicast traffic. This type of traffic can be filtered using regular firewall filters. Storm-control is not available as a configuration option on MX routers. However, the rich and extensive firewall filter possibilities allow you to create the same, or even better, type of protection from BUM traffic if needed.

DDoS features: the Junos OS offers DDoS features for both QFX as well as MX. Devices that run Junos can recognize the protocols that are being forwarded through the chassis. DDoS features work through a set of policers that can limit undesirable traffic flows before the traffic makes its way through to the RE. Control over these internal policers enables you to protect the forwarding plane at a very early state. Traffic can be discarded before is offered to the RE or before it has a chance to overrun the data plane. Both MX and QFX come preconfigured with very high values. For DDoS protection to be successful, the values will need to be tweaked. The configuration is located under [**systems ddos-protection**].

Stateful firewalling and SRX configuration

Routers and firewalls both forward packets, but the way in which they do this differs. Routers perform packet-based forwarding. This means that every packet is individually analyzed. This is followed by a route-lookup and, when a route is found, a forwarding action. Firewalls use session-based forwarding. A session is a bidirectional flow of data that can be identified based on source IP, destination IP, source port, destination port, protocol number and a session token.

For every incoming packet, the firewall checks to see if the packet is part of an existing flow. If this is not the case, the packet is sent through the firewall flow module. The following displays the packet processing order of the firewall flow module:



Simplified flow module. Screen options, ALG and NGFW are not displayed.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The flow module will first perform destination NAT. This is done before any route and forwarding lookup. After the forwarding lookup, the firewall is aware of both incoming and outgoing interfaces. This information is used to check what incoming and outgoing security zones are involved. The security zones involved determine what security policies apply to the traffic.

If traffic is permitted, the firewall applies source NAT, creates a session and installs it into the session table. It is important to note that the firewall installs a session based on this first packet. The firewall will allow the current packet, subsequent packets for the same flow and the anticipated return traffic. Sessions are only valid for a certain period of time. By default, TCP sessions will timeout after 30 minutes and UDP sessions will timeout after 1 minute.

After the session is created and installed, traffic is sent and subsequent packets that belong to the same flow can use fast processing. Fast processing uses less processing as it only involves TCP checks and NAT.

About NAT

There are three types of NAT available on SRX:

- Source NAT: translates the source IP of incoming the packet (always PAT).
- Destination NAT: translates the destination IP of an incoming packet.
- Static NAT: translates an IP address into another IP address for both incoming as well as outgoing traffic.

By default, Juniper devices perform PAT when possible. Disable PAT if this is unwanted.

About SCREENs

SCREENs are a security feature available on Junos security devices to provide for additional protection against a variety of attacks, such as flood attacks (SYN, ICMP, UDP), IP sweeps, port scans, etc. SCREENs are performed after session lookup, before any NAT translation and route lookup. SCREENs offer a means to enhance network security while saving system resources from processing the unwanted traffic at the same time.

SCREENs are defined using ids-options. Multiple ids-options can be configured for a single SCREEN. You can assign a maximum of 1 SCREEN per security zone.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: Securing the datacenter

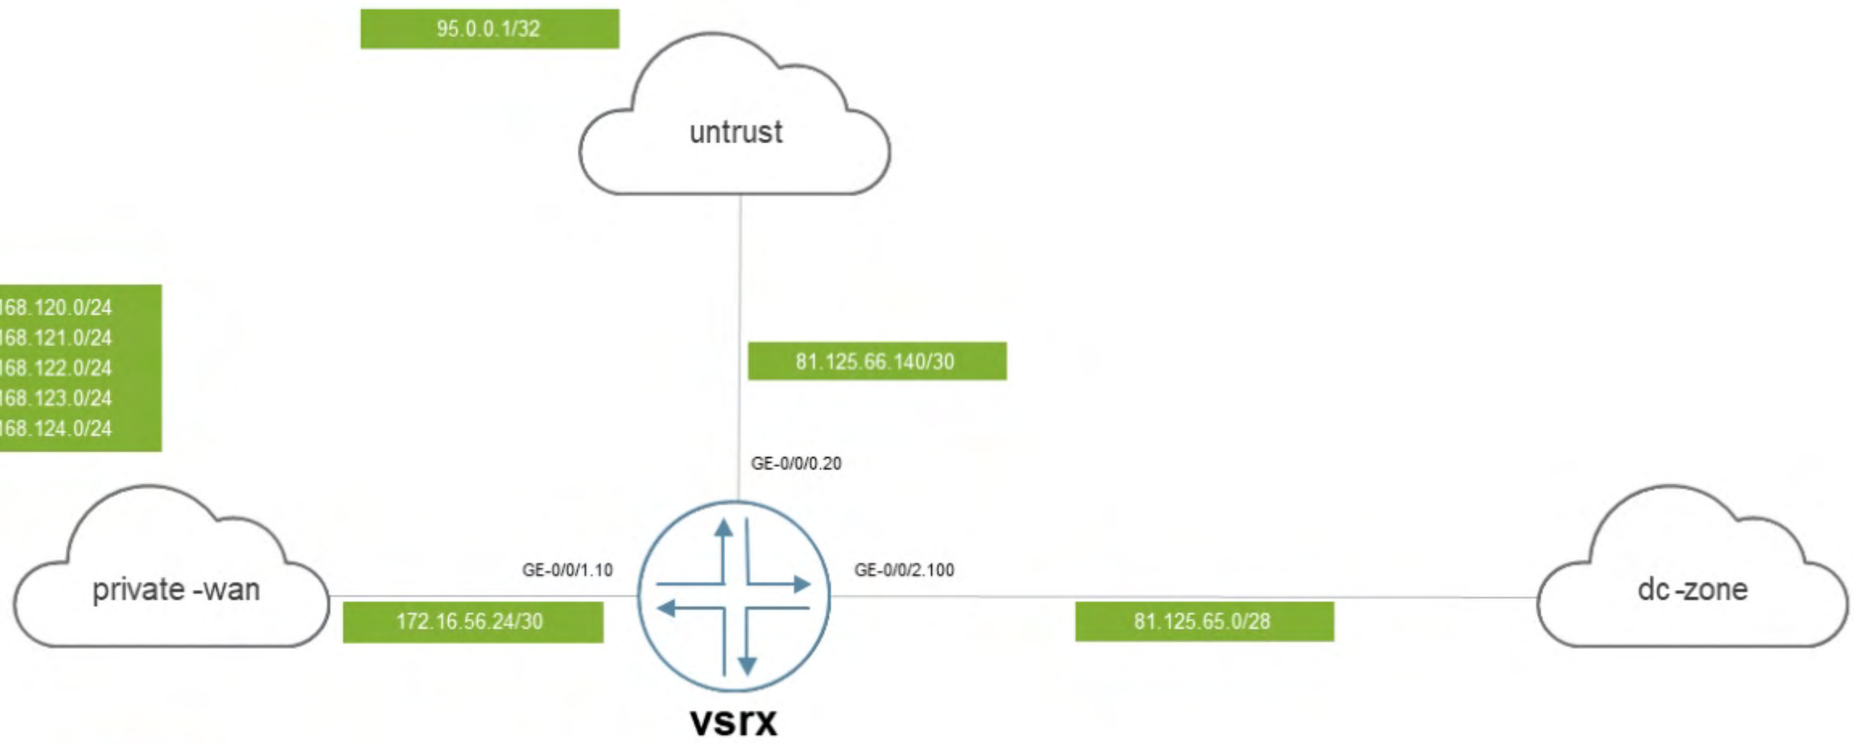

Figure 12 - Topology

Note: you can use vMX4 for verification purposes.

Task 5.1: vSRX Zone and Interface Configuration

- Configure the following zones:

Interface	IP address	VLAN-ID	Zone
GE-0/0/0	81.125.66.141/30	20	untrust
GE-0/0/1	172.16.56.25/30	10	private-wan
GE-0/0/2	81.125.65.1/28	100	dc-zone

- Users in the dc-zone and the private-wan zone should be able to ping the address configured on the vSRX.
- Configure a static default route towards the 81.125.66.142 address.
- Configure the vSRX with autonomous system 65000 and set up peering with 172.16.56.26 inside autonomous system 65010. Advertise a default route into the private-wan domain.

Task 5.2: vSRX Policy Configuration

- The hosts in the dc-zone are allowed to access the private-wan zone. The hosts can also use HTTP, HTTPS and DNS to access any resource in the untrust zone. All other traffic should be dropped and logged locally on the device.
- From the private-wan zone into the dc-zone, the management subnet 192.168.123.0/24 should be allowed to use any of the following protocols:
    - SSH, HTTP, HTTPS, ICMP, DNS and TCP 3389

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- From the private-wan zone into the dc-zone, all hosts should be allowed the following protocols:
    - HTTP, HTTPS and DNS
- From the untrust zone, any host should be able to use HTTP and HTTPS to access the webserver configured with IP 81.125.65.3.
- You are not allowed to configure any address object for a specific security-zone.

Task 5.3: vSRX NAT

- From the private-wan zone, all the 192.168.12x.0/24 subnets should be allowed to access resources inside the untrust zone. Inside the untrust zone, users should appear with the vSRX interface address.
- When users in the untrust zone try to contact 81.125.65.6, the vSRX should forward and allow that traffic towards host 192.168.123.123. The vSRX should automatically create a reverse mapping. Allow ICMP from the untrust zone towards the private-wan zone. Allow ICMP and HTTPS from the private-wan to the untrust zone.
- When users in the untrust zone try to contact a webserver located at 81.122.31.2 using HTTPS or ICMP, traffic should be forwarded to 81.125.65.2.

Task 5.4: vSRX Screens

- Protect the session table from inbound traffic from the untrust zone to any other zone. There should be no more than 500 sessions per source and per destination IP.
- The firewall should also protect resources from an ICMP flood that is greater than 500 packets per second and an IP address sweeps that are greater than 10 ICMP packets arriving at a 500ms time interval.

Task 5.5: vSRX RE protection

- Configure RE protection on the SRX using a stateless firewall filter.
- The SRX should be able set up a BGP session.
- Configure a separate filter for the monitoring and management of the switch. Allow the following protocols: SSH, FTP, NTP, Radius, DNS, SNMP, traceroute, and ICMP.
- Restrict SSH, Telnet, FTP, NTP, DNS, SNMP and Radius traffic to subnet 10.10.20.0/24 and 192.168.123.0/24.
- Rate limit both ICMP and traceroute to 1 Mbit/s. Rate limit both FTP and SSH to 5 Mbit/s.
- Discard all other traffic. A counter should be incremented for every discarded packet and the header of the discarded packet should be logged locally on the device.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 5.6: Limiting bandwidth using a Stateless Firewall Filter

- Traffic from 192.168.120.1/32 towards 81.125.65.2/32 should bypass the flow engine.
- Configure a stateless firewall filter that permits both ICMP and TCP port 4567.
- TCP traffic should be rate-limited to 100m.

Chapter 6: Class of Service

Packet loss, delay and jitter can adversely impact traffic on a packet switched network. Quality of Services, or QoS, can be used to give preferential treatment to applications and protocols. The QoS mechanisms on Juniper devices are configured through the use of the Class of Services (CoS) configuration construct. A thorough understanding of this CoS configuration construct and the order in which the CoS processing takes place are key to understanding how QoS works on Juniper devices.

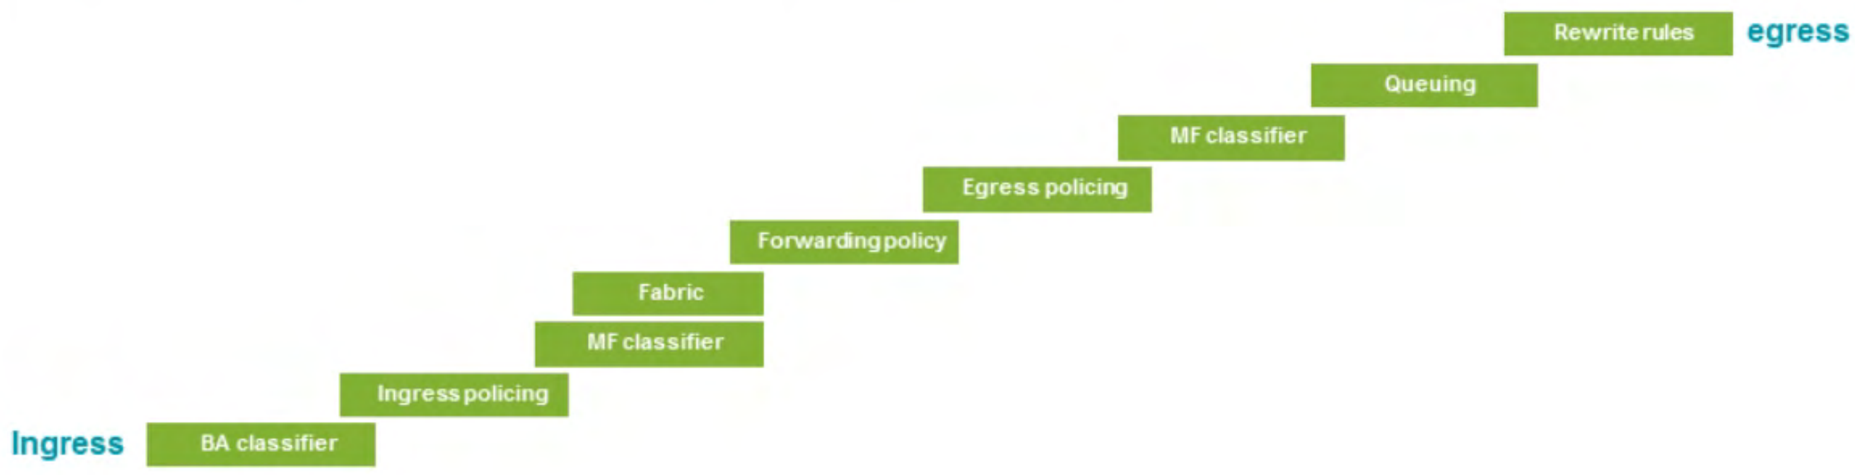

BA classifier.

The first CoS processing task is the classification of input traffic by a BA classifier. A BA classifier examines the markings on existing headers and then determines the forwarding class of those packets. The BA classifier can also be used to determine the Packet Loss Priority (PLP). There are 4 possible values: high, medium-high, medium-low and low. The 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.

Ingress policer.

After the BA classification is done, an ingress policer can be invoked to limit the incoming traffic for a configured bandwidth value. A policer uses a token bucket algorithm allowing for traffic bursts as long as there are tokens in the bucket. The tokens in the bucket are replenished during periods where there is no traffic and spend when the interface transmits packet at wire speed.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

By default, a policer used in a firewall filter is 'term specific' and instantiated specifically for the term under which the policer is configured. We can also use a 'filter specific' policer. In this case, the policer will be applied to traffic that is matched by every term inside the entire firewall filter. Apart from using a policer in a firewall filter, the policers can be applied to an interface directly.

Out-of-profile traffic can be remarked through the use of soft policers and dropped through the use of hard policers. Junos OS supports the following policers:

- Single rate two color:

This policer uses a single Committed Information Rate (CIR) and a single Committed Burst Size (CBS). Traffic exceeding the CIR and CBS can be remarked or dropped.

- Single-rate tricolor:

On top of using the CIR and CBS, this policer introduces the Excess Burst Size (EBS). Green traffic is permitted by the CIR+CBS limit. Yellow traffic exceeds the CBS but does not cross the EBS threshold. Red traffic exceeds the EBS limit. Different packet loss priorities can be applied to each of these colors.

- Two-rate tricolor:

This type of policer uses the CIR and CBS and adds to this the Peak Information Rate (PIR) and 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.

The policing capabilities vary per platform.

MF classifier.

The Multifield (MF) classifier can be used to (re)mark packets using firewall filter parameters as match criteria. The MF classifier can be used to assign packets to another forwarding class and set Packet Loss Priority (PLP) values. The MF classifier is processed after the BA classifier, enabling you to overwrite any classifications that were previously set.

Forwarding policy.

Before transit traffic is sent through the fabric of a device, it can be subjected to CoS-based forwarding. This mechanism can be used to change the way in which the outgoing interface is identified. Instead of a traditional destination IP lookup, the device can be configured to you can 'push' traffic out of a specific interface (or onto a specific LSP).

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Egress policer.

After the packets has traversed the internal switching fabric of the node, the egress CoS stage begins with the possibility to subject traffic to another policer. This policer is called the egress policer and can be used to conform traffic in the outgoing direction. The policer action takes place before all other CoS mechanisms on egress, preventing any unnecessary queuing or manipulation for traffic that is going to be discarded for being out of profile.

MF classifier.

This second MF classifier can be used to alter the way that the queue of the egress interface handles packets that were previously classified on ingress. This is useful because the ingress classification strategy cannot always consider interface from which the traffic will exit the node.

Queuing.

Every physical interface has a certain amount of hardware queues that can be mapped to forwarding classes. How these queues service traffic can be controlled through the configuration of the scheduler. Schedulers define the characteristics of a queue and can also be used to reference a drop profile. Before schedulers can be applied to an interface, a scheduler-map that connects the schedulers to a forwarding class is required. Scheduler-maps can be applied to an interface directly or they can be applied through the use of a traffic-control-profile.

- Scheduling.

Schedulers control how traffic inside a single queue is treated. Every scheduler can be configured with a **transmission rate**, **queue priority**, **buffer size** and **drop-profile**.

- Drop profiles and (W)RED.

By default, a 100% drop rate is applied when a queue fills up for 100%. This default behavior can be changed through the use of Weighted Random Early Detection (WRED). This congestion avoidance mechanism can be invoked through the use of a drop profile.

By default, a configured drop profile is **segmented**. This means the drop profile will apply the drop probability as soon as a fill level is reached, and up until the point where 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.

- The scheduler-map and the traffic control profile.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

In essence, the scheduler-map is used to connect the forwarding classes to the schedulers. The scheduler-map is nothing more than a list of forwarding classes and schedulers:

The scheduler-map can be applied to a (logical) interface directly. After doing so, the device will start to apply the listed schedulers to the forwarding classes on the interface.

Instead of applying scheduler maps directly to the interface, you can also decide to use a traffic control profile. A traffic-control profile can be used to 'unlock' the hierarchical scheduling capabilities on Juniper MX devices. This is done by configuring a traffic-control profile with an output bandwidth and applying it to a (sub)interface.

Things slightly differ on the QFX. On the QFX, the traffic-control-profile is used to be able to apply a scheduler-map to an interface.

Rewrite rules.

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 or change different CoS field values according to the packet's forwarding class and PLP. The purpose of the rewrite rule is to convey to downstream devices how each packet should be handled. After applying rewrite rules, other devices in the forwarding path should be able to perform classification using only the BA classifier.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: CoS Inside the Data Center.

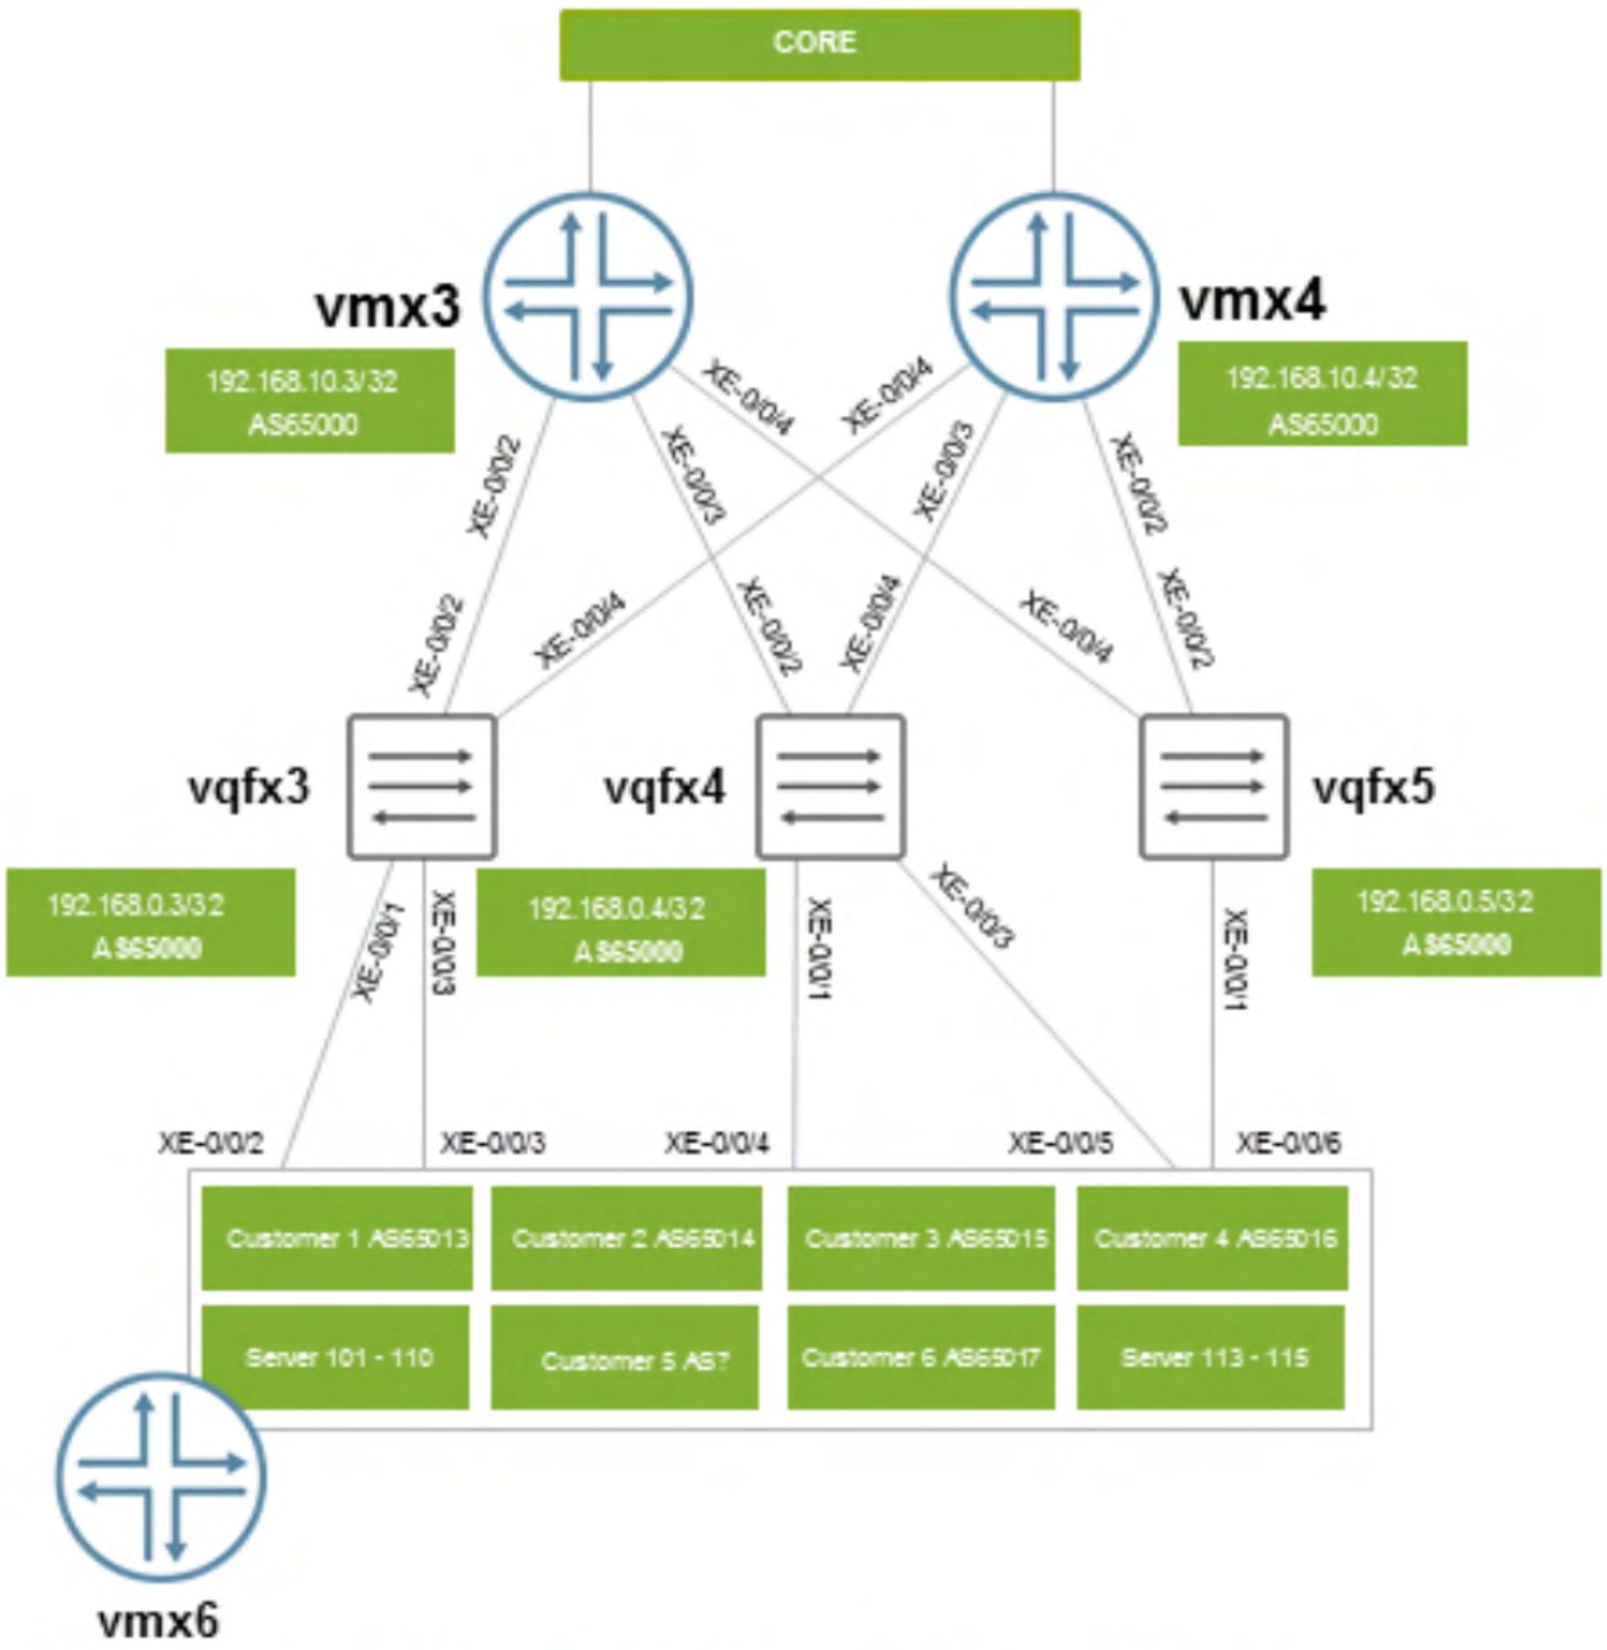

Figure 13 – DC network topology. The entire IP-Fabric has been configured for you.

Note:

There are some differences in the way a vQFX and a QFX5100 handle CoS. In the appendix, after the solutions to task 6.6, you can find a special section called 'QFX5100 Class of Service' where these differences are highlighted.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 6.1: Forwarding Classes

- Configure the following forwarding classes on both vMX as well as vQFX:

Queue number	Class name
0	be
1	data
2	data-pri
3	rt
4	control

Task 6.2: Scheduling

- Configure the following schedulers on both vQFX as well as vMX:

Scheduler name	Transmit rate	Buffer-size	Priority
be	25 percent	remainder	low
data	30 percent	30 percent	low
data-pri	30 percent	30 percent	low
rt	10 percent	10 percent	strict-high
control	5 percent	5 percent	low

- Apply the schedulers to all main interfaces.
- Ensure that it is impossible for a queue to get starved.

Task 6.3: WRED

- Configure an interpolated aggressive drop profile. When the buffer is at 50% utilization, the drop probability should be 75%. When the buffer utilization is 75%, the drop probability should be 95%.
- Configure an interpolated non-aggressive drop profile. When the buffer is at 50% utilization, the drop probability should be 30%. When the buffer utilization is 75%, the drop probability should be 50%.
- Apply the aggressive drop profile to the ‘data’ scheduler and apply the non-aggressive profile to the ‘data-pri’ scheduler.

Task 6.4: BA classifiers

- Configure classifiers for IPv4.
- Apply the classifiers to all interfaces.
- Ensure that the following classes are mapped to the mentioned code-points:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Class name	Code-points
be	be
data	af21
data-pri	af31
rt	ef
control	af11

Task 6.5: Rewrite Rules

- Make sure that when packets exit devices, their IPv4 headers have markings that correspond to the queues that serviced the packets.

Task 6.6: Traffic Classification on vQFX3 and vQFX4

- On vQFX3, rate limit inbound traffic on xe-0/0/1.10 to 2000Mbit and make sure all traffic is serviced by the be queue.
- On vQFX4 , inbound UDP traffic on xe-0/0/1.20 should be placed into the 'rt' queue. Inbound TCP traffic should be placed into the data-pri queue and rate-limited to 4000Mbit.
- On vQFX4, rate limit traffic inside VLAN 100 to 4Gbit and place traffic into the data queue. When future VLANs are added, they should not be rate-limited and traffic from those VLANs should be put into the be queue.

Task 6.7: Traffic Classification on vQFX5

- Inbound traffic on the xe-0/0/1 destined to 172.16.0.0/16 should be placed into the 'be' queue.
- Traffic destined to 192.168.0.0/16 should be placed into the 'data' queue.
- Traffic towards other destinations should be placed into the 'data-pri' queue.
- Any traffic that is destined to UDP destination port 5060 should be placed into the 'rt' queue.
- Police the 'be' and 'data' traffic to a combined total of 3gbit. Police the 'data-pri' traffic to 3gbit.

Chapter 7: Management

On box scripts

Junos OS offers different types of on-box scripting possibilities. These are 'event scripts', 'commit scripts' and 'operational scripts'. These scripts can be written in XSLT (Extensible Stylesheet Language Transformations), SLAX (Stylesheet Language Alternative Syntax), or Python. A complete understanding of the scripts is out of scope for the JNCIE-DC exam.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

What is important to the exam are the purposes of these different scripts, the way to configure them and a basic understanding of how a simple script is created.

Event scripts.

Event scripts can be used to have Junos perform an action based on an event. An event script is only executed when a certain event occurs.

To determine what event can trigger the event policy, an event policy needs to be configured. This can be done under in [**event-options policy**] stanza. The event policy contains an 'if', which is the event

(system log message, SNMP trap, configured RPM probe), and a 'then', specifying what action that Junos should take. Amongst other, an event policy can be used to have Junos take the following actions:

- The execution of operational mode commands.
- Altering the device configuration.
- Raising an SNMP trap.
- Upload a file to a specific destination.
- Executing a Junos event script.

Apart from referencing it inside an event policy, there are two things to take care of before the event script referenced inside an event-policy can be executed. First of all, the script needs to be stored on the hard disk in the '/var/db/scripts/event' folder or on a flash drive in the '/config/scripts/events' directory.

Second is that the script has to be enabled. Enabling the script happens outside the event policy inside the [event-options event-script file] hierarchy.

To configure an event policy to match an event, or a list of events, you configure the following:

```
set event-options policy policy-name events list-of-events
set event-options policy policy-name attributes-match list-of-attributes
set event-options policy policy-name then event-script script.slax
set event-options event-script file script.slax
```

Note that for both SLAX or Python scripts, you have to include the .slax or .py filename extension.

Commit scripts.

After issuing a commit, the Junos OS will check configuration for syntax and it will perform several additional sanity checks. For instance, if you have BGP configured, you must also

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

configure an autonomous-system. Because a commit script is executed right after a commit and just before the configuration is read by the daemons, commit scripts can be used to enhance your network with additional sanity checks. For instance, if a loopback interface was added to a routing-instance, was there a firewall filter applied to it as well?

A commit script can be configured to warn a user or to completely abort the commit process. In addition to this, commit scripts can also be used to make permanent or transient changes.

A commit scripts is configured at the [**edit systems scripts commit**] configuration stanza. Upon commit, the management daemon will initiate every commit script that is configured.

Commit scripts need to be stored on the hard disk in the '/var/db/scripts/commit' folder or on a flash drive in the '/config/scripts/commit' directory.

Operational scripts.

Operation scripts are used in operational mode and can be used to help automate network troubleshooting or other operational tasks. To enable an operational script, configure the name of the script under the [**system scripts op**] hierarchy. Once configured, the operational script can be executed using the 'op *script-name*' operational mode command.

Operational scripts need to be stored on the hard disk in the '/var/db/scripts/op' folder or on a flash drive in the '/config/scripts/op' directory.

Junos Space

The Junos Space network management platform provides fault, configuration, accounting, performance and security management (FCAPS) capabilities. The platform can be installed on a server or it can be deployed as a virtual appliance. Depending on the network's requirements, additional applications can be installed on the Junos Space management platform. Additional applications, such as the network director or the security director can offer additional management features. For the JNCIE-DC exam, the two most important Junos Space topics are the device discover and the configuration management.

Auto-discover

You can use device-discovery to find a device and synchronize its inventory and configuration with Junos Space. For Junos Space to be able to discover a device, there needs to be IP-connectivity between the device and the Junos Space server. For successful discovery, you start off specifying the target devices that need to be added to Space. You can insert a single IP, a range of IP addresses, an IP subnet or a host name. During the device discovery process, you will also need to add the target device's

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

administrator credentials. Junos Space will use the administrator login to establish an SSH connection with the device. In addition to this, you can also choose to use SNMP (V1/V2C or V3).

Configuration management.

A copy of the configuration is kept for every device that is managed by Junos Space. Through Space, you can modify the configuration files, keep a record of configuration changes and restore configuration files.

Zero Touch Provisioning

Initial provisioning of QFX switches using a CentOS DHCP server.

Using ZTP, you can automate the initial setup of a switch by ensuring that the correct configuration and software version are loaded to the switch. The QFX default configuration will have the switch try an obtain an IP address through the use of DHCP. The management interface, as well as the irb.0 interface that is associated with the default VLAN, are used to send out a DHCP discover message. By default, the QFX will have all non-management ports placed in the default VLAN and the DHCP discover message will be send out every port that is up.

When the DHCP discover messages reach the DHCP server, it will reply with a DHCP offer. This offer needs to contain certain parameters to instruct the switch to start ZTP and to tell the switch what parameters it should use. The following are some of the DHCP (sub)options that a QFX can use:

DHCP (sub)option	**Use case:**
42, sub option 00 or 04	To tell the switch what software version to use.
43, sub option 03	Instructions to what transfer mode the switch should use (FTP, TFTP or HTTP).
66 or 150	Informs the switch about the IP address where the files can be fetched.

After receiving the required information via DHCP, the QFX will first download and install the new software (if required). After this, the switch will apply the desired configuration file. If no configuration file was specified, the switch will keep its factory default configuration.

CentOS will be the operating system of choice on the exam. Even though this is not a Linux exam, there are a few useful things to know about running a DHCP server on CentOS.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

First of all, the directory where CentOS (or rather all Linux distributions) keeps the configuration files is the "**/etc**" directory. Inside this directory, you will find the DHCP and FTP configuration files.

The second thing is that you will need to be able to perform at least the following basic tasks on the operating system:

Linux command	
ls	To tell the switch what software version to use.
service srvcname *status*	Check the status of a daemon/process.
service srvcname *restart*	Restart the daemon/process.
cat filename	Display the content of a file on screen.
vi filename	Open a file using the vi text editor

Last but not least, you have to be able to quit vi. To exit from vi, you first have to get into vi ex mode. You can do this by pressing [Esc] followed by a : (colon). Then, you simply type 'wq:'. This will quit the vi text editor while saving changes to the file you were working on.

NETCONF

NETCONF, or the Network Configuration Protocol, is defined in IETF RFC 6241). The RFC mentions that NETCONF 'provides mechanisms to install, manipulate, and delete the configuration of network devices. It uses an Extensible Markup Language (XML)-based data encoding for the configuration data as well as the protocol messages. The NETCONF protocol operations are realized as remote procedure calls (RPCs).'

You can choose to use NETCONF, but instead of opening a direct NETCONF connection with a Junos device, you can also choose to use Junos PyEZ. PyEZ is a Python library that can act as an abstraction layer on top of NETCONF. This way, the network administrator only has to interact with PyEZ.

The PyEZ library offers you a set of modules that will enable you to manage Junos devices without having extensive programming experience. Through the use of PyEZ, you are able to connect to devices via a serial telnet, a console connection, or NETCONF over SSH.

To be able to start using PyEZ to connect to a device, you will need to enable NETCONF over SSH. This can be done under the [system services netconf] stanza.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

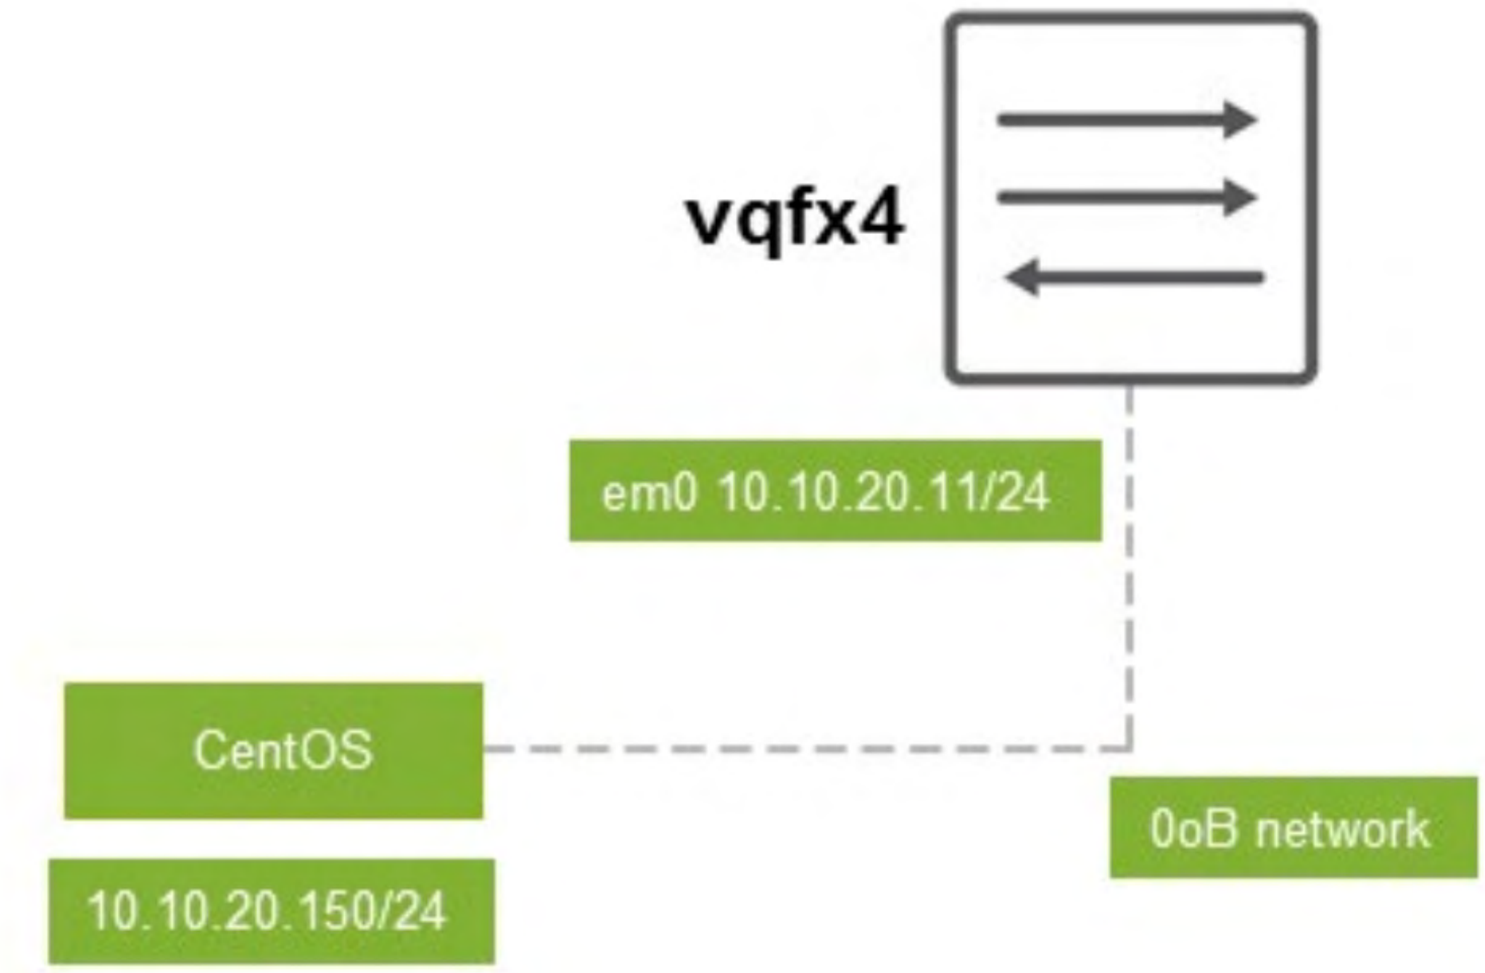

Figure 1 – ZTP and NETCONF.

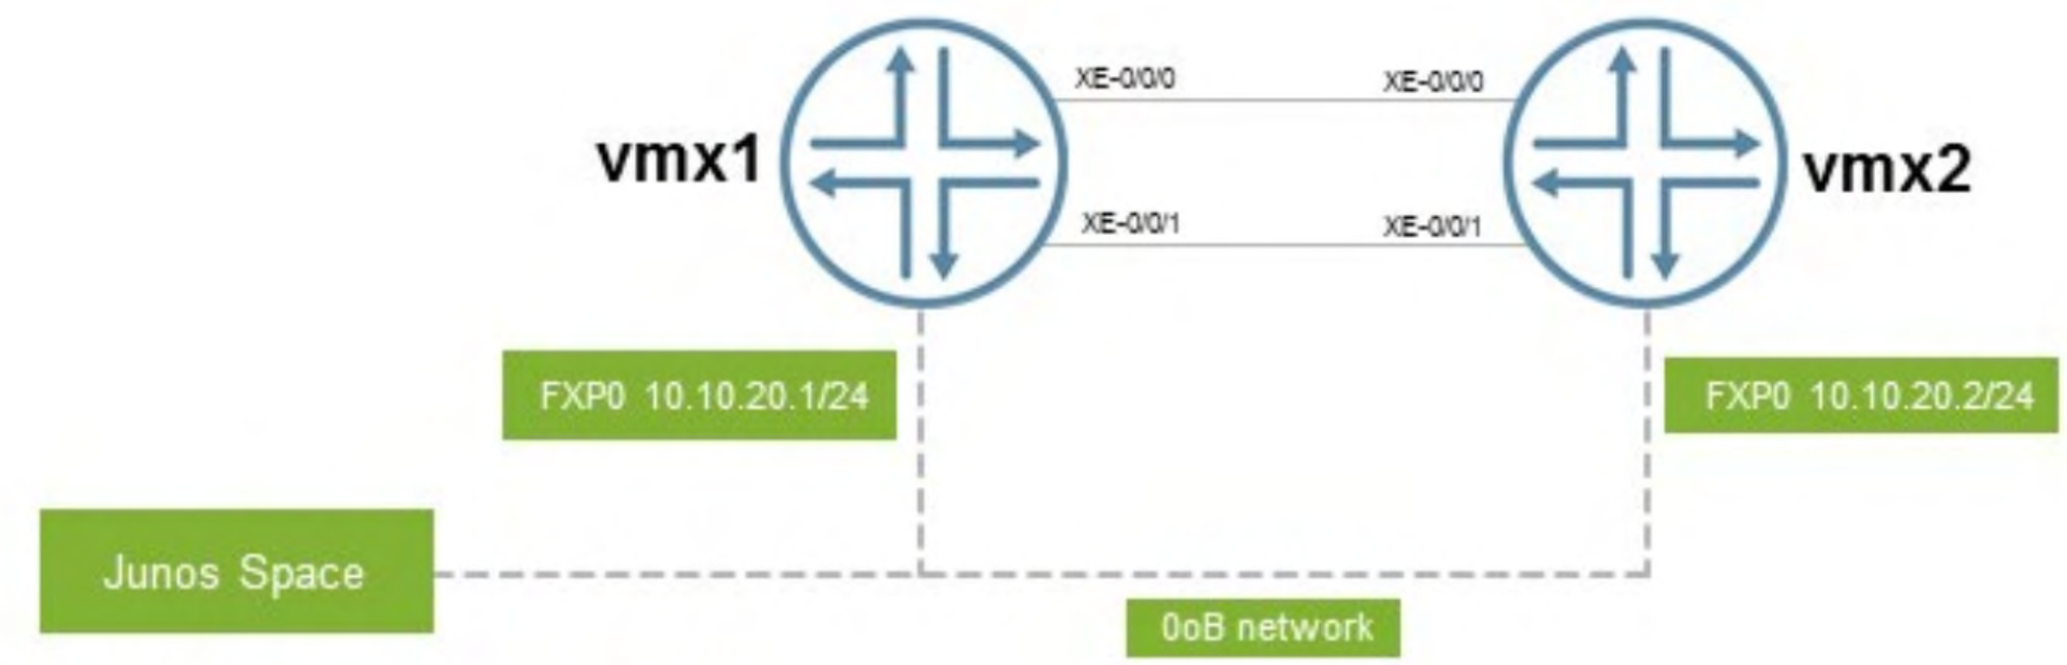

Figure 2 – Junos space and on box scripting.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 1: Data Center Management

Task 7.1: Zero Touch Provisioning (ZTP)

- Ensure that vQFX4 is able to perform ZTP using interface em0.
- The DHCP server that is running on the CentOS host should assign the 10.10.20.11 IP address based on the MAC address of vQFX6 em0 interface.
- During ZTP, vQFX4 should automatically download the configuration file /var/pub/ztp/vqfx4 using FTP from the Centos host.
- Edit the */etc/dhcp/dhcpd.conf* file using the "vi" editor. Do not use another editor.

Task 7.2: Python using NETCONF

- On the CentOS host, create a Python script called "jncie.py" that retrieves the routing table of device vQFX4. You may adjust the following example script:

```
from jnpr.junos import Device
from jnpr.junos.op.routes import RouteTable
from lxml import etree
username='root'
password='root1'
dev = Device(host='10.10.20.4', user=username, password=mypassword, gather_facts='false')
dev.open()
routes = RouteTable(dev).get()
print(routes)
dev.close()
```

- Make sure that the Python script is able to execute using the following syntax:

  /home/jncie/jncie.py

Task 7.3: Junos Space

- Ensure that vMX1 and vMX2 are discovered using SNMP version 1/2c. Use SNMP community "jncie"
- Make sure that the configuration files of vMX1 and vMX2 are stored on the Junos Space server.

Task 7.4: On Box Scripting 1

- Configure an OSPF adjacency without a DR/BDR election between vMX1 and vMX2 on interface xe-0/0/0. Hello interval and dead interval should be set to the default values / 10.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Ensure that interface xe-0/0/0 on vMX1 is disabled by the Junos OS when the OSPF adjacency on interface xe-0/0/0 fails two times in one minute.
- Ensure that, when the Junos OS disables link xe-0/0/0 in vMX1 due to OSPF failures, an XML file is created in the /tmp directory of vMX1. The file must list all active OSPF adjacencies.

Task 7.5: On Box Scripting 2

- Configure an OSPF adjacency without DR/BDR election between vMX1 and vMX2 on interface xe-0/0/0 and xe-0/0/1. Hello interval and dead interval should be set to the default values / 10.
- Enable the LDP protocol on interface xe-0/0/0 and xe-0/0/1 for vMX1 and vMX2.
- Ensure that an event-script called “jncie.slax” is executed on vMX2 when the LDP neighbor relationship on interface xe-0/0/0 or interface xe-0/0/1 goes down twice in one minute. The failed interface name must be passed to the jncie.slax script.
- Ensure that the jncie.slax script is paused for 1/2 second after each executed line when the CPU usage of vMX2 is 40% or more.

Note: *For this task you can use the predefined script named jncie.slax stored in the /var/db/scripts/events/ directory. The point for this task is to understand how to execute an event-script policy that calls a script.*

```
% pwd
/var/db/scripts/event
% cat jncie.slax
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "interface";
  }
}
param $interface;
match / {
  <event-script-results> {
    var $message = "The " _ $interface _ " interface is causing LDP to bounce more than two
times in the one minute";
    expr jcs:syslog("pfe.notice" , $message);
  }
}

%
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 7.6: On Box Scripting 3

- Add the following commit script with name “jncie2.slax” on device vMX1
- Ensure that if “family mpls” is configured on interface xe-0/0/0 that the interface is disabled

```
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    for-each (interfaces/interface[name != 'xe-0/0/0']/unit) {
      if ( (family/iso)) {
        call jcs:emit-change () {
            with $content= {
              <disable>;
            }
        }
      }
    }
}
```

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-DC lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

You are going to configure several vMX, vQFX and vSRX devices. The DCI PE routers are simulated and preconfigured on vMX5 using Logical systems with routing-instances. The server with IP address 8.8.8.8 is also preconfigured on vMX5 in the global routing table.

All devices have preconfigurations and might contain errors that you need to correct. Some tasks are time consuming. We recommend using a text editor to copy and paste configurations between devices.

It is not allowed to created static routes unless explicitly mentioned for certain tasks.

There are several servers as shown in the topology. Server 1 (S1) in DC1, Server 2 (S2), Server 3 (S3) and Server 4 (S4) in DC2. Server 5 (S5) in the Main DC. These servers are simulated on vMX6.

You can login to vMX6 to test connectivity between servers and towards the DNS server.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Example 1:

Ping towards Server 1 on VLAN 40 with IP address (**10.150.10.4**) in DC1 sourced from Server 5 on VLAN 10 with IP address (10.150.10.5) in MAIN DC

```
jncie@vMX6> ping 10.150.10.4 routing-instance S5-10

PING 10.150.10.4 (10.150.10.4): 56 data bytes
64 bytes from 10.150.10.4: icmp_seq=0 ttl=64 time=73.167 ms
```

Example 2:

Traceroute to the DNS server (8.8.8.8) from Server 1 on VLAN 30 with IP address (10.150.30.1) in DC1

```
jncie@vMX6> traceroute 8.8.8.8 routing-instance S1-30

traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 52 byte packets
 1 10.150.30.100 (10.150.30.100) 142.394 ms 165.814 ms 165.432 ms
 2 10.30.4.5 (10.30.4.5) 47.652 ms 10.30.3.5 (10.30.3.5) 80.761 ms 112.557 ms
 3 10.40.2.1 (10.40.2.1) 96.291 ms 91.697 ms 52.271 ms
 4 10.40.2.2 (10.40.2.2) 70.905 ms 41.330 ms 43.164 ms
 5 8.8.8.8 (8.8.8.8) 87.411 ms 93.586 ms 96.363 ms
```

Topology

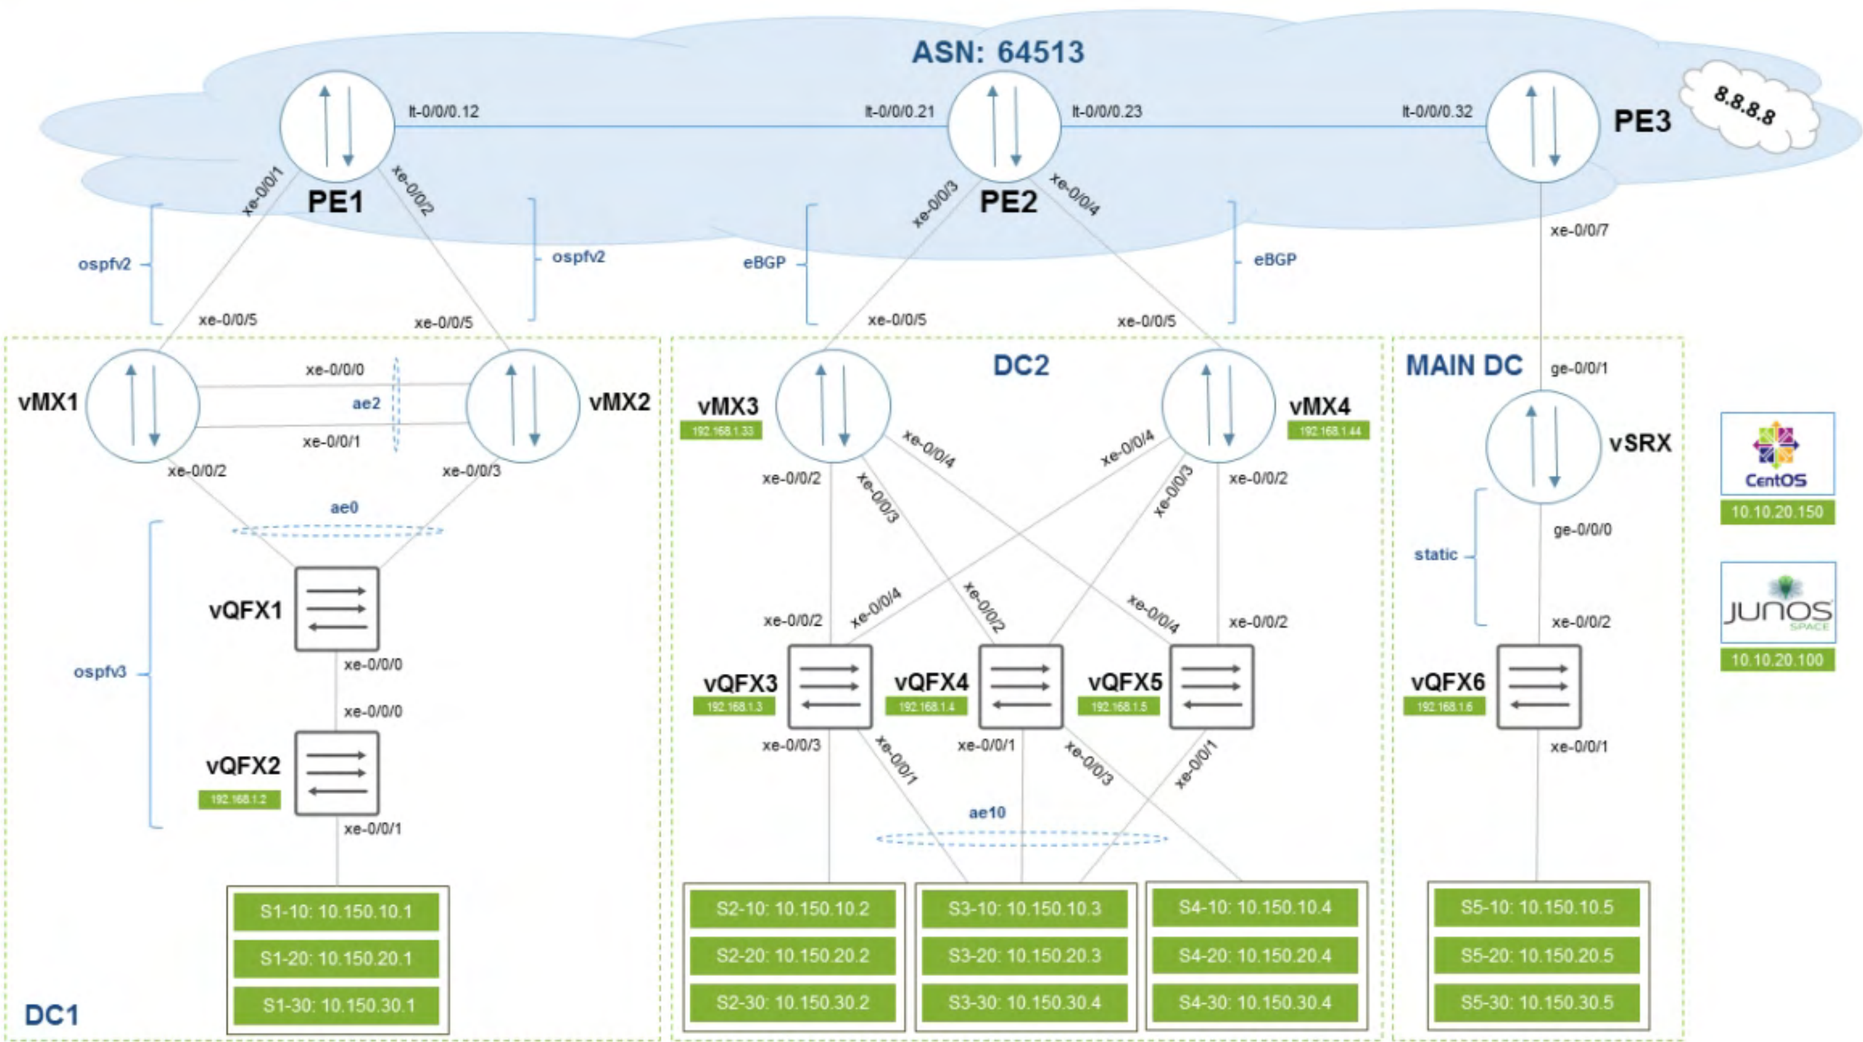

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1: Device management

Task 1.1: Zero Touch Provisioning (ZTP)

- Ensure that vQFX1 is able to perform ZTP using interface em0.
- The DHCP server that is running on the CentOS host should assign IP address 10.10.20.8 based on the MAC address of vQFX1 em0 interface.
- During ZTP, vQFX1 should automatically download the configuration file /var/pub/ztp/vqfx1 using FTP from the Centos host.
- Edit the /etc/dhcp/dhcpd.conf file using the “vi” editor. Do not use another editor.

Task 1.2: Junos Space

- Ensure that vMX1 is discovered in Junos Space using SNMP version 1/2c. Use SNMP community “jncie”
- Make sure that the configuration file of vMX1 is stored on the Junos Space server.

Task 1.3: On Box Scripting

- Ensure that interface AE1 on vMX1 is disabled by the Junos OS when the OSPF adjacency with neighbor 10.20.121.3 fails 3 times within 2 minutes. Do not use an external script
- Ensure that on-box Python scripts are allowed on all vMX devices in DC1 and DC2
- 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 1.4: NETCONF

- There is a Python script that uses the RFC 4741 compliant NETCONF protocol to manage the configuration of vMX1. Ensure that the Python script can connect to vMX1 over a secured shell session. Allow 6 concurrent connections and a maximum of 2 sessions per minute.

Task 2: Underlay

Task 2.1: Multi Chassis Link Aggregation (MC-LAG)

Configure an active/active Multi Chassis LAG on MX1 and MX2 according to the following requirements:

- Interface AE0 towards vQFX1 should use VLAN 121. vMX1 should use IP address 10.20.121.1/24, vMX2 should use IP address 10.20.121.2/24 and vQFX1 should use IP address 10.20.121.3/24
- Ensure that multi chassis protection is configured

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Configure a redundant ICL link. Use VLAN 12 for the ICCP connection. vMX1 should use IP address 10.21.12.1/30 and vMX2 should use IP address 10.21.12.2/30.
- Configure a redundant default gateway with IP address 10.20.121.254. vMX1 should always be the master, also when recovered from a failure. Use MD5 password “jncie”.

Task 2.2: Layer 2 Underlay Configuration

Configure the Layer 2 interfaces between the vQFX devices and the servers according to the following requirements:

- Configure interface xe-0/0/1 between S1 and vQFX2. VLAN 10, 20 and 30 should be enabled on this interface
- Configure interface AE10 between S3 and vQFX3, vQFX4 and vQFX5. Ensure that vQFX3, vQFX4 and vQFX5 initiate LACP sessions. VLAN 10, 20 and 30 should be enabled on this interface
- Configure interface xe-0/0/3 on vQFX3, interfaces xe-0/0/3 on vQFX4 and interface xe-0/0/1 on vQFX6. VLAN 10, 20, and 30 should be enabled on these interfaces.

Task 2.3: Layer 2 Underlay Features

- Ensure traffic towards Server 10.150.10.5 is flooded within VLAN 10.
- Configure VQFX1 VLAN 121 with a mac-address limit. Set the limit to the number of mac-addresses used by vMX1 and vMX2 in VLAN 121. Log a message and drop packets if the limit is reached.

Task 2.4: BUM Traffic Handling

- Ensure that only Broadcast and Unknown Unicast traffic is limited to 100Mbit on vQFX2
- Do not use the default profile
- Use an apply-group to activate the configuration for all 10g layer2 interfaces on vQFX2.

Task 2.5: Layer 2 Redundancy

- Ensure that interface AE10 towards Server 3 keeps forwarding traffic for VLAN10, 20 and 30 if Server3 cannot exchange LACP PDUs during a server upgrade.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.6: Layer 3 Underlay MAIN DC

- Configure an untagged IP interface between vSRX and VQFX6. vSRX1 should use IP address 10.40.3.1/30 and vQFX6 should use IP address 10.40.3.2/30. Use only static routing between vQFX6 and vSRX.

Task 2.7: Layer 3 Underlay Routing

Configure OSPFv2 between vMX1, vMX2, vQFX1 according to the following requirements:

- All interfaces should be configured in the backbone area
- Ensure that OSPFv2 adjacencies between vQFX1, vMX1 and vMX2 are in the "full" state.

Configure OSPFv2 between vQFX1 and vQFX2 according to the following requirements:

- All interfaces should be configured in the backbone area
- Ensure there is no DR/BDR election on interface between vQFX1 and vQFX2.
- vQFX1 should use IP address 10.20.1.1/30 and vQFX2 should use IP address 10.20.1.2/30
- Do not use VLANs

Task 2.8: EBGP Underlay

Configure the IP fabric in DC2 according to the following requirements:

Device	BGP AS
vMX3	64523
vMX4	64524
vQFX3	64513
vQFX4	64514
vQFX5	64515

- Configure BFD on all EBGP sessions for all devices in DC2s IP fabric.
- Ensure that BFD sessions do no support BGP sessions that peer on loopback interfaces
- Enable BFD MD5 authentication with password "jncie"
- Ensure that traffic is load-balanced if there are equal cost paths towards the same destination.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3: Controllerless Overlay

Task 3.1: EVPN Overlay Configuration

Configure an EVPN overlay between DC1, DC2 and MAIN DC according to the following requirements:

- Use BGP AS: 65535
- vQFX2, vQFX3, vQFX4, vQFX5 and vQFX6 should only have EVPN overlay sessions with vMX3 and vMX4.
- vMX3 and vMX4 should **not** initiate BGP sessions used for the EVPN overlay. Do not use the "passive" command or a firewall filter to accomplish this task.

Task 3.2: EVPN/VLAN Service Configuration

- Configure an EVPN service called "**vpn1**" for VLAN 10 and VLAN 20.
- Configure an EVPN service called "**vpn2**" for VLAN30.

EVPN/VLAN service	**VLAN**	**vMX3 gateway address**	**vMX4 gateway address**
vpn1	10	10.150.10.100 10.150.10.254	10.150.10.101 10.150.10.254
	20	10.150.20.100 10.150.20.254	10.150.20.101 10.150.20.254
vpn2	30	10.150.30.100 10.150.30.254	10.150.30.101 10.150.30.254

Ensure that the following communities are used for each EVPN/VLAN service as shown in the following table:

EVPN/VXLAN service	**Community name**	**Community value**
vpn1	evpn	1:1
	vpn1	1:1000
	vni-10	1:10
	vni-20	1:20
vpn2	evpn	1:1
	vpn2	1:2000
	vni-30	1:30

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.3: EVPN/VXLAN Service Tuning

- Configure vQFX6 to age out the address of the remote VTEP after 6 minutes, if all the MAC addresses learned from that VTEP age out. Your solution must also account for VNIs that might be added in the future.

Task 3.4: EVPN Signaling

- Ensure that vQFX2 receives the following EVPN type 1 routes:

```
1:192.168.1.3:1::11111111111111122::0/304
1:192.168.1.4:1::11111111111111122::0/304
1:192.168.1.5:1::11111111111111122::0/304
```

- Ensure that vQFX2 receives the following EVPN type 2 routes:

```
2:192.168.1.33:2::30::00:99:00:00:00:10/304
2:192.168.1.33:1::20::00:99:00:00:00:20/304
2:192.168.1.33:1::10::00:99:00:00:00:30/304
```

Task 4: Data Center Interconnect

The PE devices as shown in the topology are simulated on vMX5 using preconfigured logical systems. The interfaces between the PE devices are also preconfigured.

Logical system name	Interface	IP address	Remote end
PE1	lt-0/0/0.12	172.16.12.1/24	PE2
PE2	lt-0/0/0.21	172.16.12.2/24	PE1
	lt-0/0/0.23	172.16.23.2/24	PE3
PE3	lt-0/0/0.32	172.16.23.3/24	PE2

Task 4.1: DCI Configuration

Configure two L3VPN services using the following information:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- The first L3VPN service should be named “underlay” on vMX5 and is used for the underlay between DC1, DC2 and MAIN DC.

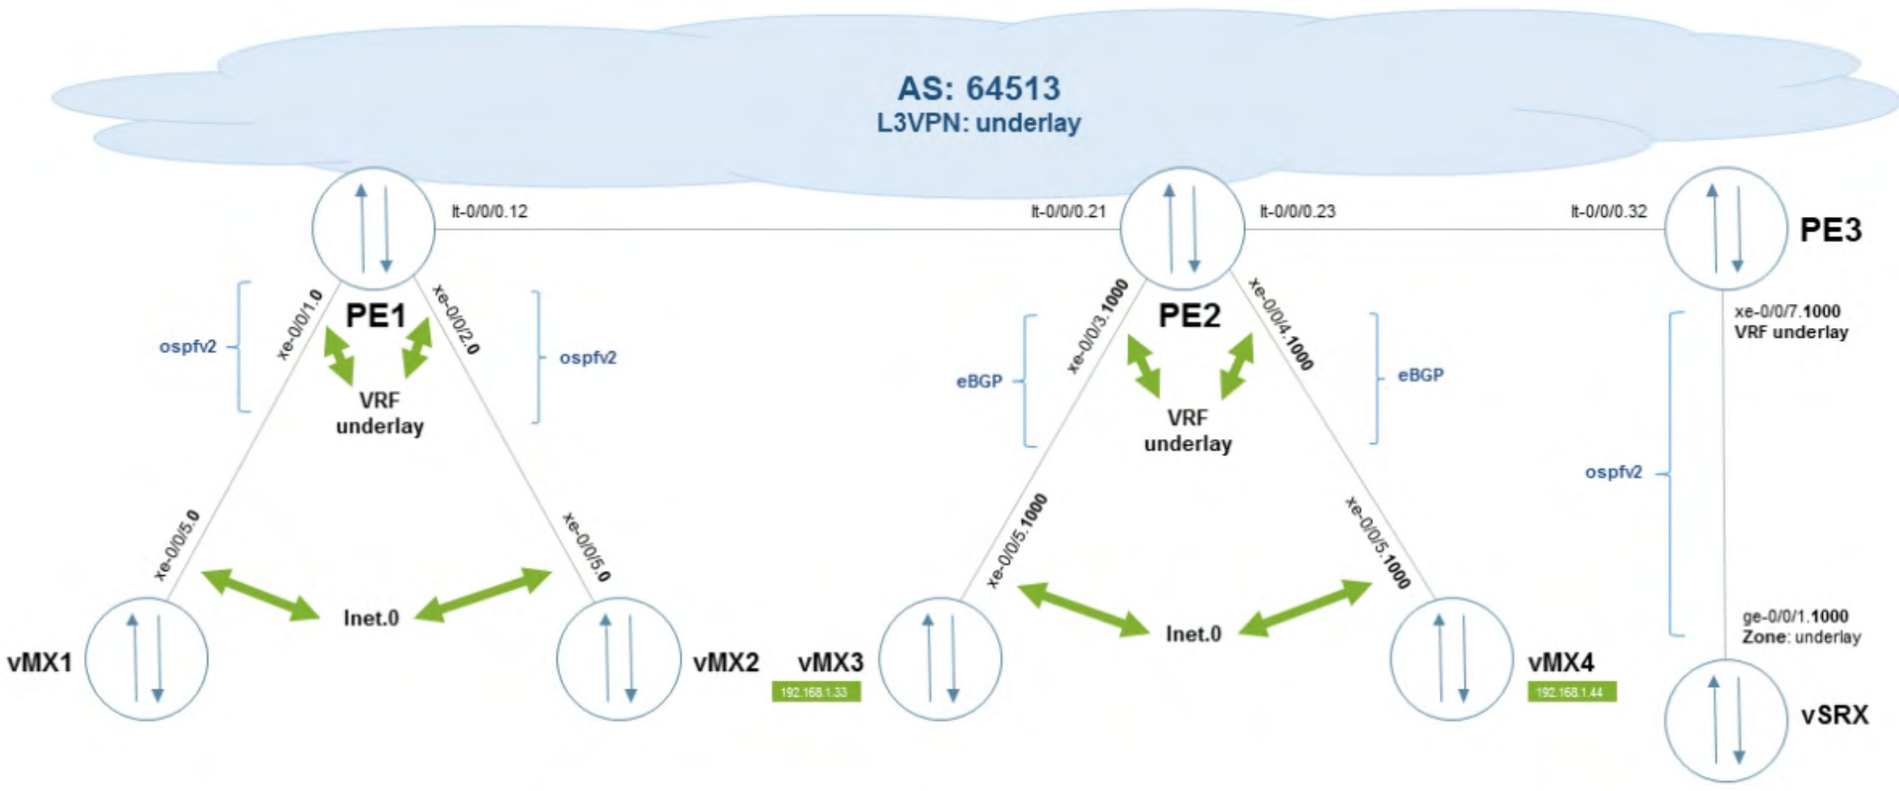

The following table shows the devices and interface configurations needed for the DCI. Some of the interfaces are already preconfigured.

Device	Interfaces VLAN = unit	PE-CE routing	IP address	Additional information
vMX1	xe-0/0/5.0 untagged	OSPFv2	10.20.1.2/30	Global routing (underlay)
vMX2	xe-0/0/5.0 untagged	OSPFv2	10.20.2.2/30	Global routing (underlay)
vMX3	xe-0/0/5.1000	EBGP	10.30.3.2/30	Global routing (underlay)
vMX4	xe-0/0/5.1000	EBGP	10.30.4.2/30	Global routing (underlay)
vSRX	ge-0/0/1.1000	EBGP	10.40.1.2/30	Security zone: (underlay)
vMX5 (PE1)	xe-0/0/1.0	OSPFv2	10.20.1.1/30	Routing instance: underlay
	xe-0/0/2.0	OSPFv2	10.20.2.1/30	Routing instance: underlay
vMX5 (PE2)	xe-0/0/3.1000	EBGP	10.30.3.1/30	Routing instance: underlay
	xe-0/0/4.1000	EBGP	10.30.4.1/30	Routing instance: underlay
vMX5 (PE3)	xe-0/0/7.1000	EBGP	10.40.1.1/30	Routing instance: underlay

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- The second L3VPN should be named “dc-vpns” on vMX5 and is used for L3 routing for EVPN service “vpn1” and “vpn2” between DC2 and MAIN DC.

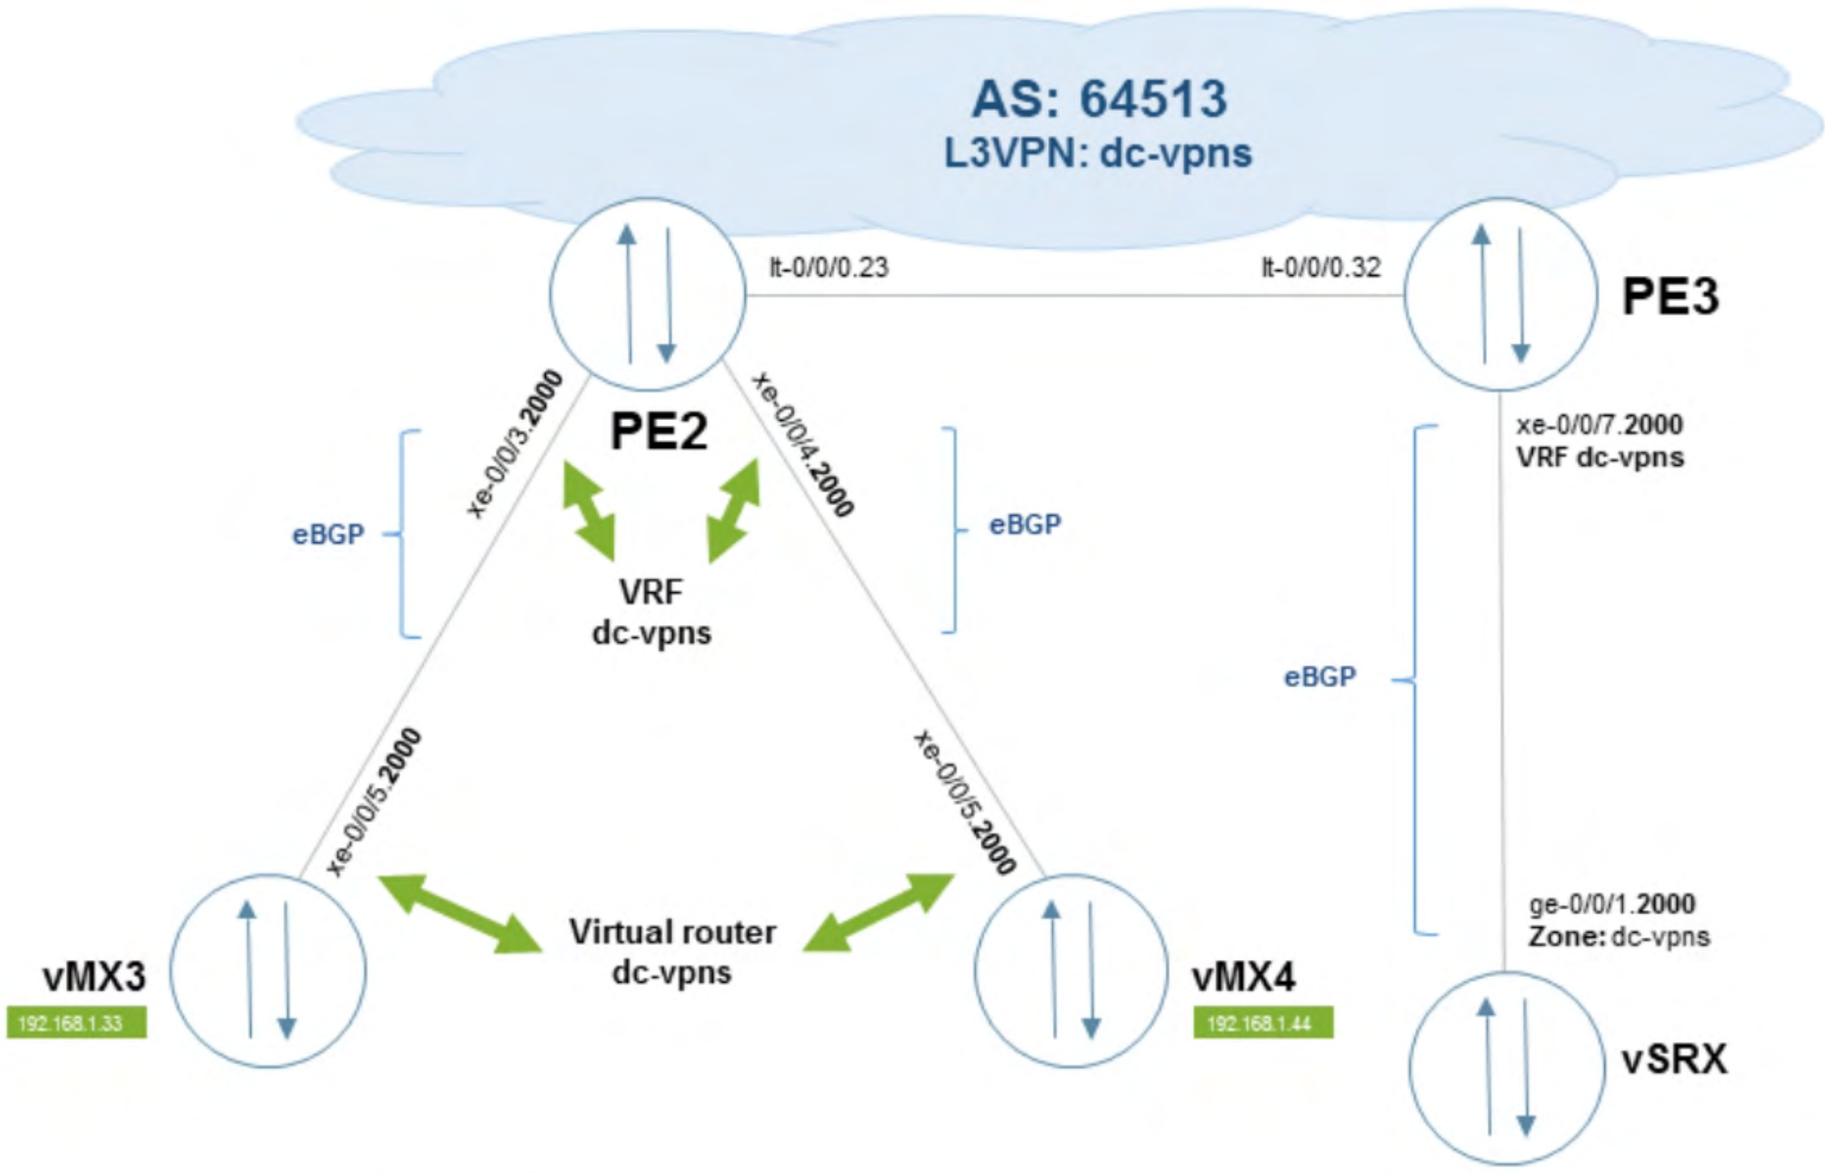

The following table shows the devices and interface configurations needed for the DCI. Some of the interfaces are already preconfigured

Device	Interfaces VLAN = unit	PE-CE routing	IP address	Additional information
vMX3	xe-0/0/5.2000	EBGP	10.30.3.6/30	Routing instance: dc-vpns
vMX4	xe-0/0/5.2000	EBGP	10.30.4.6/30	Routing instance: dc-vpns
vMX5 (PE2)	xe-0/0/3.2000	EBGP	10.30.3.5/30	VRF: dc-vpns
	xe-0/0/4.2000	EBGP	10.30.4.5/30	VRF: dc-vpns
vMX5 (PE3)	xe-0/0/7.2000	EBGP	10.40.2.1/30	VRF: dc-vpns

Task 4.2: DC1 Redundancy

Ensure that the link between PE1 and vMX1 is the primary path for traffic to and from DC1. Ensure that the link between PE1 and vMX2 is only used for traffic to and from DC1 if the primary path is down.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.3: (Inter) Data Center Connectivity

Ensure that the Servers in DC1, DC2 and MAIN DC have full connectivity between all VLANs. Use the test addresses shown in the topology for verification purposes

Task 4.4: Data Center Connectivity to External Server

Ensure that the Servers in DC1, DC2 and MAIN DC can ping the external server (8.8.8.8). The server is preconfigured on vMX5.

- Traffic between the servers in DC1, DC2 , MAIN DC and the external server should always transit the “dc-vpns” and “untrust” zones of the vSRX firewall.
- The server should use destination IP address 1.1.1.2 if it sends ICMP echo replies to Servers in DC1, DC2 and MAIN DC.

Task 5: Security

Task 5.1: User Authentication and Authorization

Configure a user called “jncie-1” on vSRX. Helpful information should be shown when user “jncie-1” logs in to the device. User “jncie-1” should only be allowed to display the device configuration and view system and interfaces statistics

Task 5.2: SRX Configuration

Configure the vSRX according to the following requirements:

- Interface ge-0/0/0.0 and ge-0/0/1.1000 should be in the “underlay” zone
- Interface ge-0/0/1.2000 should be in the “dc-vpns” zone
- Interface ge-0/0/1.3000 should be in the “untrust” zone
- Configure a static default route towards the next-hop IP address (1.1.1.1) of AS 64513 in the untrust zone.
- Allow ICMP PING traffic towards vSRX interfaces, except for interfaces in the “untrust” zone
- Allow ICMP traffic between devices in the “underlay” zone
- Do not configure routing-instances on vSRX
- Only allow traffic if required to complete other tasks in this superlab. Do not allow other traffic

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 5.3: RE Protection

Configure control-plane protection on vQFX6 with the following requirements:

- Allow incoming telnet and ssh sessions
- Allow BGP sessions to vMX3 and vMX4. Make sure that if new BGP neighbors are configured in the “protocols bgp” stanza, you do not need to change your RE protection configuration.
- Allow ICMP
- Allow specific traffic needed by other tasks in this superlab.
- Deny and count all other traffic. The control plane should not respond to denied traffic.

Task 6: Class of Service

Task 6.1: Forwarding Classes

Configure the following forwarding classes and forwarding class sets on all vQFX devices

Forwarding Class	Queue ID	No-Loss
BE	0	Disabled
FCOE	3	Enabled
NOLOSS	4	Enabled
NC	7	Disabled

Forwarding Class Set	Forwarding Class
LOSS	BE
	NC
NOLOSS	FCOE
	NOLOSS

Task 6.2: Schedulers and Traffic Profiles

- Configure a scheduler mapping named “LOSS” that includes the BE and NC schedulers on all vQFX devices.
- Configure a scheduler mapping named “NOLOSS” that includes the FCOE and NOLOSS schedulers on all vQFX devices.
- Configure traffic profiles for each scheduler

Scheduler name	Transmit rate	Buffer size	Priority

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

BE	55 %	remainder	low
FCOE	25 %	25 %	low
NOLOSS	19 %	19 %	low
NC	1 %	1 %	low

Note: vQFX does not support the usage of forwarding-class-sets and traffic profiles on xe-interfaces. For this reason, we do not ask you to activate them on the interfaces. Please make sure you understand how to apply them on physical QFX5100 devices interfaces.

Task 6.3: Policing

- Configure vQFX2 so that FTP transfers from Server 1 are policed to 1g and should use the BE forwarding class with low loss-priority.

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-DC lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

Topology

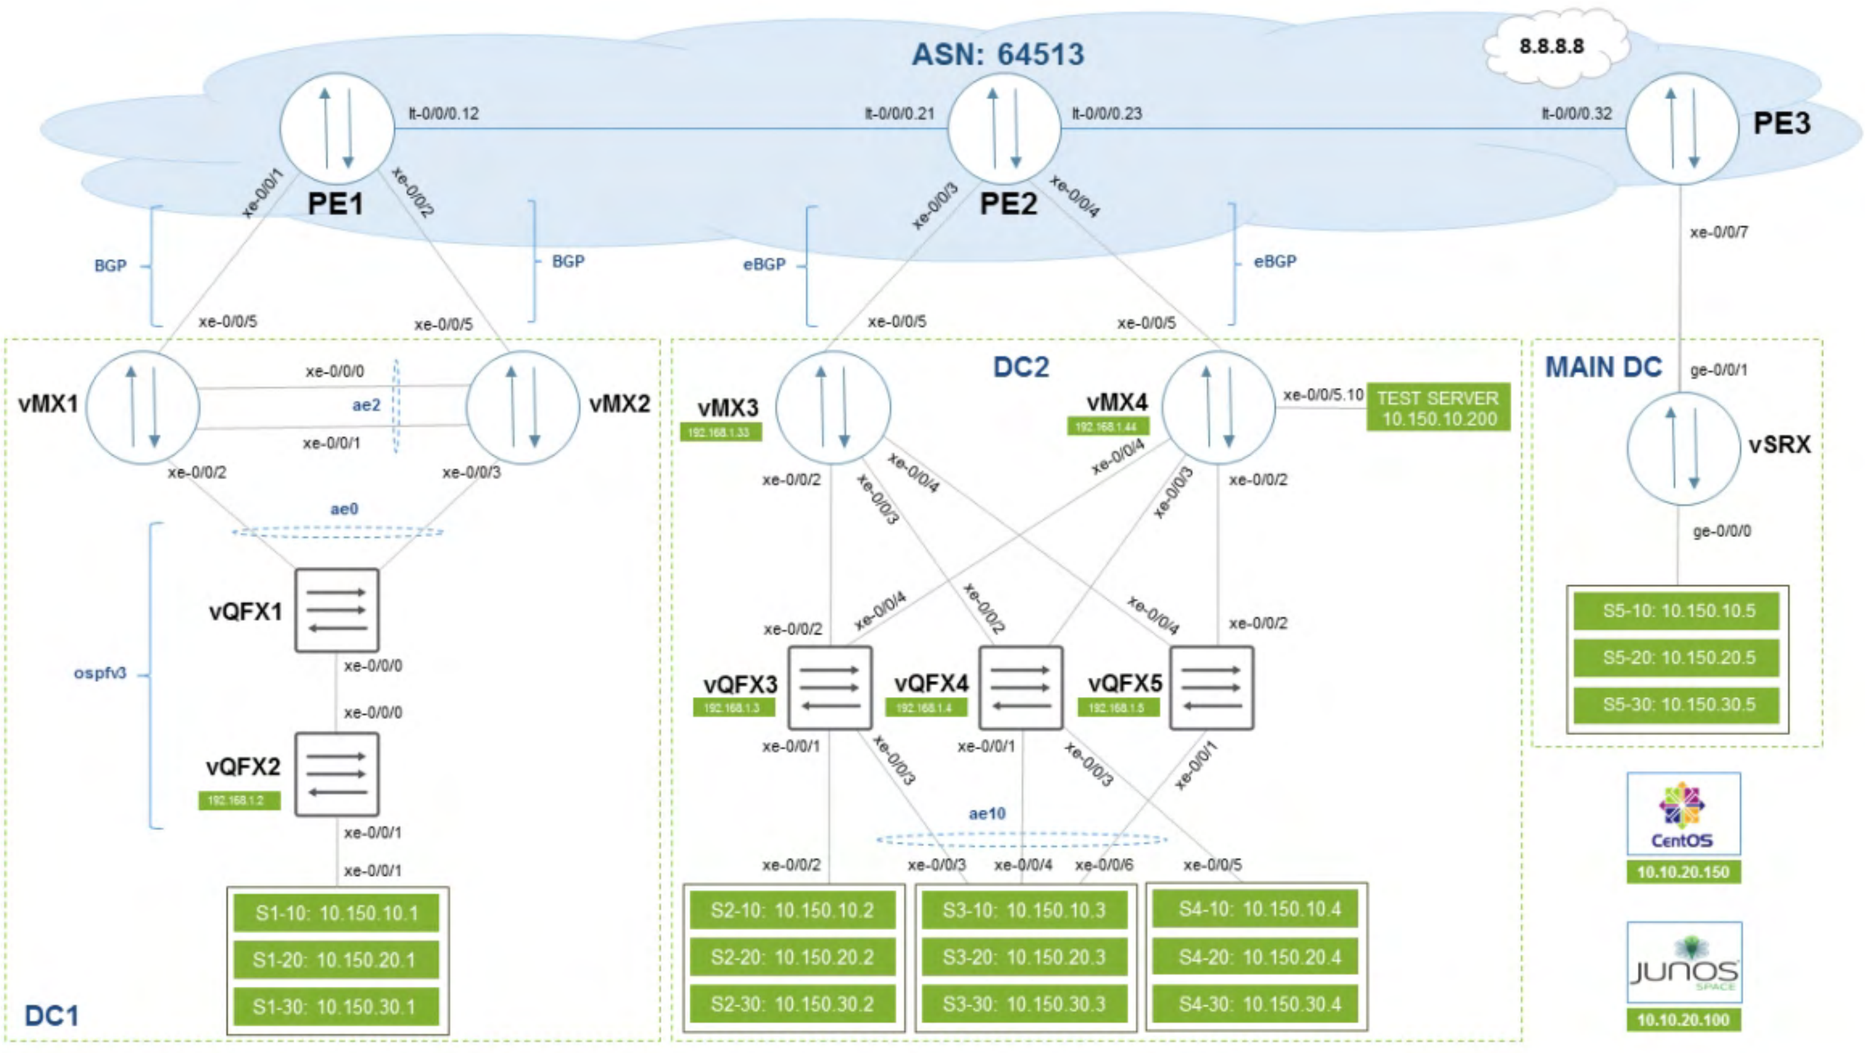

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1: Device Management

Task 1.1: Zero Touch Provisioning (ZTP)

- Ensure that vQFX1 is able to perform ZTP using interface em0.
- The DHCP server that is running on the CentOS host should assign IP address 10.10.20.8 based on the MAC address of vQFX1 em0 interface.
- During ZTP, vQFX1 should automatically download the configuration file /var/pub/ztp/vqfx1 using FTP from the Centos host.
- Edit the */etc/dhcp/dhcpd.conf* file using the "vi" editor.

Task 1.2: Junos Space

- Ensure that all vMX devices in DC1 and DC2 are discovered in Junos Space using SNMP version 1/2c. Use SNMP community "jncie"
- Make sure that the configuration file of vMX1 is stored on the Junos Space server.

Task 1.3: On Box Scripting

- Ensure that interface xe-0/0/5 on vMX1, vMX2, vMX3 and vMX4 is disabled by the Junos OS when it goes down 2 times within 1 minutes. Do not use an external script
- Ensure that event-policies are not allowed to be executed simultaneously.
- Ensure that on-box Python scripts are allowed on all vMX devices in DC1 and DC2
- 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 1.4: NETCONF

- There is a Python script that uses the RFC 4741 compliant NETCONF protocol to manage the configuration of vMX1. Ensure that the Python script can connect to vMX1 over a secured shell session. Allow 6 concurrent connections and a maximum of 2 sessions per minute.

Task 2: Underlay

Task 2.1: Layer 2 Underlay Configuration

Configure the Layer 2 interfaces between the vQFX devices and the servers according to the following requirements:

- Configure interface xe-0/0/1 between S1 and vQFX2. VLAN 10, 20 and 30 should be enabled on this interface
- Configure interface AE0 and xe-0/0/0 on vQFX1 as trunk interfaces for VLAN 10, 20 and 30.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Configure interface AE10 between S3 and vQFX3, vQFX4 and vQFX5. Ensure that vQFX3, vQFX4 and vQFX5 initiate LACP sessions. VLAN 10, 20 and 30 should be enabled on this interface
- Configure interface xe-0/0/3 on vQFX3, interfaces xe-0/0/3 on vQFX4 and interface xe-0/0/1 on vQFX6. VLAN 10, 20, and 30 should be enabled on these interfaces.

Task 2.2: Multi Chassis Link Aggregation (MC-LAG)

Configure an active/active Multi Chassis LAG on MX1 and MX2 according to the following requirements:

- Interface AE0 towards vQFX1 should use VLAN 121. vMX1 should use IP address 10.20.121.1/24, vMX2 should use IP address 10.20.121.2/24 and vQFX2 should use IP address 10.20.121.3/24 on interface xe-0/0/0
- Ensure that multi chassis protection is configured. Do not configure multi chassis protection under the interface hierarchy.
- Configure a redundant ICL link. Use VLAN 12 for the ICCP connection. vMX1 should use IP address 10.21.12.1/30 and vMX2 should use IP address 10.21.12.2/30.
- Configure a redundant default gateway with IP address 10.20.121.254. vMX1 should always be the master, also when recovered from a failure. Use MD5 password “jncie”.

Task 2.3: Layer 2 Underlay Features

- Log a message if there are more than 10 mac-flaps in vlan 121 on vQFX3
- Limit the total number of mac-addresses for VLAN 121 on VQFX1 to 256. Drop packets if the threshold is exceeded.

Task 2.4: BUM Traffic Handling

- Ensure that only Unknown Unicast traffic is limited to 150 kbps on vQFX2
- Do not use the default profile
- Enable this filter for all existing 10G layer 2 interfaces and automatically for 10G layer 2 interfaces that are added in the future.

Task 2.5: Layer 2 Redundancy

- Ensure that interface AE10 towards Server 3 keeps forwarding traffic for VLAN10, 20 and 30 if Server 3 cannot exchange LACP PDUs during a server upgrade.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.6: Layer 3 Underlay MAIN DC

- Configure interface ge-0/0/0 on vSRX for VLAN 50, VLAN 51 and VLAN 52. Use the last possible address within the /24 subnet as the default gateway address for each VLAN

Task 2.7: Layer 3 Underlay Routing DC1

Configure OSPFv2 between vMX1, vMX2, vQFX2 according to the following requirements:

- All interfaces should be configured in the backbone area
- Ensure that OSPFv2 adjacencies between vQFX2, vMX1 and vMX2 are in the “full” state.

Task 2.8: Layer 2 Underlay Switching vQFX1

Ensure the Ethernet switching table of vQF1 has exactly the same output as shown below:

```
Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
Vlan                MAC                 MAC         Age    Logical                NH        RTR
name                address             flags              interface               Index     ID
jncielag            00:00:5e:00:01:fe   D             -    ae0.0                  0         0
jncielag            00:05:01:01:01:01   D             -    ae0.0                  0         0
jncielag            00:05:02:02:02:02   D             -    ae0.0                  0         0
jncielag            00:05:03:03:03:03   D             -    xe-0/0/0.0             0         0
```

Task 2.9: EBGP Underlay

Configure the IP fabric in DC1 and DC2 according to the following requirements:

Device	**BGP AS**
vMX1	64512
vMX2	64512
vMX3	64523
vMX4	64524
vQFX3	64513
vQFX4	64514
vQFX5	64515

Configure BFD on all EBGP sessions for all devices in DC2s IP fabric.

- Ensure that BFD sessions do no support BGP sessions that peer on loopback interfaces
- Enable BFD MD5 authentication with password “jncie”

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Ensure that traffic is load-balanced if there are equal cost paths towards the same destination.

Task 3: Controllerless Overlay

Task 3.1: EVPN Overlay Configuration

Configure an EVPN overlay between DC1, DC2 and MAIN DC according to the following requirements:

- Use BGP AS: 65535
- vQFX2, vQFX3, vQFX4, vQFX5 should only have EVPN overlay sessions with vMX3 and vMX4.
- vMX3 and vMX4 should **not** initiate BGP sessions used for the EVPN overlay. Do not use the "passive" command or a firewall filter to accomplish this task.

Task 3.2: EVPN/VLAN Service Configuration

- Configure an EVPN service called "**vpn1**" for VLAN 10, VLAN 20 and VLAN30.

EVPN/VLAN service	**VLAN**	**vMX3 gateway address**	**vMX4 gateway address**
vpn1	10	10.150.10.100 10.150.10.254	10.150.10.101 10.150.10.254
	20	10.150.20.100 10.150.20.254	10.150.20.101 10.150.20.254
	30	10.150.30.100 10.150.30.254	10.150.30.101 10.150.30.254

Ensure that the following communities are used for each EVPN/VLAN service as shown in the following table:

EVPN/VXLAN service	**Community name**	**Community value**
vpn1	evpn	1:1
	vpn1	1:1000
	vni-10	1:10
	vni-20	1:20
	vni-30	1:30

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.3: EVPN/VXLAN Service Tuning

- Configure vQFX2 to age out the address of the remote VTEP after 320 seconds, if all the MAC addresses learned from that VTEP age out. Your solution must also account for VNIs that might be added in the future.

Task 3.4: EVPN Signaling

- Ensure that vQFX2, vQFX3, vQFX4 and vQFX5 receive the following EVPN type 2 routes:

```
2:192.168.1.33:1::30::00:99:00:00:00:10/304
2:192.168.1.33:1::20::00:99:00:00:00:20/304
2:192.168.1.33:1::10::00:99:00:00:00:30/304
```

Task 4: Data Center Interconnect

The PE devices as shown in the topology are simulated on vMX5 using preconfigured logical systems. The interfaces between the PE devices are also preconfigured.

Logical system name	Interface	IP address	Remote end
PE1	lt-0/0/0.12	172.16.12.1/24	PE2
PE2	lt-0/0/0.21	172.16.12.2/24	PE1
	lt-0/0/0.23	172.16.23.2/24	PE3
PE3	lt-0/0/0.32	172.16.23.3/24	PE2

Task 4.1: DCI Configuration

Configure two L3VPN services using the following information:

- The first L3VPN service should be named “underlay” on vMX5 and is used for the underlay between DC1 and DC2

The following table shows the devices and interface configurations needed for the DCI. Some of the interfaces are already preconfigured.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

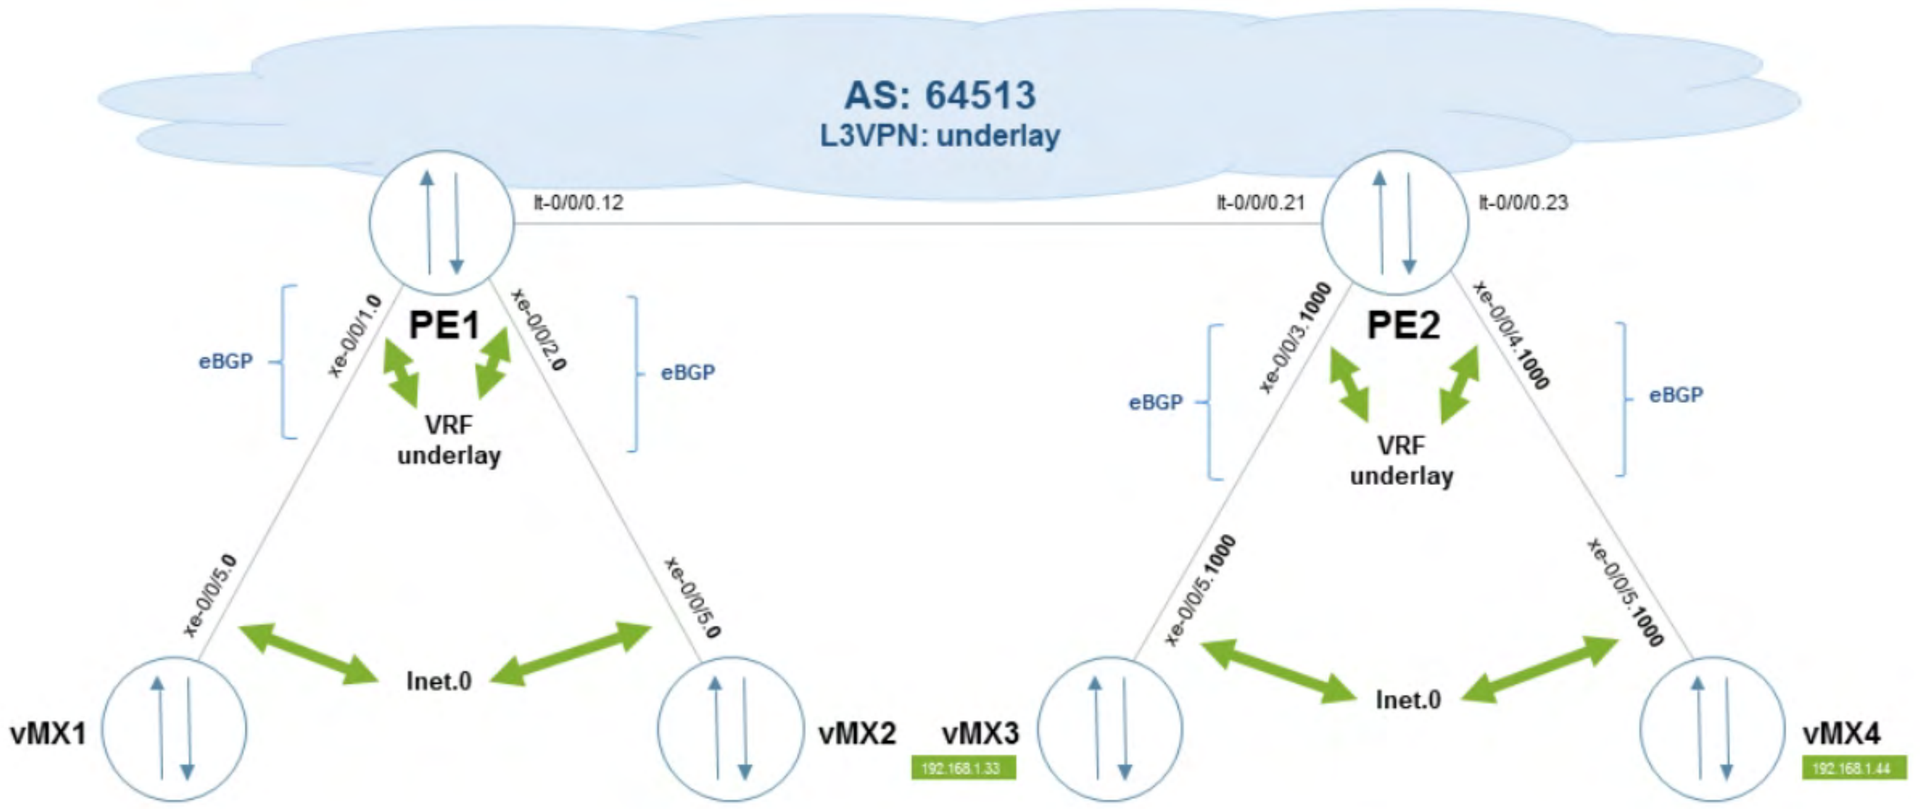

Device	Interfaces VLAN = unit	PE-CE routing	IP address	Additional information
vMX1	xe-0/0/5.0 untagged	EBGP	10.20.1.2/30	Global routing (underlay)
vMX2	xe-0/0/5.0 untagged	EBGP	10.20.2.2/30	Global routing (underlay)
vMX3	xe-0/0/5.1000	EBGP	10.30.3.2/30	Global routing (underlay)
vMX4	xe-0/0/5.1000	EBGP	10.30.4.2/30	Global routing (underlay)
vMX5 (PE1)	xe-0/0/1.0	EBGP	10.20.1.1/30	Routing instance: underlay
	xe-0/0/2.0	EBGP	10.20.2.1/30	Routing instance: underlay
vMX5 (PE2)	xe-0/0/3.1000	EBGP	10.30.3.1/30	Routing instance: underlay
	xe-0/0/4.1000	EBGP	10.30.4.1/30	Routing instance: underlay

- The second L3VPN should be named “dc-vpns” on vMX5 and is used for L3 routing for EVPN service “vpn1” between DC2 and MAIN DC.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

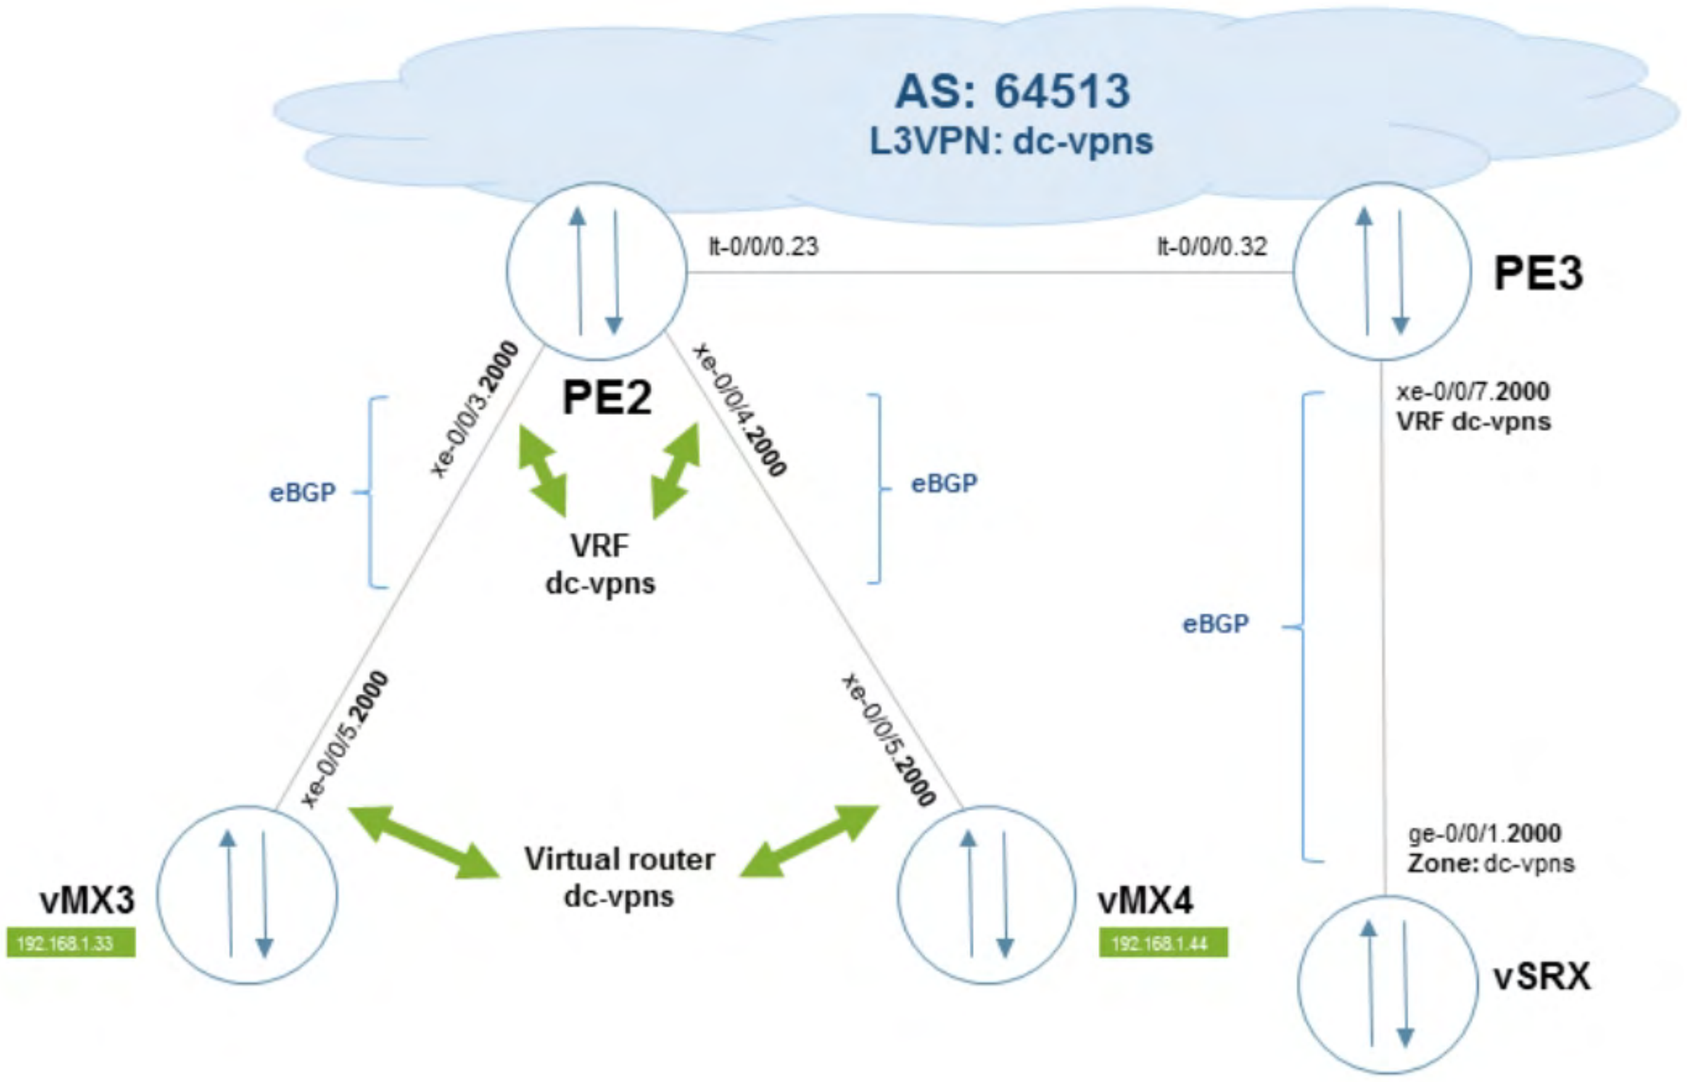

The following table shows the devices and interface configurations needed for the DCI. Some of the interfaces are already preconfigured

Device	Interfaces VLAN = unit	PE-CE routing	IP address	Additional information
vMX3	xe-0/0/5.2000	EBGP	10.30.3.6/30	Routing instance: dc-vpns
vMX4	xe-0/0/5.2000	EBGP	10.30.4.6/30	Routing instance: dc-vpns
vMX5 (PE2)	xe-0/0/3.2000	EBGP	10.30.3.5/30	VRF: dc-vpns
	xe-0/0/4.2000	EBGP	10.30.4.5/30	VRF: dc-vpns
vMX5 (PE3)	xe-0/0/7.2000	EBGP	10.40.2.1/30	VRF: dc-vpns

Task 4.2: DC1 Configuration

vMX1 and vMX2 use BGP AS 64512 to exchange routes with PE1. Ensure that vMX1 and vMX2 always prefers the shortest path towards PE1. Ensure vMX1 and vMX2 redistribute DCI underlay routes into OSPFv2. Ensure that the vQFX2 VTEP interface is learned in DC2.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.3: (Inter) Data Center Connectivity

Ensure that the Servers in DC1, DC2 and MAIN DC have full connectivity between all VLANs. Use the test addresses shown in the topology for verification purposes

Task 4.4: Data Center Connectivity to External Server

Ensure that the Servers DC2 and MAIN DC can ping the external server (8.8.8.8). The external server is preconfigured on vMX5.

- Traffic between the servers in DC2, MAIN DC and the external server should always transit the "dc-vpns" and "untrust" zones of the vSRX firewall.
- The external server should use destination IP address 1.1.1.2 if it sends ICMP echo replies to Servers in DC2 and MAIN DC.

Task 4.5: EVPN VXLAN / EVPN MPLS Interworking

Prepare vMX4 for a future EVPN MPLS DCI connection:

- Create an EVPN MPLS instance called "**evpn**".
- Interface xe-0/0/5.10 should be included in the EVPN MPLS instance.
- Ensure all servers in DC2 are able to ping TEST SERVER.

Task 5: Security

Task 5.1: User Authentication and Authorization

Configure a user called "jncie-1" on vSRX with password "jncie123". Helpful information should be shown when this user logs in to the device. User "jncie-1" should only be allowed to display the device configuration and view system and interfaces statistics.

Configure a user called "jncie-2" on vSRX with password "jncie123" and assign the operator role.

Task 5.2: SRX Configuration

Configure the vSRX according to the following requirements:

- Interface ge-0/0/1.2000 should be in the "dc-vpns" zone
- Interface ge-0/0/1.3000 should be in the "untrust" zone
- Configure a static default route towards the next-hop IP address (1.1.1.1) of AS 64513 in the untrust zone.
- Allow ICMP PING traffic towards vSRX interfaces, except for interfaces in the "untrust" zone

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Allow ICMP traffic between all zones
- Do not configure routing-instances on vSRX
- Only allow traffic if required to complete other tasks in this superlab. Do not allow other traffic

Task 5.3: RE Protection

Configure control-plane protection on vQFX2 with the following requirements:

- Allow incoming telnet and ssh sessions from all destinations.
- Allow BGP sessions to vMX3 and vMX4. Make sure that if new BGP neighbors are configured in the "protocols bgp" stanza, you do not need to change your RE protection configuration.
- Allow ICMP from all destinations
- Allow specific traffic needed by other tasks in this superlab.
- Deny, count and log all other traffic. The control plane should not respond to denied traffic.

Task 6: Class of Service

Task 6.1: Forwarding Classes

Configure the following forwarding classes and forwarding class sets on all vQFX devices

Forwarding Class	Queue ID	No-Loss
BE	0	Disabled
FCOE	3	Enabled
NOLOSS	4	Enabled
NC	7	Disabled

Forwarding Class Set	Forwarding Class
LOSS	BE
	NC
NOLOSS	FCOE
	NOLOSS

Task 6.2: Schedulers and Traffic Profiles

- Configure a scheduler mapping named "LOSS" that includes the BE and NC schedulers on all vQFX devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Configure a scheduler mapping named “NOLOSS” that includes the FCOE and NOLOSS schedulers on all vQFX devices.
- Configure traffic profiles for each scheduler

Scheduler name	Transmit rate	Buffer size	Priority
BE	55 %	remainder	low
FCOE	25 %	25 %	low
NOLOSS	19 %	19 %	low
NC	1 %	1 %	low

Note: vQFX does not support the usage of forwarding-class-sets and traffic profiles on xe- interfaces. For this reason, we do not ask you to activate them on the interfaces. Please make sure you understand how to apply them on physical QFX5100 devices interfaces.

Task 6.3: Policing

Configure vQFX2 so that NFS traffic from server 1 is policed to 2g and should use the NOLOSS forwarding class with low loss-priority

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Appendices

Appendix – Chapter 1: Layer 2 Underlay

Part 1: MC-LAG

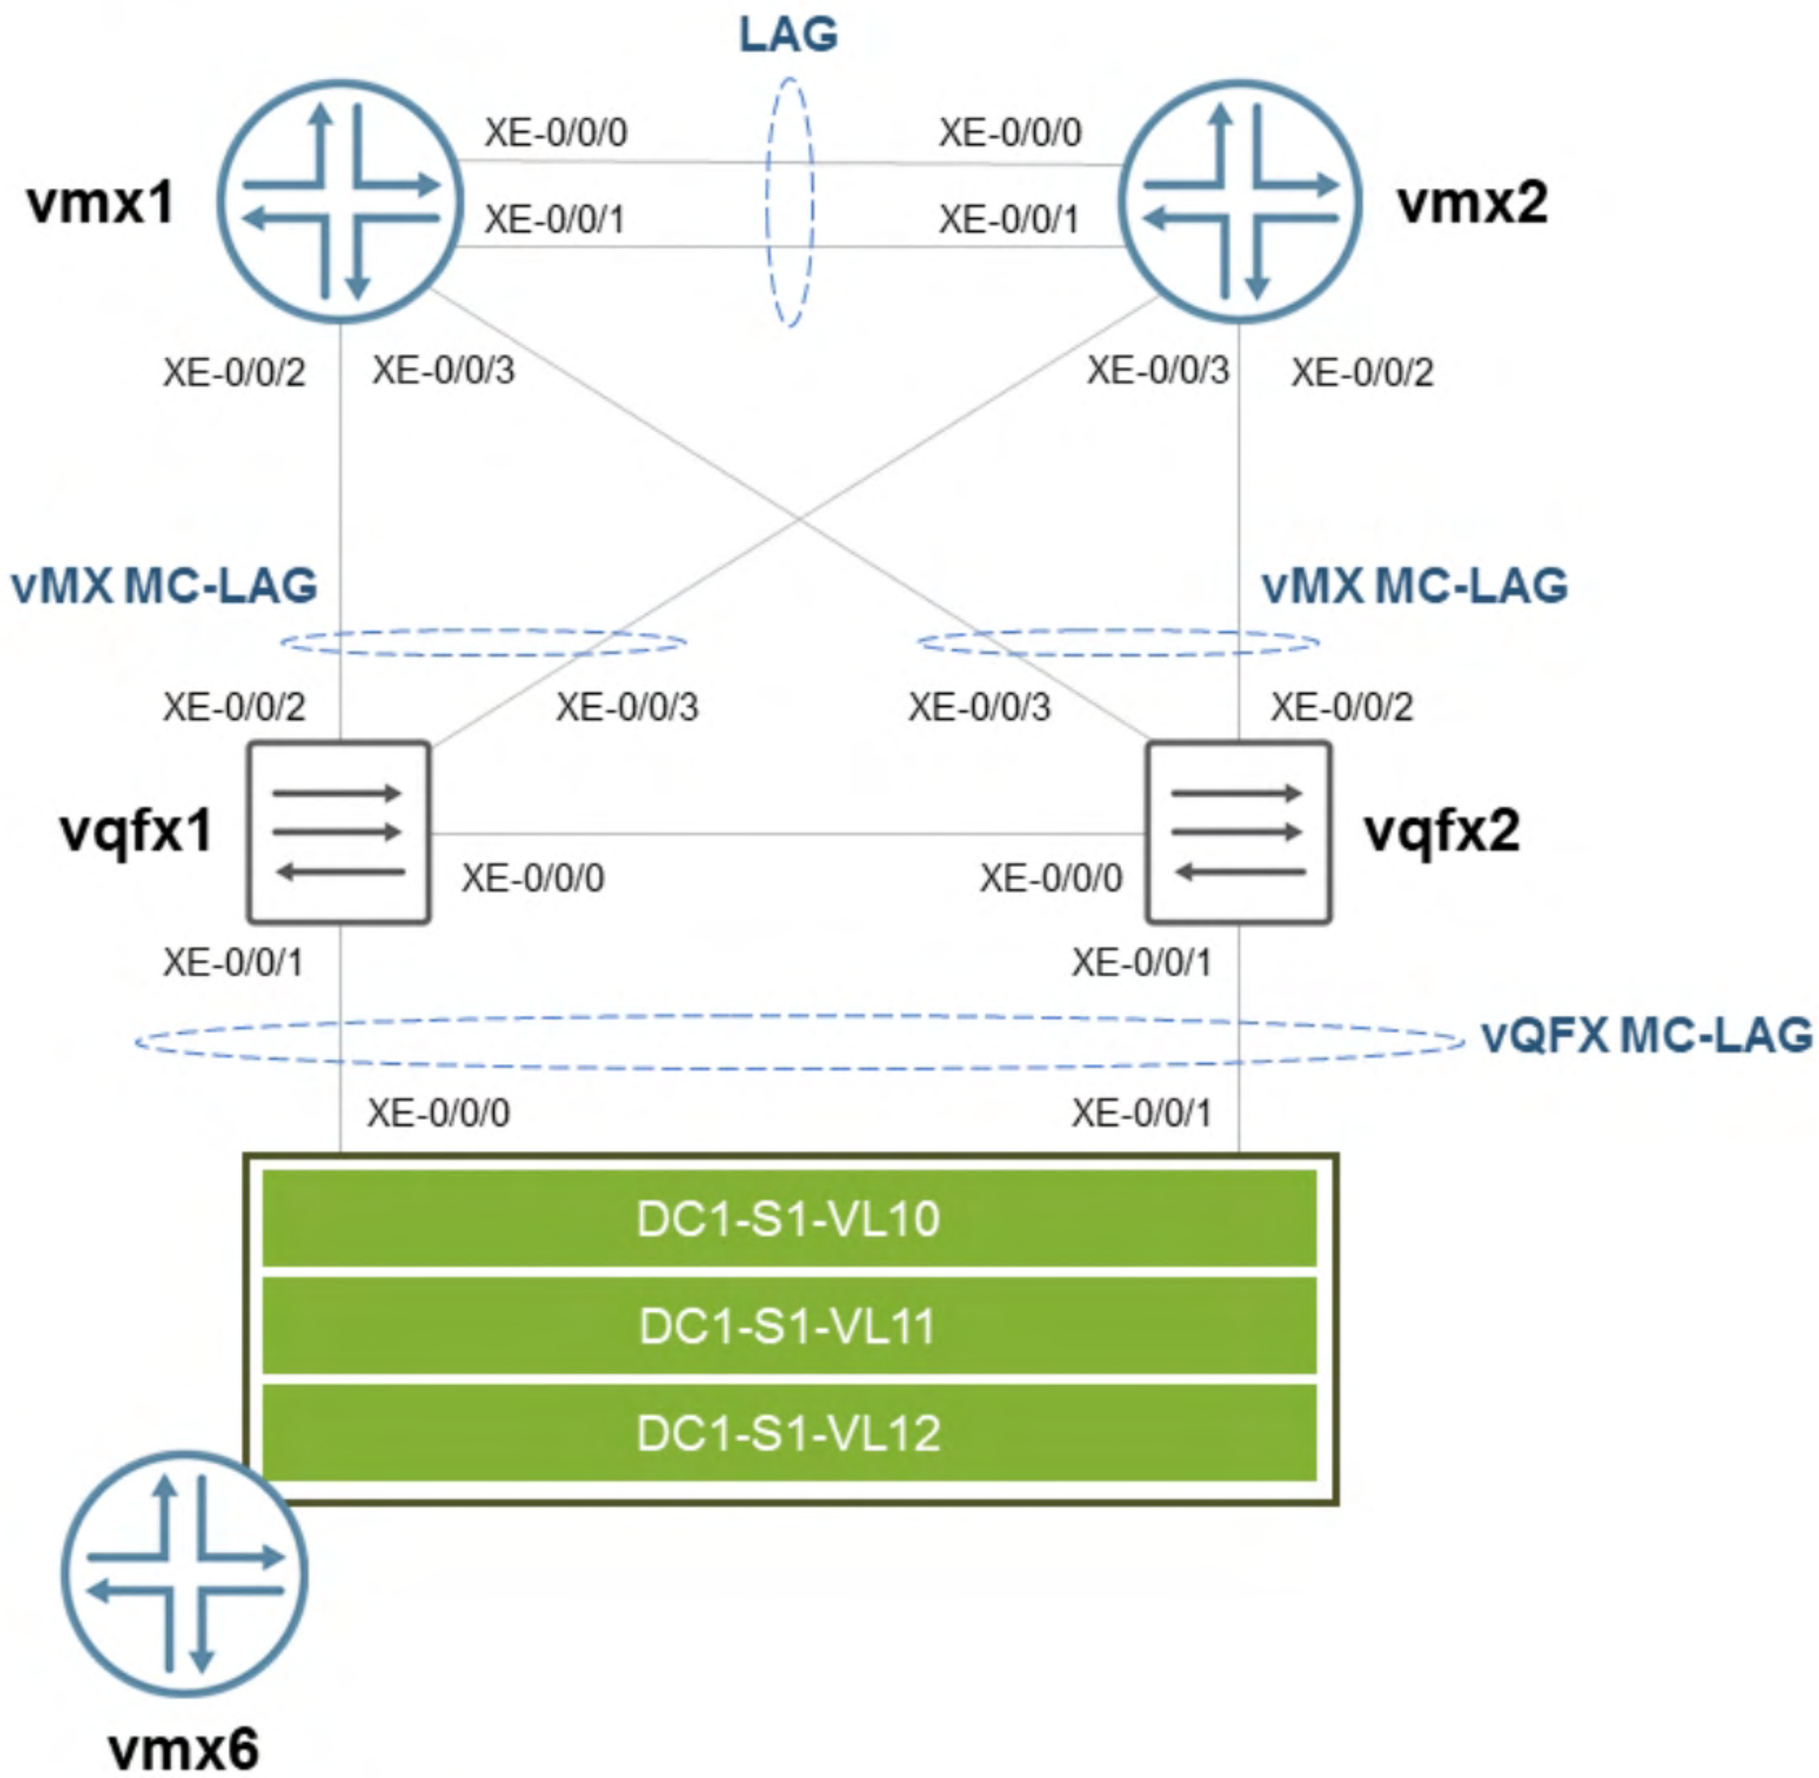

Figure 14 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.1: ToR Configuration

- Configure VLAN 100 on interface xe-0/0/0 on both vQFX1 as well as vQFX2.
- Configure an IRB interface for the VLAN. Use the following IP addresses:

vQFX1	192.168.100.1/30
vQFX2	192.168.100.2/30

- Traffic between the switches should be tagged.

Solution

vQFX1:

The configuration of the IRB interface:

```
set interfaces irb unit 100 family inet address 192.168.100.1/30
```

Ensuring traffic between the two switches is tagged by setting the interface mode to 'trunk':

```
set interfaces xe-0/0/0 unit 0 family ethernet-switching interface-mode trunk
```

Statically assigning the VLAN to the link:

```
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members 100
```

Configuring the VLAN with a VLAN-ID and an IRB interface:

```
set vlans vlan-100 vlan-id 100
set vlans vlan-100 l3-interface irb.100
```

vQFX2:

```
set interfaces irb unit 100 family inet address 192.168.100.2/30

set interfaces xe-0/0/0 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members 100

set vlans vlan-100 vlan-id 100
set vlans vlan-100 l3-interface irb.100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

This is a pretty straightforward assignment. It is still important to properly verify your work. This is because other parts of the infrastructure are going to rely on this link.

We first check the VLAN:

```
jncie@vQFX1> show vlans 100

Routing instance        VLAN name               Tag         Interfaces
default-switch          vlan-100                100         xe-0/0/0.0*
```

Besides the VLAN name and tag, this command tells us that the VLAN is active on interface xe-0/0/0. The asterisk next to the interface is an indication that the interface is UP.

The previous command does not tell us if the link is tagged or whether or not an IRB interface is configured for it. To display that additional information, we issue the following command:

```
jncie@vQFX1> show vlans 100 detail

Routing instance: default-switch
  VLAN Name: vlan-100                       State: Active
Tag: 100
Internal index: 3, Generation Index: 5, Origin: Static
MAC aging time: 300 seconds
Layer 3 interface: irb.100
VXLAN Enabled : No
Interfaces:
    xe-0/0/0.0*,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0
```

To verify what the physical interface is doing:

```
jncie@vQFX1> show ethernet-switching interface xe-0/0/0.0

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,
                         SCTL - shutdown by Storm-control )

Logical          Vlan          TAG     MAC          STP          Logical           Tagging
interface        members               limit        state        interface flags
xe-0/0/0.0                             8192                                        tagged
                 vlan-100      100     1024         Forwarding                     tagged
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Check the IRB interface:

```
jncie@vQFX1> show interfaces irb.100

  Logical interface irb.100 (Index 547) (SNMP ifIndex 528)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: default-switch Bridging Domain: vlan-100
    Input packets : 0
    Output packets: 1
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.100.0/30, Local: 192.168.100.1, Broadcast: 192.168.100.3
```

Verify whether or not we can send traffic across the trunk to the IRB interface on vQFX2:

```
jncie@vQFX1> ping 192.168.100.2 rapid

PING 192.168.100.2 (192.168.100.2): 56 data bytes
!!!!!
--- 192.168.100.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.637/14.363/16.138/0.908 ms
```

Check and see where we learned the ARP entry:

```
jncie@vQFX1> show arp interface irb.100

MAC Address               Name                  Interface               Flags
02:05:86:71:41:00 192.168.100.2   192.168.100.2          irb.100 [xe-0/0/0.0]    none
```

In addition to an ARP entry, we can also see the MAC table was updated with vQFX2's MAC address:

```
jncie@vQFX1> show ethernet-switching table

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 : 1 entries, 1 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical           NH        RTR
    name                address             flags              interface         Index     ID
    vlan-100            02:05:86:71:41:00   D             -    xe-0/0/0.0        0         0
```

In order to find out what MAC a Juniper device is using on the IRB interface, we can issue the following command:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX2> show interfaces irb

Physical interface: irb    , Enabled, Physical link is Up
  Interface index: 642, SNMP ifIndex: 504
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:05:86:71:41:00, Hardware address: 02:05:86:71:41:00
  Last flapped   : Never
    Input packets : 0
    Output packets: 0
```

...

The current address is the address in use. If a MAC address is configured, then that address will be displayed here. The hardware address refers to the MAC address of the hardware.

Task 1.2: MC-LAG Peers: Server-Facing Configuration

- Configure vQFX1 and vQFX2 as MC-LAG peers.
- Use IRB 100 to source the ICCP session. Use 'jncie' as the MD5 authentication key.
- Set the mode to 'active-active' and determine the other configuration parameters yourself.
- Use the OoB network for additional checks between the MC-LAG peers, enabling the network to handle split-brain scenarios.
- During a split brain, vQFX2 should become the inactive MC-LAG peer.
- Enable the MC-LAG interface for VLANs 10, 11 and 12.

Solution

vQFX1:

The switch has to be enabled for LAGs. Since a later task will require one additional LAG, we enable the switch to support two AE interfaces:

```
set chassis aggregated-devices ethernet device-count 2
```

We assign the server-facing interface to the LAG:

```
set interfaces xe-0/0/1 ether-options 802.3ad ae0
```

We configure the LAG to actively send out LACP packets:

```
set interfaces ae0 aggregated-ether-options lacp active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

By default, the switch will use its base MAC address. We need to present the two MC-LAG peers as one device. Therefore, we configure an identical system-ID on both MC-LAG peers. In addition to specifying the system-ID, we must also configure both MC-LAG peers with the same admin-key:

```
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:11:22
set interfaces ae0 aggregated-ether-options lacp admin-key 1
```

The AE 0 interface has to be configured as an MC-LAG interface. We do that in the [interfaces ae0 aggregated-ether-options mc-ae] stanza. First, we specify the same MC-LAG ID on both peers:

```
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
```

We can assign the MC-AE interface to a redundancy-group. Without this configuration statement, the MC-AE interface is assigned the default redundancy group (1):

```
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
```

We need to configure both MC-LAG peers with a different chassis ID for the MC-AE interface:

```
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
```

Next, we specify the MC-AE mode. You must configure this, even though the only possible mode on the QFX is active-active:

```
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```

The last MC-AE parameter is the status-control parameter. This parameter decides what happens during split-brain scenarios. When a split-brain occurs, one device has to remain active and the other must become standby. In this scenario, vQFX2 should become standby. Therefore, vQFX1 is configured as active:

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
```

We need to assign the VLANs to the interface. When we specify the interface mode prior to the VLAN members, we can apply the following configuration in a single line:

```
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members 10-12
```

We start the ICCP configuration with the specification of the source IP and the authentication key:

```
set protocols iccp local-ip-addr 192.168.100.1
set protocols iccp authentication-key jncie
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We need to specify what redundancy groups this peer is part of:

```
set protocols iccp peer 192.168.100.2 redundancy-group-id-list 1
```

Configure the MC-LAG peer to check the vQFX2 OoB management interface for additional checks (on vQFX2 we specify the vQFX1 OoB IP):

```
set protocols iccp peer 192.168.100.2 backup-liveness-detection backup-peer-ip 10.10.20.9
```

Activate BFD for the ICCP session:

```
set protocols iccp peer 192.168.100.2 liveness-detection minimum-interval 3000
```

Activate the vQFX for MC-LAG. Here, we specify the MC-LAG peer IP (vQFX2) as well as the physical interface that vQFX1 uses to reach the MC-LAG peer (vQFX2):

```
set multi-chassis multi-chassis-protection 192.168.100.2 interface xe-0/0/0
```

Finally, for the MC-LAG part of the configuration, we need to specify the same service-id on both switches:

```
set switch-options service-id 1
```

Do not forget to configure the VLANs that are added to the trunk. On EX-switches without the ELS configuration, a commit error would appear when VLANs were assigned to interfaces without being configured. The ELS issues no such warning. When the VLAN configuration is skipped, the VLANs will not handle any traffic:

```
set vlans vlan-10 vlan-id 10
set vlans vlan-11 vlan-id 11
set vlans vlan-12 vlan-id 12
```

vQFX2:

The configuration on the vQFX2 is similar to the one required for vQFX1. The highlighted parts of the configuration emphasize what the differences are from the vQFX1 configuration:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:11:22
set interfaces ae0 aggregated-ether-options lacp admin-key 1

set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby

set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members 10-12

set protocols iccp local-ip-addr 192.168.100.2
set protocols iccp authentication-key jncie
set protocols iccp peer 192.168.100.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.100.1 backup-liveness-detection backup-peer-ip 10.10.20.8

set protocols iccp peer 192.168.100.1 liveness-detection minimum-interval 3000

set multi-chassis multi-chassis-protection 192.168.100.1 interface xe-0/0/0

set switch-options service-id 1

set vlans vlan-10 vlan-id 10
set vlans vlan-11 vlan-id 11
set vlans vlan-12 vlan-id 12
```

Verification

There are quite a few things that require verification for this task. In this example, we'll start our verification at the interface level and then move on to check the MC-AE together with ICCP.

Checking the AE interface:

```
jncie@vQFX1> show interfaces terse | match ae

xe-0/0/1.0              up    up   aenet    --> ae0.0
ae0                     up    up
ae0.0                   up    up   eth-switch
```

This output tells us that AE0 is using 1 physical link (xe-0/0/1) and that the AE0 interface itself is up. The first 'up' tells us the link is administratively enabled and the second 'up' shows us that the link is up.

Additional verification of the AE interface can be done as follows:

```
jncie@vQFX1> show interfaces ae0 extensive

Physical interface: ae0 ) (MC-AE-1, active), Enabled, Physical link is Up
  Interface index: 662, SNMP ifIndex: 537, Generation: 1245
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE ...

    Aggregate member links: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    LACP info:         Role     System              System       Port     Port     Port
                          priority         identifier   priority   number     key

  xe-0/0/1.0      Actor         127   00:00:00:00:11:22       127        1     1
  xe-0/0/1.0    Partner         127   00:05:86:71:25:c0       127        1     5
LACP Statistics:       LACP Rx      LACP Tx   Unknown Rx   Illegal Rx
  xe-0/0/1.0               980          941            0            0
Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
  xe-0/0/1.0                 0            0            0            0
Protocol eth-switch, MTU: 1514, Generation: 793, Route table: 5
  Flags: Trunk-Mode
```

...

This command outputs additional information over the 'terse' option. In relation to the AE interface, the interesting information here is located at the bottom of the output. Here we see the local system identifier used in LACP packets by the Actor (local device), the LACP key in use and the amount of LACP PDUs send and received on the physical interface.

To check what LACP role the devices on either side of the links assume (active or passive) we issue the following command:

```
jncie@vQFX1> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing
```

Here we see the Actor (local device) as well as the partner (remote device) is actively trying to form a LAG using this link.

Another very nice command to consider during troubleshooting is the following:

```
jncie@vQFX1> show lacp statistics interfaces

Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/1                6739        6480            0            0
```

On vQFX2, we check the following:

- the link status of the AE interface
- the status of the physical link used as member of the LAG
- the LACP system ID and the LACP key ID

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX2> show interfaces ae0 extensive

Physical interface: ae0 ) (MC-AE-1, active), Enabled, Physical link is Up
  Interface index: 662, SNMP ifIndex: 542, Generation: 1805
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE ...
    Aggregate member links: 1

        LACP info:        Role     System             System       Port     Port     Port
                             priority        identifier   priority   number     key

      xe-0/0/1.0     Actor         127  00:00:00:00:11:22       127   32769     1
      xe-0/0/1.0   Partner         127  00:05:86:71:25:c0       127       2     5
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/1.0              6394        6123            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      xe-0/0/1.0                 0           0            0            0
    Protocol eth-switch, MTU: 1514, Generation: 1114, Route table: 5
      Flags: Trunk-Mode
```

After this, we check what VLANs are active on this link:

```
jncie@vQFX1> show ethernet-switching interface ae0.0

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,
                         SCTL - shutdown by Storm-control )

Logical          Vlan          TAG     MAC          STP          Logical           Tagging
interface        members               limit        state        interface flags
ae0.0                                  8192                                        tagged
                 vlan-10       10      1024         Forwarding                     tagged
                 vlan-11       11      1024         Forwarding                     tagged
                 vlan-12       12      1024         Forwarding                     tagged
```

```
jncie@vQFX2> show ethernet-switching interface ae0.0

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         MMAS - Mac-move action shutdown,
                         SCTL - shutdown by Storm-control )

Logical          Vlan          TAG     MAC          STP          Logical           Tagging
interface        members               limit        state        interface flags
ae0.0                                  8192                                        tagged
                 vlan-10       10      1024         Forwarding                     tagged
                 vlan-11       11      1024         Forwarding                     tagged
                 vlan-12       12      1024         Forwarding                     tagged
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

After verifying the normal LAG operations, we move over to the MC-LAG verification. The primary verification commands at our disposal are:

- **show iccp** - used to check status information about the ICCP protocol
- **show interface mc-ae** - used to check status information about the MC-AE interface

In addition to these commands, we can also use the 'show ether-switching table' command to verify whether or not MAC addresses are being learned on both of the MC-LAG peers.

We start out verifying the ICCP protocol on vQFX1:

```
jncie@vQFX1> show iccp

Redundancy Group Information for peer 192.168.100.2
  TCP Connection       : Established
  Liveliness Detection : Up

Backup liveness peer status: Up
  Redundancy Group ID          Status
    1                           Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1
```

The output basically tells us everything we need to know. vQFX1 has an ICCP session with 192.168.100.2 that is established. The 'Liveness detection' tells us that the BFD session was established. The 'Backup liveness' status informs us over the fact that vQFX1 can 'see' vQFX2 over an alternative path (in this case the OoB network).

Note that we can also verify BFD through the use of the following command:

```
jncie@vQFX1> show bfd session

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.100.2            Up                       9.000     3.000         3
```

Here, we see a BFD session formed with the 192.168.100.2 address, which is the configured ICCP peer. The BFD status is up and the detect time is 9 seconds.

Next, we check the MC-AE interface status information:

```
jncie@vQFX1> show interfaces mc-ae extensive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Member Link                     : ae0
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : up
    Logical Interface           : ae0.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.100.2 xe-0/0/0.0 up

MCAE Configuration
    Redundancy Group            : 1
    MCAE ID                     : 1
    MCAE Mode                   : active_active
    Status Control              : active
    Chassis ID                  : 0
LACP Configuration
    System ID                   : 00:00:00:00:11:22
    Admin Key                   : 1
```

We check the command output for the 'mcae active state', indicating that the MC-AE is active and working.

The 'Local Status' and the Peer Status' tell us whether or not the MC-LAG peers are able to function as a part of this MC-LAG. When everything is operational, the status is 'active'.

The 'Local State' and Peer State' indicate whether or not the physical link on the device is operational. For instance, if all configuration parameters on both vQFX1 and vQFX2 would have been in order, but the link between vQFX2 and the server would not have been operational, the output would have looked like this:

```
jncie@vQFX1> show interfaces mc-ae

Member Link                     : ae0
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : down
    Logical Interface           : ae0.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.100.2 xe-0/0/0.0 up
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX2> show interfaces mc-ae

 Member Link                  : ae0
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : down
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.100.1 xe-0/0/0.0 up
```

If you see that the peer status is 'Unknown', check the ICCP configuration and status on both devices.

Back to the lower part of the 'show interfaces mc-ae extensive' output:

```
...
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.100.2 xe-0/0/0.0 up
...
 MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 1
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 0
...
 LACP Configuration
     System ID                : 00:00:00:00:11:22
     Admin Key                : 1
```

Here, we see three different parts of the MC-LAG operations. In the first part, we see the status of the multi-chassis protection link. The second part tells us about the MC-AE configuration parameters and the third part displays the LACP configuration parameters.

Additionally, we can (and should) examine the MAC table on both devices:

```
jncie@vQFX1> show ethernet-switching table

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)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical              NH        RTR
    name                address             flags              interface             Index     ID
    vlan-10             00:05:86:71:25:f0   DL            -    ae0.0                0         0
    vlan-11             00:05:86:71:25:f0   DL            -    ae0.0                0         0
    vlan-12             00:05:86:71:25:f0   DL            -    ae0.0                0         0

jncie@vQFX2> show ethernet-switching table

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 : 3 entries, 3 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical              NH        RTR
    name                address             flags              interface             Index     ID
    vlan-10             00:05:86:71:25:f0   DR            -    ae0.0                0         0
    vlan-11             00:05:86:71:25:f0   DR            -    ae0.0                0         0
    vlan-12             00:05:86:71:25:f0   DR            -    ae0.0                0         0
```

When we examine the MAC table on the devices, we not only notice the MAC-addresses, we also notice the MAC flags: 'DL' and 'DR'. The D indicates that the MAC address was learned dynamically. L is an indication that the MAC address was learned on the local interface and R is an indication that MAC was learned remotely.

Task 1.3: MC-LAG Peers: Core-Facing Configuration

- Configure an AE uplink towards the vMX routers.
- Use interfaces xe-0/0/2 and xe-0/0/3 on both vQFX devices.
- Enable the uplink for the server VLANs.
- Make sure that communications inside the server VLANs can tolerate the loss of an uplink.

Solution

To ensure that the communication inside the server VLANs can tolerate the loss of a single uplink, we configure an AE interface on both QFX devices.

vQFX1:

```
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae1 description core-facing
set interfaces ae1 aggregated-ether-options lacp active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 10-12
```

vQFX2:

```
set interfaces xe-0/0/2 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae1 description core-facing
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 10-12
```

Verification

To verify the configuration, we issue the '`show interfaces terse | match ae`' command on all devices involved.

Task 1.4: vMX LAG Configuration

- Configure VLAN 101 on both vMX1 as well as vMX2.
- Configure an IRB interface for the VLAN. Use the following IP addresses:

vMX1	192.168.101.1/30
vMX2	192.168.101.2/30

- Traffic between the routers should be tagged.
- The links between the vMX routers should be bundled into a LAG.

Solution

Everything in this task has been done before, the only difference here is that we now have to make things work on a pair of vMX routers. Configuration wise, we encounter only two minor differences between vMX and vQFX in this task:

- we assign physical interfaces to the LAG using 'gigether-options' and not 'ether-options'.
- we configure the trunk between the vMX routers with the family bridge and the 'vlan-id-list' instead of using family ethernet-switching and 'vlan members'.

When working on MX and QFX, do not lose yourself in semantics. The syntaxes may differ, the concepts of what we aim to achieve does not.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX1:

```
set chassis aggregated-devices ethernet device-count 3

set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0

set interfaces irb.101 family inet address 192.168.101.1/30

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 101

set bridge-domains bd-101 vlan-id 101 routing-interface irb.101
```

vMX2:

```
set chassis aggregated-devices ethernet device-count 3

set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0

set interfaces irb.101 family inet address 192.168.101.2/30

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 101

set bridge-domains bd-101 vlan-id 101 routing-interface irb.101
```

Verification

The verification commands and output on vMX hardly differ from those on vQFX. First, we verify the LAG. Then we verify the bridge-domain, IRB interface and the trunk configuration. We end our verification generating some traffic between the two vMX routers.

Verifying the AE membership, mode, state and statistics is exactly the same as on the vQFX:

```
jncie@vMX1> show lacp interfaces ae0

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/0                Current   Fast periodic Collecting distributing
      xe-0/0/1                Current   Fast periodic Collecting distributing

jncie@vMX1> show lacp statistics interfaces

Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/0                1246        1289            0            0
      xe-0/0/1                1246        1284            0            0
```

As an alternative, we can also check the AE interface by issuing the following command:

```
jncie@vMX1> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 530, Generation: 159
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:72:fb:c6, Hardware address: 00:05:86:72:fb:c6
  Last flapped   : 2017-04-14 11:47:23 UTC (00:21:50 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                 0 bps
   Output bytes  :               354380              2128 bps
   Input  packets:                    0                 0 pps
   Output packets:                 2726                 2 pps
   IPv6 transit statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  3                                   0                    0                    0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                                   0                    0                    0
  1                                   0                    0                    0
  2                                   0                    0                    0
  3                                   0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
```

```
Logical interface ae0.0 (Index 321) (SNMP ifIndex 531) (Generation 175)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Statistics        Packets        pps         Bytes           bps
  Bundle:
      Input :             0          0             0             0
      Output:             0          0             0             0
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    xe-0/0/0.0
      Input :             0          0             0             0
      Output:             0          0             0             0
    xe-0/0/1.0
      Input :             0          0             0             0
      Output:             0          0             0             0

  Aggregate member links: 2

  LACP info:        Role     System             System       Port     Port    Port
                             priority         identifier   priority   number     key
    xe-0/0/0.0     Actor         127  00:05:86:72:fb:c0        127        1        1
    xe-0/0/0.0   Partner         127  00:05:86:94:9b:c0        127        1        1
    xe-0/0/1.0     Actor         127  00:05:86:72:fb:c0        127        2        1
    xe-0/0/1.0   Partner         127  00:05:86:94:9b:c0        127        2        1
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    xe-0/0/0.0                 0           0            0            0
    xe-0/0/1.0                 0           0            0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    xe-0/0/0.0                 0           0            0            0
    xe-0/0/1.0                 0           0            0            0
  Protocol bridge, MTU: 1514, Generation: 201, Route table: 7, Mesh Group: __all_ces__
    Flags: Is-Primary, Trunk-Mode
```

The commands to analyze the bridge-domain, the MAC table, the IRB interface and the trunk link are as follows:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show bridge domain detail

Routing instance: default-switch
Bridge domain: bd-101                                State: Active
Bridge VLAN ID: 101
Interfaces:
    ae0.0
Total MAC count: 0

jncie@vMX1> show bridge mac-table

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-101, VLAN : 101
   MAC                 MAC      Logical          NH     RTR
   addresssss          flags    interface        Index  ID
   00:05:86:94:9b:f0   D        ae0.0

jncie@vMX1> show interfaces irb

Physical interface: irb, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 504
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:05:86:72:fb:f0, Hardware address: 00:05:86:72:fb:f0
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.101 (Index 327) (SNMP ifIndex 533)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: default-switch Bridging Domain: bd-101
    Input packets : 125
    Output packets: 126
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH
drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.101.0/30, Local: 192.168.101.1, Broadcast: 192.168.101.3
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show interfaces ae0.0

  Logical interface ae0.0 (Index 321) (SNMP ifIndex 531)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes           bps
    Bundle:
        Input :             0          0             0             0
        Output:           298          1         24994           880
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol bridge, MTU: 1514
      Flags: Is-Primary, Trunk-Mode
```

To conclude the verification for this task, we issue a ping between the two vMX routers:

```
jncie@vMX1> ping 192.168.101.2 rapid

PING 192.168.101.2 (192.168.101.2): 56 data bytes
!!!!!
--- 192.168.101.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.035/3.589/4.460/0.481 ms
```

Task 1.5: vMX MC-LAG Configuration

- Configure MC-LAGs on the vMX routers.
- The MC-LAGs should connect to the LAGs that were configured in task 1.3.
- Use the OoB network for additional checks between the MC-LAG peers.
- During a split brain, vMX2 should become the inactive MC-LAG peer.
- Enable the server VLANs on the MC-AE interface.

Solution

vMX1:

We start out configuring the first MC-AE interface:

```
set interfaces xe-0/0/2 gigether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:11:00:00:11
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae1 unit 0 family bridge interface-mode trunk vlan-id-list 10-12
```

After this, we configure a second MC-AE interface:

```
set interfaces xe-0/0/3 gigether-options 802.3ad ae2

set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:00:11:00:00:22
set interfaces ae2 aggregated-ether-options lacp admin-key 1
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 10-12
```

The vMX routers van sustain multiple MC-AEs using the same ICCP session, so we only need to configure a single neighbor relationship:

```
set protocols iccp local-ip-addr 192.168.101.1
set protocols iccp peer 192.168.101.2 redundancy-group-id-list 1
set protocols iccp peer 192.168.101.2 backup-liveness-detection backup-peer-ip 10.10.20.2
set protocols iccp peer 192.168.101.2 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.101.2 authentication-key jncie
```

Multi-chassis configuration:

```
set multi-chassis multi-chassis-protection 192.168.101.2 interface ae0
```

Switch-options configuration:

```
set switch-options service-id 1
```

Bride domains:

```
set bridge-domains bd-10 vlan-id 10
set bridge-domains bd-11 vlan-id 11
set bridge-domains bd-12 vlan-id 12
```

vMX2:

```
set interfaces xe-0/0/3 gigether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:11:00:00:11
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 unit 0 family bridge interface-mode trunk vlan-id-list 10-12


set interfaces xe-0/0/2 gigether-options 802.3ad ae2

set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:00:11:00:00:22
set interfaces ae2 aggregated-ether-options lacp admin-key 1
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 10-12

set protocols iccp local-ip-addr 192.168.101.2
set protocols iccp peer 192.168.101.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.101.1 backup-liveness-detection backup-peer-ip 10.10.20.1
set protocols iccp peer 192.168.101.1 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.101.1 authentication-key jncie

set multi-chassis multi-chassis-protection 192.168.101.1 interface ae0

set switch-options service-id 1

set bridge-domains bd-10 vlan-id 10
set bridge-domains bd-11 vlan-id 11
set bridge-domains bd-12 vlan-id 12
```

Verification

Since the verification on the vMX MC-LAG peers is practically the same as on the vQFX MC-LAG peers, we do will not cover the verification in detail here.

Checking the ICCP session, together with BFD and the additional backup liveness detection via the OoB:

```
jncie@vMX1> show iccp

Redundancy Group Information for peer 192.168.101.2
  TCP Connection       : Established
  Liveliness Detection : Up

Backup liveness peer status: Up
  Redundancy Group ID          Status
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    1                                Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1

Checking the two MC-AE interfaces:

jncie@vMX1> show interfaces mc-ae

 Member Link                    : ae1
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : up
    Logical Interface           : ae1.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.101.2 ae0.0 up

Member Link                     : ae2
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : up
    Logical Interface           : ae2.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.101.2 ae0.0 up
```

Adding extensive to the command reveals the LACP system ID:

```
jncie@vMX1> show interfaces mc-ae extensive

 Member Link                    : ae1
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : up
    Logical Interface           : ae1.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.101.2 ae0.0 up
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
MCAE Configuration
    Redundancy Group          : 1
    MCAE ID                   : 1
    MCAE Mode                 : active_active
    Status Control            : active
    Chassis ID                : 0
LACP Configuration
    System ID                 : 00:00:11:00:00:11
    Admin Key                 : 1

Member Link                   : ae2
Current State Machine's State: mcae active state
Local Status                  : active
Local State                   : up
Peer Status                   : active
Peer State                    : up
    Logical Interface         : ae2.0
    Topology Type             : bridge
    Local State               : up
    Peer State                : up
    Peer Ip/MCP/State         : 192.168.101.2 ae0.0 up

MCAE Configuration
    Redundancy Group          : 1
    MCAE ID                   : 2
    MCAE Mode                 : active_active
    Status Control            : active
    Chassis ID                : 0
LACP Configuration
    System ID                 : 00:00:11:00:00:22
    Admin Key                 : 1
```

We issue the same commands on vMX2. After this, we can verify the MC-LAG from the vQFX point of view.

First, we check LACP:

```
jncie@vQFX1> show lacp interfaces ae1

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/3                  Current   Fast periodic Collecting distributing
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Then we check the AE interface and use the keyword extensive. This way, we see the link status, LACP parameters and LACP statistics:

```
jncie@vQFX1> show interfaces ae1 extensive

Physical interface: ae1    , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 524, Generation: 132
  Description: core-facing
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, ...

  Logical interface ae1.0 (Index 551) (SNMP ifIndex 525) (Generation 136)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :           495          2         33628         1624
      Output:           425          1         54208         2040
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    xe-0/0/2.0
      Input :             0          0             0            0
      Output:             0          0             0            0
    xe-0/0/3.0
      Input :             0          0             0            0
      Output:             0          0             0            0

  Aggregate member links: 2

      LACP info:        Role     System             System       Port     Port     Port
                                    priority       identifier   priority   number     key
    xe-0/0/2.0     Actor         127  02:05:86:71:91:00         127        2      2
    xe-0/0/2.0   Partner         127  00:00:11:00:00:11         127        3      1
    xe-0/0/3.0     Actor         127  02:05:86:71:91:00         127        3      2
    xe-0/0/3.0   Partner         127  00:00:11:00:00:11         127    32771      1
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    xe-0/0/2.0               219         205            0            0
    xe-0/0/3.0               231         207            0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    xe-0/0/2.0                 0           0            0            0
    xe-0/0/3.0                 0           0            0            0
  Protocol eth-switch, MTU: 1514, Generation: 157, Route table: 5
    Flags: Trunk-Mode
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX2> show lacp interfaces ae1

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/3                  Current   Fast periodic Collecting distributing

jncie@vQFX2> show interfaces ae1 extensive

Physical interface: ae1    , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 520, Generation: 132
  Description: core-facing
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, ...
  Logical interface ae1.0 (Index 548) (SNMP ifIndex 523) (Generation 133)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           347          1         23560         1080
        Output:           291          1         37248         2032
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      xe-0/0/2.0
        Input :              0          0             0            0
        Output:              0          0             0            0
      xe-0/0/3.0
        Input :              0          0             0            0
        Output:              0          0             0            0

  Aggregate member links: 2

    LACP info:         Role     System             System       Port     Port     Port
                          priority         identifier   priority   number      key

      xe-0/0/2.0       Actor        127  02:05:86:71:ac:00        127        2      2
      xe-0/0/2.0     Partner        127  00:00:11:00:00:22        127    32772      1
      xe-0/0/3.0       Actor        127  02:05:86:71:ac:00        127        3      2
      xe-0/0/3.0     Partner        127  00:00:11:00:00:22        127        4      1
  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/2.0             138         123            0            0
      xe-0/0/3.0             135         122            0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  xe-0/0/2.0                0          0          0          0
  xe-0/0/3.0                0          0          0          0
Protocol eth-switch, MTU: 1514, Generation: 151, Route table: 5
  Flags: Trunk-Mode
```

Task 1.6: vMX MC-LAG Gateway Configuration

- Configure the following gateways on the vMX routers:

IRB / VLAN	Gateway address
IRB10	10.0.10.254/24
IRB11	10.0.11.254/24
IRB12	10.0.12.254/24

- Use VRRP for VLAN 10 and VLAN 11.
- The vMX1 router should be VRRP master for VLAN 10 and the vMX2 router should be the master for VLAN 11.
- Make sure that you can ping the virtual-router.
- Authenticate the messages using password 'jncie'.
- Use an anycast IRB for VLAN 12.

Solution

vMX1:

We add the server VLANs to the ICL. This way, we can tolerate the loss of a downlink:

```
set interfaces ae0 unit 0 family bridge vlan-id-list 10-12
```

We configure VRRP for VLAN 10 and make sure that vMX1 is the master by setting the priority to 150. We also configure the VRRP group with the 'accept-data' keyword. This enables the master router to accept packets destined to the virtual router. After specifying the authentication type and key, we associate the irb.10 interface with bridge-domain 'bd-10':

```
set interfaces irb unit 10 family inet address 10.0.10.253/24 vrrp-group 10 virtual-address
10.0.10.254
set interfaces irb unit 10 family inet address 10.0.10.253/24 vrrp-group 10 priority 150
set interfaces irb unit 10 family inet address 10.0.10.253/24 vrrp-group 10 accept-data
set interfaces irb unit 10 family inet address 10.0.10.253/24 vrrp-group 10 authentication-
type md5
set interfaces irb unit 10 family inet address 10.0.10.253/24 vrrp-group 10 authentication-key
jncie

set bridge-domains bd-10 routing-interface irb.10
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The default priority for VRRP is 100. For the VRRP group in VLAN 11, we specify no priority and configure the vMX2 router with a priority of 150 later on:

```
set interfaces irb unit 11 family inet address 10.0.11.253/24 vrrp-group 11 virtual-address
10.0.11.254
set interfaces irb unit 11 family inet address 10.0.11.253/24 vrrp-group 11 accept-data
set interfaces irb unit 11 family inet address 10.0.11.253/24 vrrp-group 11 authentication-
type md5
set interfaces irb unit 11 family inet address 10.0.11.253/24 vrrp-group 11 authentication-key
jncie

set bridge-domains bd-11 routing-interface irb.11
```

For VLAN 12, we configure an anycast-gateway. We configure both vMX routers with the same MAC and IP address:

```
set interfaces irb unit 12 family inet address 10.0.12.254/24
set interfaces irb unit 12 mac aa:aa:aa:aa:aa:aa

set bridge-domains bd-12 routing-interface irb.12
```

vMX2:

```
set interfaces ae0 unit 0 family bridge vlan-id-list 10-12

set interfaces irb unit 10 family inet address 10.0.10.252/24 vrrp-group 10 virtual-address
10.0.10.254
set interfaces irb unit 10 family inet address 10.0.10.252/24 vrrp-group 10 accept-data
set interfaces irb unit 10 family inet address 10.0.10.252/24 vrrp-group 10 authentication-
type md5
set interfaces irb unit 10 family inet address 10.0.10.252/24 vrrp-group 10 authentication-key
jncie

set bridge-domains bd-10 routing-interface irb.10

set interfaces irb unit 11 family inet address 10.0.11.252/24 vrrp-group 11 virtual-address
10.0.11.254
set interfaces irb unit 11 family inet address 10.0.11.252/24 vrrp-group 11 priority 150
set interfaces irb unit 11 family inet address 10.0.11.252/24 vrrp-group 11 accept-data
set interfaces irb unit 11 family inet address 10.0.11.252/24 vrrp-group 11 authentication-
type md5
set interfaces irb unit 11 family inet address 10.0.11.252/24 vrrp-group 11 authentication-key
jncie

set bridge-domains bd-11 routing-interface irb.11

set interfaces irb unit 12 family inet address 10.0.12.254/24
set interfaces irb unit 12 mac aa:aa:aa:aa:aa:aa

set bridge-domains bd-12 routing-interface irb.12
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

We start out verifying the VRRP mastership on vMX1 and vMX2:

```
jncie@vMX1> show vrrp

Interface     State       Group   VR state VR Mode    Timer    Type   Address
irb.10        up             10   master   Active       A  0.724 lcl    10.0.10.253
                                                                vip    10.0.10.254
irb.11        up             11   backup   Active       D  3.419 lcl    10.0.11.253
                                                                vip    10.0.11.254
                                                                mas    10.0.11.252
```

```
jncie@vMX2> show vrrp

Interface     State       Group   VR state VR Mode    Timer    Type   Address
irb.10        up             10   backup   Active       D  3.293 lcl    10.0.10.252
                                                                vip    10.0.10.254
                                                                mas    10.0.10.253
irb.11        up             11   master   Active       A  0.144 lcl    10.0.11.252
                                                                vip    10.0.11.254
```

Verifying the use of MD5 authentication and checking whether or not the vMX routers will respond when an ICMP is sent to the VIP address:

```
jncie@vMX1> show vrrp detail

Physical interface: irb, Unit: 10, Address: 10.0.10.253/24
  Index: 339, SNMP ifIndex: 538, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 10, State: master, VRRP Mode: Active
  Priority: 150, Advertisement interval: 1, Authentication type: md5
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 10.0.10.254
  Advertisement Timer: 0.376s, Master router: 10.0.10.253
  Virtual router uptime: 02:27:04, Master router uptime: 00:03:26
  Virtual Mac: 00:00:5e:00:01:0a
  Preferred: yes
  Tracking: disabled

Physical interface: irb, Unit: 11, Address: 10.0.11.253/24
  Index: 341, SNMP ifIndex: 551, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 11, State: backup, VRRP Mode: Active
  Priority: 100, Advertisement interval: 1, Authentication type: md5
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 10.0.11.254
  Dead timer: 3.291s, Master priority: 150, Master router: 10.0.11.252
  Virtual router uptime: 02:27:04
  Preferred: yes
  Tracking: disabled
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We conclude our verification by running several pings from vMX6:

```
jncie@vMX6> ping routing-instance DC1-S1-VL10 10.0.10.254 rapid

PING 10.0.10.254 (10.0.10.254): 56 data bytes
!!!!!
--- 10.0.10.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 9.979/15.044/23.221/5.859 ms

jncie@vMX6> ping routing-instance DC1-S1-VL11 10.0.11.254 rapid
PING 10.0.11.254 (10.0.11.254): 56 data bytes
!!!!!
--- 10.0.11.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.867/37.473/94.985/30.787 ms

jncie@vMX6> ping routing-instance DC1-S1-VL12 10.0.12.254 rapid

PING 10.0.12.254 (10.0.12.254): 56 data bytes
!!!!!
--- 10.0.12.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.230/8.620/11.135/1.320 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: Multistage MC-LAG

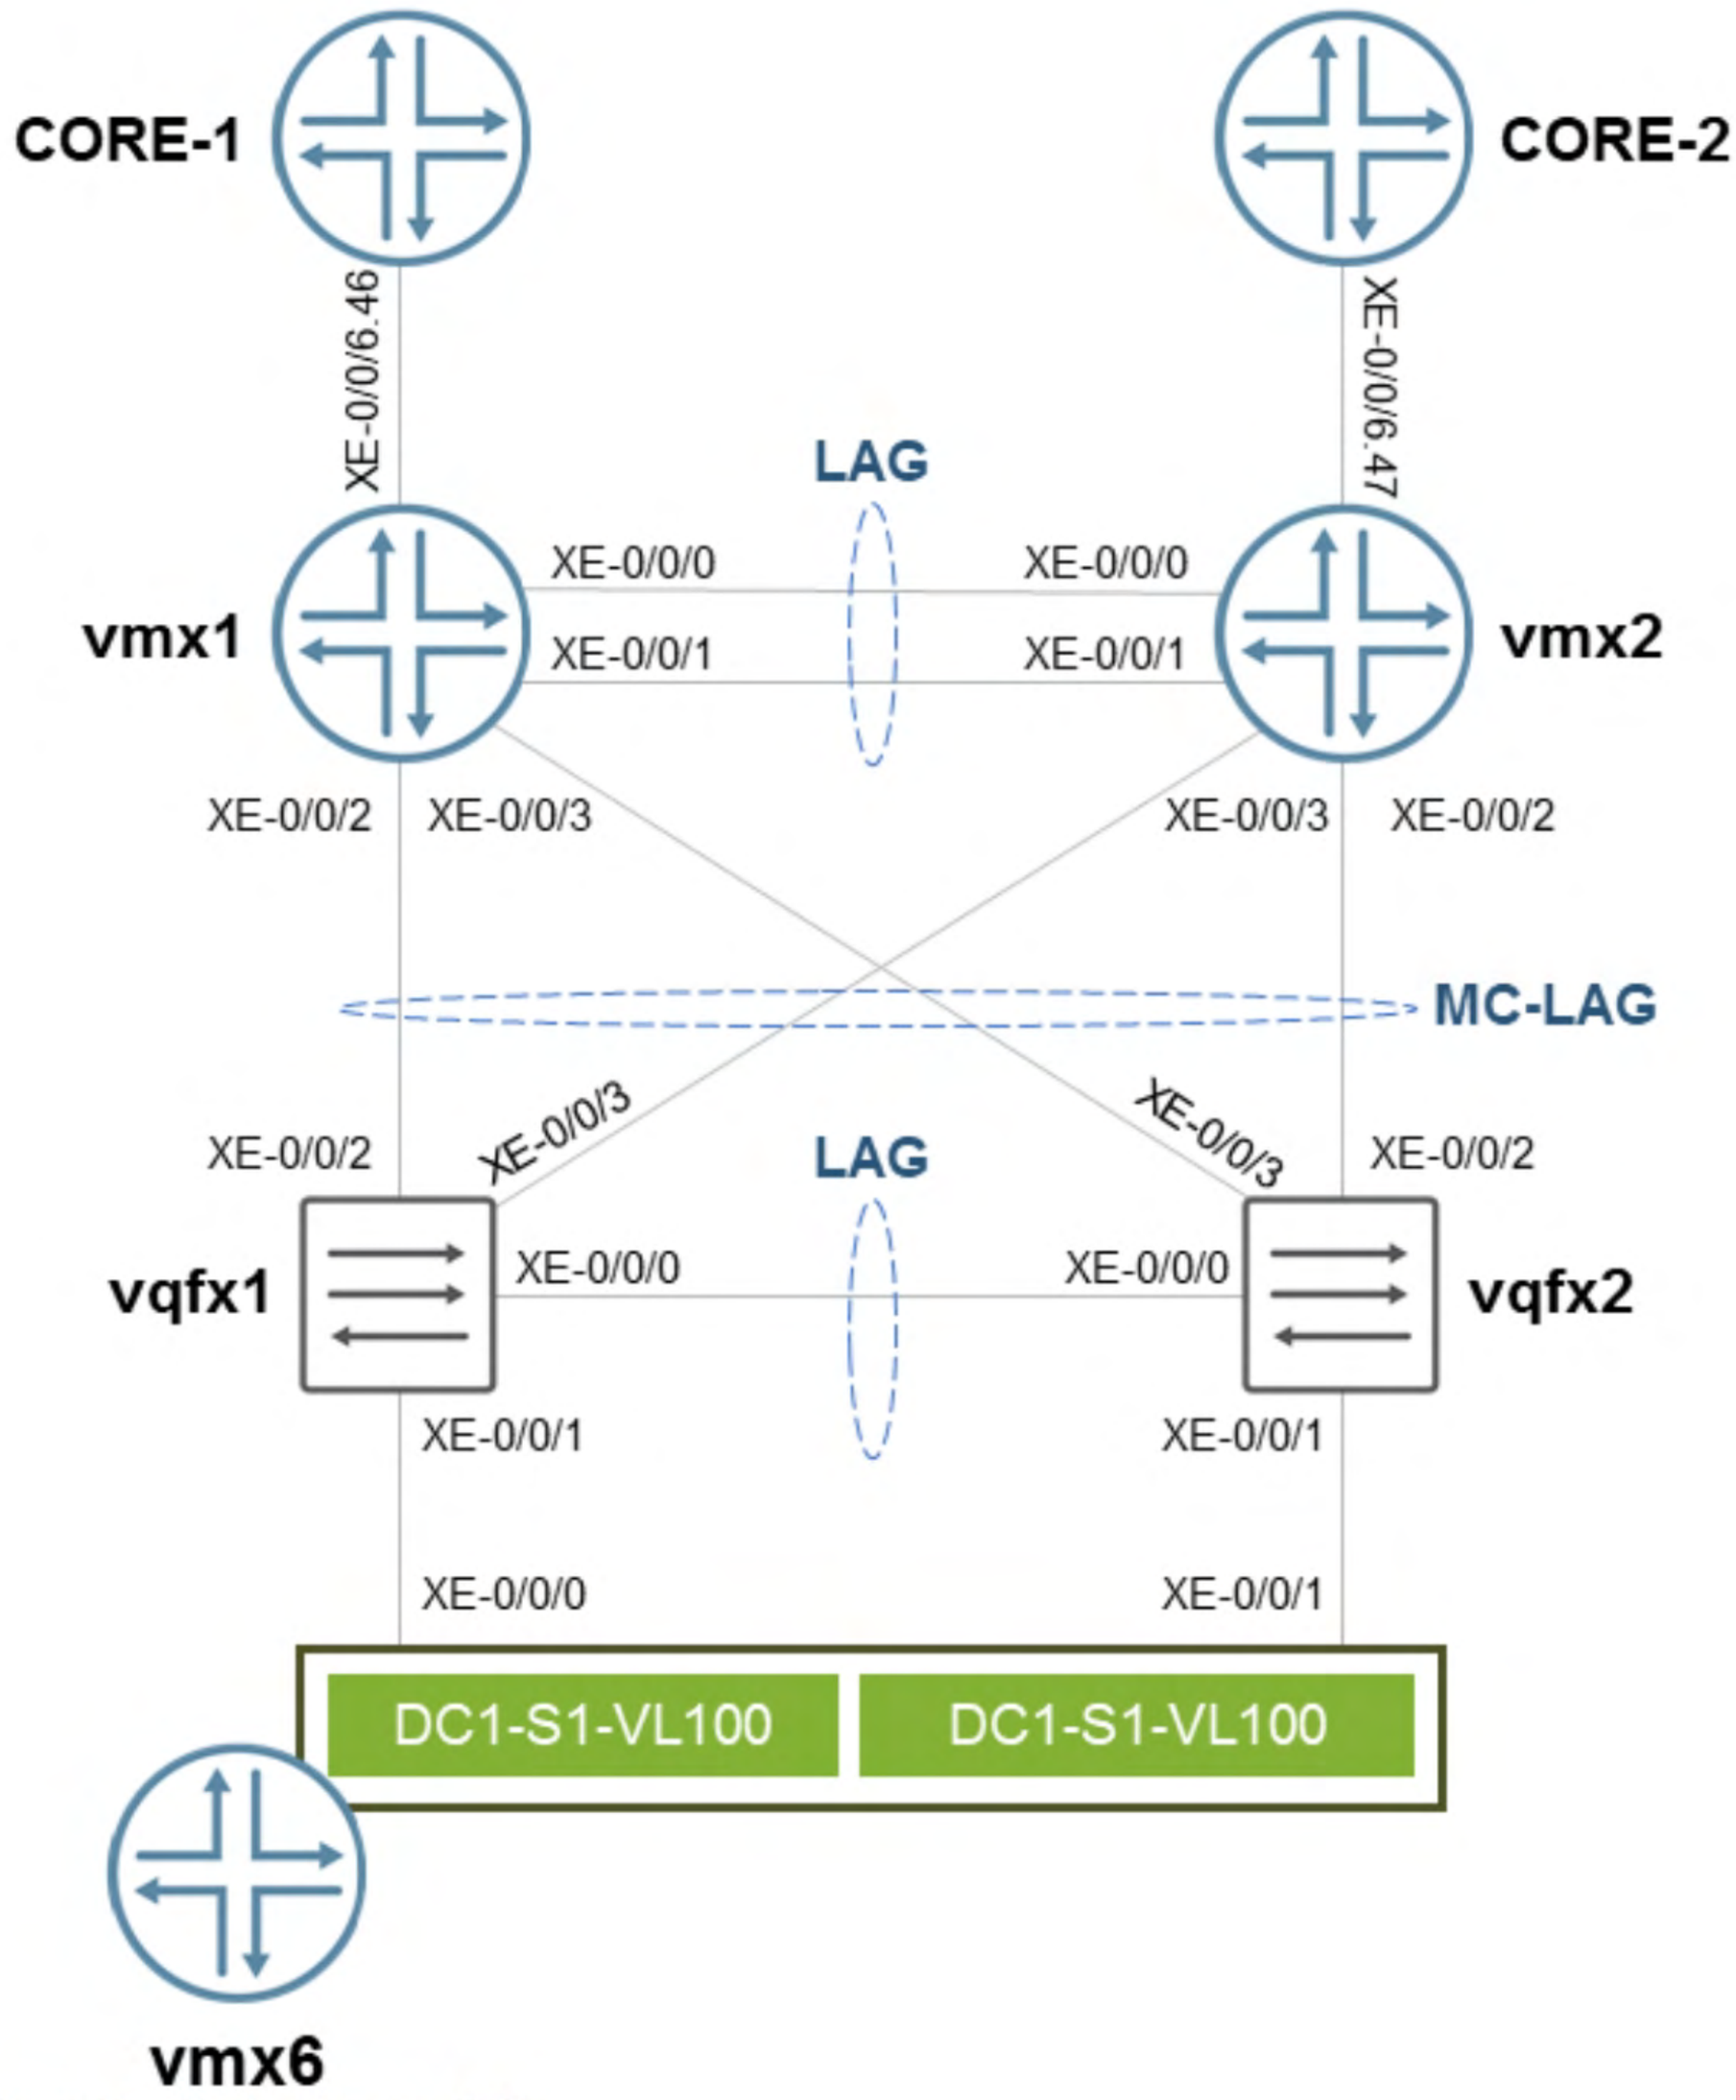

Figure 15 - Physical topology DC1

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Solution

vQFX1:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 ether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 44

set interfaces irb unit 44 family inet address 192.168.44.1/30

set vlans vlan-44 vlan-id 44
set vlans vlan-44 l3-interface irb.44
```

vQFX2:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/0 ether-options 802.3ad ae1

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 44

set interfaces irb unit 44 family inet address 192.168.44.2/30

set vlans vlan-44 vlan-id 44
set vlans vlan-44 l3-interface irb.44
```

vMX1:

```
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 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
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 45

set interfaces irb unit 45 family inet address 192.168.45.0/31
set interfaces irb unit 45 family inet6

set bridge-domains bd-45 vlan-id 45
set bridge-domains bd-45 routing-interface irb.45
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX2:

```
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 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
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 45

set interfaces irb unit 45 family inet address 192.168.45.1/31
set interfaces irb unit 45 family inet6

set bridge-domains bd-45 vlan-id 45
set bridge-domains bd-45 routing-interface irb.45
```

Verification

Checking the vQFX:

```
jncie@vQFX1> show lacp interfaces

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/0                  Current   Fast periodic Collecting distributing

jncie@vQFX1> ping 192.168.44.2 rapid

PING 192.168.44.2 (192.168.44.2): 56 data bytes
!!!!!
--- 192.168.44.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 32.947/46.719/96.647/24.979 ms

{master:0}
jncie@vQFX1> show vlans

Routing instance        VLAN name             Tag          Interfaces
default-switch          default               1
                                                           ae1.0*
default-switch          vlan-44               44
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Checking the vMX:

```
jncie@vMX1> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/0                  Current   Fast periodic Collecting distributing
      xe-0/0/1                  Current   Fast periodic Collecting distributing

jncie@vMX1> ping 192.168.45.1 rapid

PING 192.168.45.1 (192.168.45.1): 56 data bytes
!!!!!
--- 192.168.45.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.406/4.720/6.355/1.014 ms
jncie@vMX1> show bridge domain

Routing instance        Bridge domain            VLAN ID     Interfaces
default-switch          bd-45                    45
                                                             ae0.0
```

Task 1.8: DC1 Multistage MC-LAG

- The interfaces between the vMX and vQFX devices in DC1 should be part of an MC-LAG interface.
- Use the IP addresses that were configured in the previous task to source ICCP sessions.
- The vMX routers should only forward traffic to the vQFX devices when there is a minimum of 20Gbit is available towards the vQFX.
- Enable the MC-LAG and LAG interfaces for VLAN 100.
- Enable the server-facing interfaces for VLAN 100.
- All interfaces should be configured as trunk links.

Solution

vQFX1:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:01:02
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set interfaces ae1 unit 0 family ethernet-switching vlan members 100

set multi-chassis multi-chassis-protection 192.168.44.2 interface ae1

set protocols iccp local-ip-addr 192.168.44.1
set protocols iccp peer 192.168.44.2 redundancy-group-id-list 1
set protocols iccp peer 192.168.44.2 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.44.2 authentication-key jncie

set switch-options service-id 1
set vlans vlan-100 vlan-id 100
```

vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 100

set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:01:02
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set interfaces ae1 unit 0 family ethernet-switching vlan members 100

set multi-chassis multi-chassis-protection 192.168.44.1 interface ae1

set protocols iccp local-ip-addr 192.168.44.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols iccp peer 192.168.44.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.44.1 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.44.1 authentication-key jncie

set switch-options service-id 1

set vlans vlan-100 vlan-id 100
```

vMX1:

```
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1

set interfaces ae0 unit 0 family bridge vlan-id-list 100

set interfaces ae1 aggregated-ether-options minimum-links 2
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 01:02:01:02:00:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 100

set multi-chassis multi-chassis-protection 192.168.45.1 interface ae0

set protocols iccp local-ip-addr 192.168.45.0
set protocols iccp peer 192.168.45.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.45.1 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.45.1 authentication-key jncie

set bridge-domains bd-100 vlan-id 100

set switch-options service-id 1
```

vMX2:

```
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1

set interfaces ae0 unit 0 family bridge vlan-id-list 100

set interfaces ae1 aggregated-ether-options minimum-links 2
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 01:02:01:02:00:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 100

set multi-chassis multi-chassis-protection 192.168.45.0 interface ae0

set protocols iccp local-ip-addr 192.168.45.1
set protocols iccp peer 192.168.45.0 redundancy-group-id-list 1
set protocols iccp peer 192.168.45.0 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.45.0 authentication-key jncie

set bridge-domains bd-100 vlan-id 100

set switch-options service-id 1
```

Verification

Verifying the vQFX:

```
jncie@vQFX1> show vlans 100 detail

Routing instance: default-switch
  VLAN Name: vlan-100                          State: Active
Tag: 100
Internal index: 4, Generation Index: 10, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ae0.0*,tagged,trunk
    ae1.0*,tagged,trunk
    xe-0/0/1.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 2

jncie@vQFX1> show iccp detail

Redundancy Group Information for peer 192.168.44.2
  TCP Connection       : Established
  Liveliness Detection : Up
  Redundancy Group ID          Status
    1                           Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1

jncie@vQFX1> show interfaces mc-ae extensive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Member Link                    : ae0
Current State Machine's State: mcae active state
Local Status                   : active
Local State                    : up
Peer Status                    : active
Peer State                     : up
    Logical Interface          : ae0.0
    Topology Type              : bridge
    Local State                : up
    Peer State                 : up
    Peer Ip/MCP/State          : 192.168.44.2 ae1.0 up

 MCAE Configuration
    Redundancy Group           : 1
    MCAE ID                    : 1
    MCAE Mode                  : active_active
    Status Control             : active
    Chassis ID                 : 0
 LACP Configuration
    System ID                  : 00:00:00:00:01:02
    Admin Key                  : 1
```

```
jncie@vQFX1> show lacp interfaces ae0

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/3                  Current   Fast periodic Collecting distributing
```

Verifying the vMX:

```
jncie@vMX1> show bridge domain bd-100

Routing instance        Bridge domain            VLAN ID     Interfaces
default-switch          bd-100                   100
                                                             ae0.0
                                                             ae1.0
```

```
jncie@vMX1> show iccp detail

Redundancy Group Information for peer 192.168.45.1
  TCP Connection       : Established
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Liveliness Detection : Up
  Redundancy Group ID           Status
    1                            Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1

jncie@vMX1> show interfaces mc-ae extensive

 Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.45.1 ae0.0 up

 MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 1
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 0
 LACP Configuration
     System ID                : 01:02:01:02:00:01
     Admin Key                : 1

jncie@vMX1> show interfaces ae1 extensive

Physical interface: ae1 (MC-AE-1, active), Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 529, Generation: 159
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 2, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:72:fb:c6, Hardware address: 00:05:86:72:fb:c6
  Last flapped   : 2017-05-28 11:07:48 UTC (00:01:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              325620                 560 bps
   Output bytes  :             2112370                2080 bps
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 Input  packets:                3803                    0 pps
 Output packets:               16249                    2 pps
 IPv6 transit statistics:
  Input  bytes  :                  0
  Output bytes  :                  0
  Input  packets:                  0
  Output packets:                  0
Dropped traffic statistics due to STP State:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                              0                    0                    0
    1                              0                    0                    0
    2                              0                    0                    0
    3                              0                    0                    0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                              0                    0                    0
    1                              0                    0                    0
    2                              0                    0                    0
    3                              0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
```

```
Logical interface ae1.0 (Index 344) (SNMP ifIndex 532) (Generation 176)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    xe-0/0/2.0
      Input :             0          0             0            0
      Output:             0          0             0            0
    xe-0/0/3.0
      Input :             0          0             0            0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Output:              0          0              0              0


Aggregate member links: 2

LACP info:          Role     System            System         Port      Port     Port
                             priority      identifier   priority    number      key
  xe-0/0/2.0       Actor          127  01:02:01:02:00:01        127         3        1
  xe-0/0/2.0     Partner          127  00:00:00:00:01:02        127         2        1
  xe-0/0/3.0       Actor          127  01:02:01:02:00:01        127         4        1
  xe-0/0/3.0     Partner          127  00:00:00:00:01:02        127     32771        1
LACP Statistics:         LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
  xe-0/0/2.0                   0           0            0            0
  xe-0/0/3.0                   0           0            0            0
Marker Statistics:     Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
  xe-0/0/2.0                   0           0            0            0
  xe-0/0/3.0                   0           0            0            0
Protocol bridge, MTU: 1514, Generation: 207, Route table: 7, Mesh Group: __all_ces__
  Flags: Trunk-Mode
```

Task 1.9: DC1 Gateway and OSPF Configuration

- Configure 192.168.1.254/24 as the IPv4 gateway address on the vMX devices for VLAN 100.
- Configure 2001:100:0:100::254 as the IPv6 default gateway for VLAN 100.
- For IPv4, configure an anycast gateway and for IPv6 configure VRRP. Allow hosts to ping the IPv6 default gateway.
- Enable the vMX uplinks for OSPF and OSPFv3 and advertise all IRB and lo0 interfaces.
- Ensure that the loss of a single link between the vQFX and vQFX devices will not impact reachability.
- OSPF links carrying neighbor relationships should be configured as point-to-point. It should not be possible to establish OSPF neighbor adjacencies across server facing interfaces.

Solution

vMX1:

```
set interfaces irb unit 100 family inet address 192.168.1.254/24
set interfaces irb unit 100 mac 00:00:00:00:11:22

set interfaces irb unit 100 family inet6 address 2001:100:0:100::252/64 vrrp-inet6-group 100
virtual-inet6-address 2001:100:0:100::254
set interfaces irb unit 100 family inet6 address 2001:100:0:100::252/64 vrrp-inet6-group 100
virtual-link-local-address fe80::0000:0:100:1
set interfaces irb unit 100 family inet6 address 2001:100:0:100::252/64 vrrp-inet6-group 100
priority 200
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces irb unit 100 family inet6 address 2001:100:0:100::252/64 vrrp-inet6-group 100
accept-data

set routing-options router-id 10.0.0.1

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/6.46 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.45 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.100 passive

set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface xe-0/0/6.46 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.45 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.100 passive

set bridge-domains bd-100 routing-interface irb.100
```

vMX2:

```
set interfaces irb unit 100 family inet address 192.168.1.254/24
set interfaces irb unit 100 mac 00:00:00:00:11:22

set interfaces irb unit 100 family inet6 address 2001:100:0:100::253/64 vrrp-inet6-group 100
virtual-inet6-address 2001:100:0:100::254
set interfaces irb unit 100 family inet6 address 2001:100:0:100::253/64 vrrp-inet6-group 100
virtual-link-local-address fe80::0000:0:100:1
set interfaces irb unit 100 family inet6 address 2001:100:0:100::253/64 vrrp-inet6-group 100
priority 200
set interfaces irb unit 100 family inet6 address 2001:100:0:100::253/64 vrrp-inet6-group 100
accept-data

set routing-options router-id 10.0.0.2

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/6.47 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.45 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.100 passive

set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface xe-0/0/6.47 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.45 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.100 passive

set bridge-domains bd-100 routing-interface irb.100
```

Verification

Verifying OSPF and OSPFv3:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
192.168.45.1     irb.45                 Full      10.0.0.2         128    37
192.168.46.2     xe-0/0/6.46            Full      10.0.0.5         128    35

jncie@vMX1> show ospf3 neighbor

ID               Interface              State     Pri   Dead
10.0.0.2         irb.45                 Full      128     33
  Neighbor-address fe80::205:8600:2d94:9bf0
10.0.0.5         xe-0/0/6.46            Full      128     39
  Neighbor-address fe80::20c:2900:2eb9:e97d

jncie@vMX1> show route protocol ospf

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

0.0.0.0/0          *[OSPF/150] 00:00:44, metric 0, tag 0
                    > to 192.168.46.2 via xe-0/0/6.46
10.0.0.2/32        *[OSPF/10] 00:00:31, metric 1
                    > to 192.168.45.1 via irb.45
10.0.0.5/32        *[OSPF/10] 00:00:44, metric 1
                    > to 192.168.46.2 via xe-0/0/6.46
10.0.0.6/32        *[OSPF/10] 00:00:26, metric 2
                      to 192.168.46.2 via xe-0/0/6.46
                    > to 192.168.45.1 via irb.45
192.168.0.0/30     *[OSPF/10] 00:00:44, metric 2
                    > to 192.168.46.2 via xe-0/0/6.46
192.168.47.0/30    *[OSPF/10] 00:00:31, metric 2
                    > to 192.168.45.1 via irb.45
192.168.48.0/30    *[OSPF/10] 00:00:44, metric 2
                    > to 192.168.46.2 via xe-0/0/6.46
192.168.49.0/30    *[OSPF/10] 00:00:26, metric 3
                      to 192.168.46.2 via xe-0/0/6.46
                    > to 192.168.45.1 via irb.45
224.0.0.5/32       *[OSPF/10] 00:00:49, metric 1
                      MultiRecv

inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::2/128    *[OSPF3/10] 00:00:32, metric 1
                    > to fe80::205:8600:2d94:9bf0 via irb.45
2001:100::5/128    *[OSPF3/10] 00:00:44, metric 1
                    > to fe80::20c:2900:2eb9:e97d via xe-0/0/6.46
2001:100::6/128    *[OSPF3/10] 00:00:22, metric 2
                    > to fe80::20c:2900:2eb9:e97d via xe-0/0/6.46
                      to fe80::205:8600:2d94:9bf0 via irb.45
ff02::5/128        *[OSPF3/10] 00:00:49, metric 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    MultiRecv

jncie@vMX1> show ospf database advertising-router self extensive

    OSPF database, Area 0.0.0
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router  *10.0.0.1          10.0.0.1        0x80000003   632  0x22 0x16c9  96
  bits 0x0, link count 6
  id 192.168.1.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 1
…

jncie@vMX1> show ospf3 database advertising-router self extensive

    OSPF3 database, Area 0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          10.0.0.1        0x80000003   666  0x3b7e  56
…
IntraArPfx *0.0.0.1          10.0.0.1        0x80000002   673  0xede1  64
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.0.0.1
  Prefix-count 2
  Prefix 2001:100:0:100::/64
    Prefix-options 0x0, Metric 1
  Prefix 2001:100::1/128
    Prefix-options 0x2, Metric 0
  Gen timer 00:27:35
  Aging timer 00:48:46
  Installed 00:11:13 ago, expires in 00:48:47, sent 00:11:07 ago
  Last changed 00:11:13 ago, Change count: 2, Ours

...
```

Verifying VRRP:

```
jncie@vMX1> show vrrp

Interface     State       Group   VR state VR Mode   Timer    Type   Address
irb.100       up            100   master   Active      A  0.677 lcl    2001:100:0:100::252
                                                               vip    fe80::100:1
                                                               vip    2001:100:0:100::254

jncie@vMX2> show vrrp

Interface     State       Group   VR state VR Mode   Timer    Type   Address
irb.100       up            100   backup   Active      D  2.672 lcl    2001:100:0:100::253
                                                               vip    fe80::100:1
                                                               vip    2001:100:0:100::254
                                                               mas   fe80::205:8600:6472:fbf0
```

Verifying IPV4 connectivity from both servers to the lo0 IP of the core routers:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.0.0.5

PING 10.0.0.5 (10.0.0.5): 56 data bytes
!!!!!
--- 10.0.0.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 23.058/28.301/36.584/4.522 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.0.0.6

PING 10.0.0.6 (10.0.0.6): 56 data bytes
!!!!!
--- 10.0.0.6 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 30.395/35.648/48.077/6.335 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.0.0.5

PING 10.0.0.5 (10.0.0.5): 56 data bytes
!!!!!
--- 10.0.0.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 31.803/43.348/68.820/13.600 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.0.0.6

PING 10.0.0.6 (10.0.0.6): 56 data bytes
!!!!!
--- 10.0.0.6 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 36.339/80.292/145.735/40.611 ms

Verifying IPV6 connectivity from both servers to the lo0 IP of the core routers:

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 2001:100::5

PING6(56=40+8+8 bytes) 2001:100:0:100::1 --> 2001:100::5
!!!!!
--- 2001:100::5 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 23.448/30.877/51.997/10.838 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 2001:100::6

PING6(56=40+8+8 bytes) 2001:100:0:100::1 --> 2001:100::6
.!!!!
--- 2001:100::6 ping6 statistics ---
5 packets transmitted, 4 packets received, 20% packet loss
round-trip min/avg/max/std-dev = 19.589/197.421/712.252/297.262 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 2001:100::5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
PING6(56=40+8+8 bytes) 2001:100:0:100::2 --> 2001:100::5
!!!!!
--- 2001:100::5 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 30.075/32.708/37.084/2.606 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 2001:100::6

PING6(56=40+8+8 bytes) 2001:100:0:100::2 --> 2001:100::6
!!!!!
--- 2001:100::6 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 28.305/30.749/35.337/2.412 ms
```

Verifying IPV6 connectivity to the VIP:

```
jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 2001:100:0:100::254

PING6(56=40+8+8 bytes) 2001:100:0:100::2 --> 2001:100:0:100::254
!!!!!
--- 2001:100:0:100::254 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 31.621/36.296/45.131/4.952 ms
```

Task 1.10: DC1 Security

- Configure storm control to drop all BUM traffic over 5Mbit on the server-facing links.
- Allow only 1Gbit of traffic for VLAN 100 on the vQFX devices. Ensure that VLANs added in the future will not have their bandwidth restricted.
- Allow for a maximum of 20 MAC addresses to be learned in VLAN 100. Traffic for additional hosts should be dropped.
- Limit the movement of MAC addresses to 2 moves per second.
- Cap the cumulative of IPv4 and IPv6 traffic for VLAN 100 on the vMX routers to 1Gbit.

Solution

We can rate limit the cumulative of IPv4 and IPv6 traffic through the use of a logical interface policer. This type of policer will rate limit traffic for multiple protocol families on the same logical interface. We can reference the logical interface policer from a firewall filter, or we can configure it on the interface directly.

Since this task only requires us to police the traffic, we do not need any of the additional options a firewall filter could offer us, and we configure the policer on the interface directly.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX1 and vMX2:

```
set interfaces irb unit 100 family inet policer input 1g
set interfaces irb unit 100 family inet policer output 1g

set interfaces irb unit 100 family inet6 policer input 1g
set interfaces irb unit 100 family inet6 policer output 1g

set firewall policer 1g logical-interface-policer
set firewall policer 1g if-exceeding bandwidth-limit 1g
set firewall policer 1g if-exceeding burst-size-limit 50m
set firewall policer 1g then discard
```

vQFX1 and vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching filter input access
set interfaces xe-0/0/1 unit 0 family ethernet-switching storm-control dc1

set forwarding-options storm-control-profiles dc1 all bandwidth-level 5000

set firewall family ethernet-switching filter access term 1g-vlan-100 from user-vlan-id 100
set firewall family ethernet-switching filter access term 1g-vlan-100 then policer 1g
set firewall family ethernet-switching filter access term default-accept then accept

set firewall policer 1g if-exceeding bandwidth-limit 1g
set firewall policer 1g if-exceeding burst-size-limit 50m
set firewall policer 1g then discard

set vlans vlan-100 switch-options mac-table-size 20
set vlans vlan-100 switch-options mac-table-size packet-action drop
set vlans vlan-100 switch-options interface-mac-limit 10
set vlans vlan-100 switch-options interface-mac-limit packet-action drop
set vlans vlan-100 switch-options mac-move-limit 2
```

Verification

Verifying the logical policer on the vMX router:

```
jncie@vMX1> show interfaces irb.100 extensive | match Policer

      Policer: Input: 1g-irb.100-log_int-i, Output: 1g-irb.100-log_int-o
      Policer: Input: 1g-irb.100-log_int-i, Output: 1g-irb.100-log_int-o
      Policer: Input: __default_arp_policer__

jncie@vMX1> show policer 1g-irb.100-log_int-i

Policers:
Name                                                Bytes              Packets
1g-irb.100-log_int-i                                    0                    0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show policer 1g-irb.100-log_int-o

Policers:
Name                                                Bytes              Packets
1g-irb.100-log_int-o                                    0                    0
```

Verifying the configuration that was applied to the vQFX:

```
jncie@vQFX1> show interfaces xe-0/0/1 detail | match Filters

      Input Filters: access
```

```
jncie@vQFX1> show firewall filter access

Filter: access
Policers:
Name                                                Bytes              Packets
1g-1g-vlan-100                                          0                    0
```

Task 1.11: DC2 LAG and MC-LAG Configuration

- Configure the vQFX uplinks towards the vMX devices as LAG interfaces.
- Configure the downlinks from the vMX towards the vQFX as MC-LAGs.
- Enable the LAG and MC-LAG interfaces for VLAN 100.
- Enable the server-facing interfaces on the vQFX for VLAN 100.
- All links should transport tagged traffic.
- On vQFX3 and vQFX4, configure the server-facing interfaces as a LAG interface.

Solution

vMX3:

```
set chassis aggregated-devices ethernet device-count 4

set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0

set interfaces xe-0/0/2 gigether-options 802.3ad ae1

set interfaces xe-0/0/3 gigether-options 802.3ad ae2

set interfaces xe-0/0/4 gigether-options 802.3ad ae3

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
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 50
set interfaces ae0 unit 0 family bridge vlan-id-list 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 03:04:03:04:00:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 100

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 03:04:03:04:00:02
set interfaces ae2 aggregated-ether-options lacp admin-key 1
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 100

set interfaces ae3 flexible-vlan-tagging
set interfaces ae3 encapsulation flexible-ethernet-services
set interfaces ae3 aggregated-ether-options lacp active
set interfaces ae3 aggregated-ether-options lacp system-id 03:04:03:04:00:03
set interfaces ae3 aggregated-ether-options lacp admin-key 1
set interfaces ae3 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae3 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae3 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae3 aggregated-ether-options mc-ae mode active-active
set interfaces ae3 aggregated-ether-options mc-ae status-control active
set interfaces ae3 unit 0 family bridge interface-mode trunk
set interfaces ae3 unit 0 family bridge vlan-id-list 100

set interfaces irb unit 50 family inet address 192.168.50.0/31
set interfaces irb unit 50 family inet6

set multi-chassis multi-chassis-protection 192.168.50.1 interface ae0

set protocols iccp local-ip-addr 192.168.50.0
set protocols iccp peer 192.168.50.1 redundancy-group-id-list 1
set protocols iccp peer 192.168.50.1 backup-liveness-detection backup-peer-ip 10.10.20.4
set protocols iccp peer 192.168.50.1 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.50.1 authentication-key jncie
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set bridge-domains bd-100 vlan-id 100

set bridge-domains bd-50 vlan-id 50
set bridge-domains bd-50 routing-interface irb.50

set switch-options service-id 1
```

vMX4:

```
set chassis aggregated-devices ethernet device-count 4

set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0

set interfaces xe-0/0/2 gigether-options 802.3ad ae3

set interfaces xe-0/0/3 gigether-options 802.3ad ae2

set interfaces xe-0/0/4 gigether-options 802.3ad ae1

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
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 50
set interfaces ae0 unit 0 family bridge vlan-id-list 100

set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 03:04:03:04:00:01
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 100

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 03:04:03:04:00:02
set interfaces ae2 aggregated-ether-options lacp admin-key 1
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 2
set interfaces ae2 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control standby
set interfaces ae2 unit 0 family bridge interface-mode trunk
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae2 unit 0 family bridge vlan-id-list 100

set interfaces ae3 flexible-vlan-tagging
set interfaces ae3 encapsulation flexible-ethernet-services
set interfaces ae3 aggregated-ether-options lacp active
set interfaces ae3 aggregated-ether-options lacp system-id 03:04:03:04:00:03
set interfaces ae3 aggregated-ether-options lacp admin-key 1
set interfaces ae3 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae3 aggregated-ether-options mc-ae redundancy-group 1
set interfaces ae3 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae3 aggregated-ether-options mc-ae mode active-active
set interfaces ae3 aggregated-ether-options mc-ae status-control standby
set interfaces ae3 unit 0 family bridge interface-mode trunk
set interfaces ae3 unit 0 family bridge vlan-id-list 100

set interfaces irb unit 50 family inet address 192.168.50.1/31
set interfaces irb unit 50 family inet6

set multi-chassis multi-chassis-protection 192.168.50.0 interface ae0

set protocols iccp local-ip-addr 192.168.50.1
set protocols iccp peer 192.168.50.0 redundancy-group-id-list 1
set protocols iccp peer 192.168.50.0 backup-liveness-detection backup-peer-ip 10.10.20.3
set protocols iccp peer 192.168.50.0 liveness-detection minimum-interval 3000
set protocols iccp peer 192.168.50.0 authentication-key jncie

set bridge-domains bd-100 vlan-id 100

set bridge-domains bd-50 vlan-id 50
set bridge-domains bd-50 routing-interface irb.50

set switch-options service-id 1
```

vQFX3:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/4 ether-options 802.3ad ae0

set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 100
set interfaces ae1 unit 0 family ethernet-switching vlan members 1000
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set vlans vlan-100 vlan-id 100
```

vQFX4:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/4 ether-options 802.3ad ae0

set interfaces xe-0/0/1 ether-options 802.3ad ae1
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 01:01:01:01:01:01
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 100
set interfaces ae1 unit 0 family ethernet-switching vlan members 1000

set vlans vlan-100 vlan-id 100
```

vQFX5:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 100

set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/4 ether-options 802.3ad ae0

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set vlans vlan-100 vlan-id 100
```

Verification

Checking the LAG and MC-LAGs on the vMX:

vMX3:

```
jncie@vMX3> show bridge domain

Routing instance        Bridge domain            VLAN ID     Interfaces
default-switch          bd-100                   100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                                                                  ae0.0
                                                                  ae1.0
                                                                  ae2.0
                                                                  ae3.0
default-switch          bd-50                       50
                                                                  ae0.0

jncie@vMX3> show iccp

Redundancy Group Information for peer 192.168.50.1
  TCP Connection       : Established
  Liveliness Detection : Up
Backup liveness peer status: Up
  Redundancy Group ID          Status
    1                           Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1

jncie@vMX3> show bfd session

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.50.1             Up                       9.000     3.000        3

1 sessions, 1 clients
Cumulative transmit rate 0.3 pps, cumulative receive rate 0.3 pps

jncie@vMX3> show interfaces mc-ae

 Member Link                  : ae1
  Current State Machine's State: mcae active state
  Local Status                 : active
  Local State                  : up
  Peer Status                  : active
  Peer State                   : up
      Logical Interface        : ae1.0
      Topology Type            : bridge
      Local State              : up
      Peer State               : up
      Peer Ip/MCP/State        : 192.168.50.1 ae0.0 up

  Member Link                  : ae2
  Current State Machine's State: mcae active state
  Local Status                 : active
  Local State                  : up
  Peer Status                  : active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Peer State                      : up
    Logical Interface           : ae2.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.50.1 ae0.0 up

Member Link                     : ae3
Current State Machine's State: mcae active state
Local Status                    : active
Local State                     : up
Peer Status                     : active
Peer State                      : up
    Logical Interface           : ae3.0
    Topology Type               : bridge
    Local State                 : up
    Peer State                  : up
    Peer Ip/MCP/State           : 192.168.50.1 ae0.0 up

jncie@vMX3> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/0                Current   Fast periodic Collecting distributing
      xe-0/0/1                Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                Current   Fast periodic Collecting distributing

Aggregated interface: ae2
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3                Current   Fast periodic Collecting distributing

Aggregated interface: ae3
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/4                Current   Fast periodic Collecting distributing
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
jncie@vMX4> show bridge domain

Routing instance        Bridge domain            VLAN ID     Interfaces
default-switch          bd-100                   100
                                                             ae0.0
                                                             ae1.0
                                                             ae2.0
                                                             ae3.0
default-switch          bd-50                    50
                                                             ae0.0

jncie@vMX4> show iccp

Redundancy Group Information for peer 192.168.50.0
  TCP Connection       : Established
  Liveliness Detection : Up

Backup liveness peer status: Up
  Redundancy Group ID          Status
    1                           Up

Client Application: l2ald_iccpd_client
  Redundancy Group IDs Joined: 1

Client Application: lacpd
  Redundancy Group IDs Joined: 1

jncie@vMX4> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.50.0             Up                       9.000     3.000        3

1 sessions, 1 clients
Cumulative transmit rate 0.3 pps, cumulative receive rate 0.3 pps

jncie@vMX4> show interfaces mc-ae

Member Link                  : ae1
 Current State Machine's State: mcae active state
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae1.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Peer Ip/MCP/State          : 192.168.50.0 ae0.0 up

Member Link                    : ae2
Current State Machine's State: mcae active state
Local Status                   : active
Local State                    : up
Peer Status                    : active
Peer State                     : up
    Logical Interface          : ae2.0
    Topology Type              : bridge
    Local State                : up
    Peer State                 : up
    Peer Ip/MCP/State          : 192.168.50.0 ae0.0 up

 Member Link                   : ae3
 Current State Machine's State: mcae active state
 Local Status                  : active
 Local State                   : up
 Peer Status                   : active
 Peer State                    : up
     Logical Interface         : ae3.0
     Topology Type             : bridge
     Local State               : up
     Peer State                : up
     Peer Ip/MCP/State         : 192.168.50.0 ae0.0 up

jncie@vMX4> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/0                  Current   Fast periodic Collecting distributing
      xe-0/0/1                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/4                  Current   Fast periodic Collecting distributing

Aggregated interface: ae2
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3                  Current   Fast periodic Collecting distributing
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Aggregated interface: ae3
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
```

Verifying the LAGs and VLAN on vQFX:

vQFX3:

```
jncie@vQFX3> show vlans 100

Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-100              100
                                                           ae0.0*
                                                           ae1.0*

{master:0}
jncie@vQFX3> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/4                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing
      xe-0/0/3                  Current   Fast periodic Collecting distributing
```

vQFX4:

```
jncie@vQFX4> show vlans 100

Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-100              100
                                                           ae0.0*
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                                                                ae1.0*

{master:0}
jncie@vQFX4> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/4                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing
      xe-0/0/3                  Current   Fast periodic Collecting distributing
```

vQFX5:

```
jncie@vQFX5> show vlans 100

Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-100              100
                                                           ae0.0*
                                                           xe-0/0/1.0*

{master:0}
jncie@vQFX5> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/2                  Current   Fast periodic Collecting distributing
      xe-0/0/4                  Current   Fast periodic Collecting distributing
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 1.12: DC2 Gateway, VRRP, and OSPF Configuration

- Configure the following gateways on the vMX routers:

IRB / VLAN	Gateway address IPv4	Gateway address IPv6
IRB100	192.168.100.254	2001:100:100:100::254

- Use VRRP and ensure that vMX3 is the active router.
- Make sure that you can ping the virtual-router and authenticate the VRRP-message for IPv4 with the MD5 hash 'jncie'.
- Ensure that vMX3 and vMX4 advertise the IRB as well as the loopback interface to the core.
- It should not be possible to form OSPF neighbor relationships across IRB 100.
- OSPF links carrying neighbor relationships should be configured as point-to-point. Use OSPFv2 for IPv4 and OSPFv3 for IPv6.

Solution

vMX3:

```
set interfaces irb unit 100 family inet address 192.168.100.252/24 vrrp-group 100 virtual-
address 192.168.100.254
set interfaces irb unit 100 family inet address 192.168.100.252/24 vrrp-group 100 priority 200
set interfaces irb unit 100 family inet address 192.168.100.252/24 vrrp-group 100 accept-data
set interfaces irb unit 100 family inet address 192.168.100.252/24 vrrp-group 100
authentication-key jncie

set interfaces irb unit 100 family inet6 address 2001:100:100:100::252/64 vrrp-inet6-group 100
virtual-inet6-address 2001:100:100:100::254
set interfaces irb unit 100 family inet6 address 2001:100:100:100::252/64 vrrp-inet6-group 100
virtual-link-local-address fe80::0000:0:100:1
set interfaces irb unit 100 family inet6 address 2001:100:100:100::252/64 vrrp-inet6-group 100
priority 200
set interfaces irb unit 100 family inet6 address 2001:100:100:100::252/64 vrrp-inet6-group 100
accept-data

set routing-options router-id 10.0.0.3

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/6.49 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.50 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.100 passive

set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface xe-0/0/6.49 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.50 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.100 passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set bridge-domains bd-100 routing-interface irb.100
```

vMX4:

```
set interfaces irb unit 100 family inet address 192.168.100.253/24 vrrp-group 100 virtual-
address 192.168.100.254
set interfaces irb unit 100 family inet address 192.168.100.253/24 vrrp-group 100 priority 50
set interfaces irb unit 100 family inet address 192.168.100.253/24 vrrp-group 100 accept-data
set interfaces irb unit 100 family inet address 192.168.100.253/24 vrrp-group 100
authentication-key jncie

set interfaces irb unit 100 family inet6 address 2001:100:100:100::253/64 vrrp-inet6-group 100
virtual-inet6-address 2001:100:100:100::254
set interfaces irb unit 100 family inet6 address 2001:100:100:100::253/64 vrrp-inet6-group 100
virtual-link-local-address fe80::0000:0:100:1
set interfaces irb unit 100 family inet6 address 2001:100:100:100::253/64 vrrp-inet6-group 100
priority 200
set interfaces irb unit 100 family inet6 address 2001:100:100:100::253/64 vrrp-inet6-group 100
accept-data

set routing-options router-id 10.0.0.4

set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface irb.50 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/6.48 interface-type p2p
set protocols ospf area 0.0.0.0 interface irb.100 passive

set protocols ospf3 area 0.0.0.0 interface lo0.0
set protocols ospf3 area 0.0.0.0 interface irb.50 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/6.48 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface irb.100 passive

set bridge-domains bd-100 routing-interface irb.100
```

Verification

Verifying OSPF and OSPFv3 on vMX3:

```
jncie@vMX3> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
192.168.50.1     irb.50                 Full      10.0.0.4         128    39
192.168.49.2     xe-0/0/6.49            Full      10.0.0.6         128    35

jncie@vMX3> show ospf3 neighbor

ID               Interface              State     Pri   Dead
10.0.0.4         irb.50                 Full      128     39
  Neighbor-address fe80::205:8600:3282:68f0
10.0.0.6         xe-0/0/6.49            Full      128     32
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Neighbor-address fe80::20c:2900:31b9:e97d

jncie@vMX3> show ospf interface detail irb.100

Interface          State   Area            DR ID           BDR ID          Nbrs
irb.100            DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 192.168.100.252, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 0, Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 1

jncie@vMX3> show ospf3 interface detail irb.100

Interface          State   Area            DR ID           BDR ID          Nbrs
irb.100            DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Address fe80::205:8600:644d:34f0, Prefix-length 64
  OSPF3-Intf-index 4, Type LAN, MTU 1500, Cost 1
  Adj count: 0, Router LSA ID: -, Passive
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

jncie@vMX3> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *10.0.0.3         10.0.0.3        0x800000a2   557  0x22 0xa221  96
  bits 0x0, link count 6
  id 192.168.100.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 1
  id 10.0.0.4, data 192.168.50.0, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 192.168.50.0, data 255.255.255.254, Type Stub (3)
    Topology count: 0, Default metric: 1
  id 10.0.0.3, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  id 10.0.0.6, data 192.168.49.1, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 192.168.49.0, data 255.255.255.252, Type Stub (3)
    Topology count: 0, Default metric: 1
  Topology default (ID 0)
    Type: PointToPoint, Node ID: 10.0.0.6
      Metric: 1, Bidirectional
    Type: PointToPoint, Node ID: 10.0.0.4
      Metric: 1, Bidirectional

jncie@vMX3> show ospf3 database advertising-router self detail

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Router     *0.0.0.0          10.0.0.3         0x800000a8   547  0xcf44  56
  bits 0x0, Options 0x33
  Type PointToPoint (1), Metric 1
      Loc-If-Id 1, Nbr-If-Id 1, Nbr-Rtr-Id 10.0.0.4
  Type PointToPoint (1), Metric 1
      Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.0.0.6
  Type: PointToPoint, Node ID: 10.0.0.6, Metric: 1, Bidirectional
  Type: PointToPoint, Node ID: 10.0.0.4, Metric: 1, Bidirectional
IntraArPfx *0.0.0.1          10.0.0.3         0x8000006e   552  0x98c3  64
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.0.0.3
  Prefix-count 2
  Prefix 2001:100:100:100::/64
    Prefix-options 0x0, Metric 1
  Prefix 2001:100::3/128
    Prefix-options 0x2, Metric 0

    OSPF3 Link-Local database, interface irb.50 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          10.0.0.3         0x80000002   552  0xba1b  44
  fe80::205:8600:324d:34f0
  Options 0x33, Priority 128
  Prefix-count 0

    OSPF3 Link-Local database, interface xe-0/0/6.49 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.3          10.0.0.3         0x80000001   694  0xd879  44
  fe80::20c:2900:31cf:e19a
  Options 0x33, Priority 128
  Prefix-count 0
```

Verifying OSPF and OSPFv3 on vMX4:

```
jncie@vMX4> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
192.168.50.0     irb.50                 Full      10.0.0.3         128    36
192.168.48.2     xe-0/0/6.48            Full      10.0.0.5         128    37

jncie@vMX4> show ospf3 neighbor

ID               Interface              State     Pri   Dead
10.0.0.3         irb.50                 Full      128     36
  Neighbor-address fe80::205:8600:324d:34f0
10.0.0.5         xe-0/0/6.48            Full      128     36
  Neighbor-address fe80::20c:2900:30b9:e97d

jncie@vMX4> show ospf interface detail irb.100

Interface          State   Area            DR ID           BDR ID          Nbrs
irb.100            DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 192.168.100.253, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 0, Passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 1

jncie@vMX4> show ospf3 interface detail irb.100

Interface          State   Area            DR ID           BDR ID          Nbrs
irb.100            DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Address fe80::205:8600:6482:68f0, Prefix-length 64
  OSPF3-Intf-index 4, Type LAN, MTU 1500, Cost 1
  Adj count: 0, Router LSA ID: -, Passive
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None

jncie@vMX4> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *10.0.0.4         10.0.0.4         0x800000ab   608  0x22 0xf5c4  96
  bits 0x0, link count 6
  id 192.168.100.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 1
  id 10.0.0.3, data 192.168.50.1, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 192.168.50.0, data 255.255.255.254, Type Stub (3)
    Topology count: 0, Default metric: 1
  id 10.0.0.4, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  id 10.0.0.5, data 192.168.48.1, Type PointToPoint (1)
    Topology count: 0, Default metric: 1
  id 192.168.48.0, data 255.255.255.252, Type Stub (3)
    Topology count: 0, Default metric: 1
  Topology default (ID 0)
    Type: PointToPoint, Node ID: 10.0.0.5
      Metric: 1, Bidirectional
    Type: PointToPoint, Node ID: 10.0.0.3
      Metric: 1, Bidirectional

jncie@vMX4> show ospf3 database advertising-router self detail

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          10.0.0.4         0x800000ce  596  0x41ad  56
  bits 0x0, Options 0x33
  Type PointToPoint (1), Metric 1
      Loc-If-Id 1, Nbr-If-Id 1, Nbr-Rtr-Id 10.0.0.3
  Type PointToPoint (1), Metric 1
      Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.0.0.5
  Type: PointToPoint, Node ID: 10.0.0.5, Metric: 1, Bidirectional
  Type: PointToPoint, Node ID: 10.0.0.3, Metric: 1, Bidirectional
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
IntraArPfx *0.0.0.1          10.0.0.4         0x80000099   601  0x78b5  64
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.0.0.4
  Prefix-count 2
  Prefix 2001:100:100:100::/64
    Prefix-options 0x0, Metric 1
  Prefix 2001:100::4/128
    Prefix-options 0x2, Metric 0

    OSPF3 Link-Local database, interface irb.50 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          10.0.0.4         0x80000002   601  0x2546  44
  fe80::205:8600:3282:68f0
  Options 0x33, Priority 128
  Prefix-count 0

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          10.0.0.4         0x80000002   692  0x3ebd  44
  fe80::20c:290f:fc3c:5b2f
  Options 0x33, Priority 128
  Prefix-count 0

    OSPF3 Link-Local database, interface xe-0/0/6.48 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.3          10.0.0.4         0x80000001   732  0x9f2d  44
  fe80::20c:2900:307d:2034
  Options 0x33, Priority 128
  Prefix-count 0
```

Verifying VRRP on vMX3:

```
jncie@vMX3> show vrrp detail

Physical interface: irb, Unit: 100, Address: 192.168.100.252/24
  Index: 354, SNMP ifIndex: 569, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: master, VRRP Mode: Active
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 192.168.100.254
  Advertisement Timer: 0.343s, Master router: 192.168.100.252
  Virtual router uptime: 00:07:47, Master router uptime: 00:07:42
  Virtual Mac: 00:00:5e:00:01:64
  Preferred: yes
  Tracking: disabled

Physical interface: irb, Unit: 100, Address: 2001:100:100:100::252/64
  Index: 354, SNMP ifIndex: 569, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: master, VRRP Mode: Active
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 2, VIP: fe80::100:1, 2001:100:100:100::254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Advertisement Timer: 0.134s, Master router: fe80::205:8600:644d:34f0
  Virtual router uptime: 00:07:47, Master router uptime: 00:07:42
  Virtual Mac: 00:00:5e:00:02:64
  Preferred: yes
  Tracking: disabled
```

Verifying VRRP on vMX4:

```
jncie@vMX4> show vrrp detail

Physical interface: irb, Unit: 100, Address: 192.168.100.253/24
  Index: 354, SNMP ifIndex: 570, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: backup, VRRP Mode: Active
  Priority: 50, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 192.168.100.254
  Dead timer: 3.484s, Master priority: 200, Master router: 192.168.100.252
  Virtual router uptime: 00:07:53
  Preferred: yes
  Tracking: disabled

Physical interface: irb, Unit: 100, Address: 2001:100:100:100::253/64
  Index: 354, SNMP ifIndex: 570, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: backup, VRRP Mode: Active
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 2, VIP: fe80::100:1, 2001:100:100:100::254
  Dead timer: 2.642s, Master priority: 200, Master router: fe80::205:8600:644d:34f0
  Virtual router uptime: 00:07:52
  Preferred: yes
  Tracking: disabled
```

Verifying IP-connectivity with the default gateway in DC2:

```
jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 192.168.100.254

PING 192.168.100.254 (192.168.100.254): 56 data bytes
!!!!!
--- 192.168.100.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 18.467/48.194/84.971/24.857 ms

jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 2001:100:100:100::254

PING6(56=40+8+8 bytes) 2001:100:100:100::1 --> 2001:100:100:100::254
!!!!!
--- 2001:100:100:100::254 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 34.693/84.347/141.557/43.018 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verifying IP-connectivity with the CORE lo0 addresses:

```
jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 10.0.0.5

PING 10.0.0.5 (10.0.0.5): 56 data bytes
!!!!!
--- 10.0.0.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 24.399/69.123/123.803/37.821 ms

jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 10.0.0.6
PING 10.0.0.6 (10.0.0.6): 56 data bytes
!!!!!
--- 10.0.0.6 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 63.400/131.196/317.492/93.859 ms

jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 2001:100::5

PING6(56=40+8+8 bytes) 2001:100:100:100::1 --> 2001:100::5
!!!!!
--- 2001:100::5 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 34.480/61.231/92.610/22.333 ms

jncie@vMX6> ping routing-instance DC2-S1-VL100 rapid 2001:100::6

PING6(56=40+8+8 bytes) 2001:100:100:100::1 --> 2001:100::6
!!!!!
--- 2001:100::6 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 26.109/27.809/33.407/2.805 ms
```

Task 1.13: MC-LAG Verification

- Perform ping from routing-instance DC1-S1-VL100 towards 192.168.100.1. Along the path of the packet, find:
    - all relevant destination MAC entries.
    - the ARP entries that exist as a result of this traffic.
    - the routing table entries required to get the packet from DC1 to DC2.
    - the interfaces that the traffic traverses.
- Perform a ping from routing-instance DC2-S3-VL100 towards 2001:100:0:100::1. Along the path of the packet, find:
    - all relevant destination MAC entries.
    - the NDP entries that exist as a result of this traffic.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- the routing table entries required to get the packet from DC1 to DC2.
- the interfaces that the traffic traverses.

Solution

Checking IPv4:

We start on vMX6, issue a ping and check the ARP entry for the default gateway:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 192.168.100.1 rapid

PING 192.168.100.1 (192.168.100.1): 56 data bytes
!!!!!
--- 192.168.100.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 44.808/51.576/63.984/6.680 ms

jncie@vMX6> show arp vpn DC1-S1-VL100

MAC Address       Address         Name                      Interface               Flags
00:00:00:00:11:22 192.168.1.254   192.168.1.254             xe-0/0/0.100            none
```

Next, we take note of the MAC address we used to source the ping:

```
jncie@vMX6> show interfaces xe-0/0/0 | match Current

  Current address: 00:0c:29:8c:00:80, Hardware address: 00:0c:29:17:77:db
```

We move to vQFX1 and verify the port on which the previously sent traffic arrived:

```
jncie@vQFX1> show ethernet-switching table 00:0c:29:8c:00:80
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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                 NH
RTR
    name                address             flags              interface               Index
ID
    vlan-100            00:0c:29:8c:00:80   DL              -   xe-0/0/1.0              0
0
```

Then we check where the packet is sent. We previously learned the MAC address the gateway is using. We use this information to check the switching table for the MAC address that the IPv4 gateway is configured with:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX1> show ethernet-switching table 00:00:00:00:11:22

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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags              interface              Index
ID
    vlan-100            00:00:00:00:11:22   DLR                -   ae0.0                  0
0
```

We move to the vMX and check the MAC table entries for VLAN 100:

```
jncie@vMX1> show bridge mac-table

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-100, VLAN : 100
   MAC                 MAC      Logical          NH     RTR
   addresssss          flags    interface        Index  ID
   00:0c:29:8c:00:80   DL       ae1.0
   00:0c:29:8c:00:8a   DLR      ae1.0


jncie@vMX2> show bridge mac-table

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-100, VLAN : 100
   MAC                 MAC      Logical          NH     RTR
   addresssss          flags    interface        Index  ID
   00:0c:29:8c:00:80   DR       ae1.0
   00:0c:29:8c:00:8a   DLR      ae1.0
```

The host we performed the ping from has MAC 00:0c:29:8c:00:80. We see that the vMX1 has learned this Locally (indicated by the MAC flags DL). vMX2 learned the MAC entry through ICCP (indicated by the MAC flags DR).

We do not see the destination MAC; this is because the vMX is acting as default gateway:

```
jncie@vMX2> show interfaces irb.100 | match MAC
  MAC: 00:00:00:00:11:22
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We can check the ARP entries on the vMX and see the entry that is related to the initial ICMP messages we send:

```
jncie@vMX1> show arp interface irb.100

MAC Address       Address         Name                    Interface               Flags
00:0c:29:8c:00:80 192.168.1.1     192.168.1.1             irb.100 [ae1.0]         none
00:0c:29:8c:00:8a 192.168.1.2     192.168.1.2             irb.100 [ae1.0]         none

jncie@vMX2> show arp interface irb.100

MAC Address       Address         Name                    Interface               Flags
00:0c:29:8c:00:80 192.168.1.1     192.168.1.1             irb.100 [ae1.0]         none
00:0c:29:8c:00:8a 192.168.1.2     192.168.1.2             irb.100 [ae1.0]         none
```

When the vMX1 router receives a packet with a destination MAC of the IRB interface, it examines the layer 3 destination IP address and it does a lookup in the routing-table:

```
jncie@vMX1> show route 192.168.100.1

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

192.168.100.0/24   *[OSPF/10] 00:24:48, metric 3
                    > to 192.168.46.2 via xe-0/0/6.46
```

We see the outgoing interface. In order to find out what node the packet is sent to next we check the OSPF neighbor:

```
jncie@vMX1> show ospf neighbor interface xe-0/0/6.46
Address          Interface              State     ID               Pri  Dead
192.168.46.2     xe-0/0/6.46            Full      10.0.0.5         128    35
```

The vMX1 router will forward the traffic out to CORE-1 so we move to vMX5 that has CORE-1 configured as a logical system. We can check where the traffic is coming from by examining on what interface CORE-1 neighbors vMX1:

```
jncie@vMX5> show ospf neighbor 10.0.0.1 logical-system CORE-1

Address          Interface              State     ID               Pri  Dead
192.168.46.1     xe-0/0/6.46            Full      10.0.0.1         128    32
```

Then we check where vMX5 will forward the packet to:

```
jncie@vMX5> show route logical-system CORE-1 192.168.100.1

inet.0: 20 destinations, 21 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
192.168.100.0/24    *[OSPF/10] 00:28:31, metric 2
                     > to 192.168.48.1 via xe-0/0/6.48
```

To see what router that is, we check the OSPF neighbor:

```
jncie@vMX5> show ospf neighbor logical-system CORE-1 interface xe-0/0/6.48

Address          Interface              State     ID               Pri  Dead
192.168.48.1     xe-0/0/6.48            Full      10.0.0.4         128    32
```

The neighbor across the interface is vMX4. because CORE-1 will forward the traffic to vMX4, we move to that router next. To check on what interface vMX4 receives the traffic:

```
jncie@vMX4> show ospf neighbor 10.0.0.5

Address          Interface              State     ID               Pri  Dead
192.168.48.2     xe-0/0/6.48            Full      10.0.0.5         128    31
```

To see where traffic is forwarded:

```
jncie@vMX4> show route 192.168.100.1

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

192.168.100.0/24    *[Direct/0] 00:13:47
                     > via irb.100
```

The host is reachable within the locally configured subnet. To check the destination MAC, we first check the corresponding ARP entry:

```
jncie@vMX4> show arp interface irb.100

MAC Address       Address         Name                Interface               Flags
00:05:86:4d:34:c2 192.168.100.1   192.168.100.1       irb.100 [ae1.0]         none
00:0c:29:8c:00:bc 192.168.100.3   192.168.100.3       irb.100 [ae3.0]         none
00:05:86:4d:34:f0 192.168.100.252 192.168.100.252     irb.100 [ae0.0]         permanent
remote
```

We can see the MAC address the destination host is using. Now we check the MAC table:

```
jncie@vMX4> show bridge mac-table 00:05:86:4d:34:c2

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-100, VLAN : 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
MAC                MAC      Logical          NH      RTR
addresssss         flags    interface        Index   ID
00:05:86:4d:34:c2  DL       ae1.0
```

The packet is forwarded out interface ae1. Next, we check the vQFX devices in DC2:

```
jncie@vQFX3> show ethernet-switching table 00:05:86:4d:34:c2

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 : 7 entries, 7 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags              interface              Index
ID
    vlan-100            00:05:86:4d:34:c2   D             -    ae1.0                  0
0
```

To see on what interface the traffic arrived, we check the entire MAC table for VLAN 100:

```
jncie@vQFX3> show ethernet-switching table vlan-id 100

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 : 7 entries, 7 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags              interface              Index
ID
    vlan-100            00:00:5e:00:01:64   D             -    ae0.0       0          0
    vlan-100            00:00:5e:00:02:64   D             -    ae0.0       0          0
    vlan-100            00:05:86:4d:34:c1   D             -    ae0.0       0          0
    vlan-100            00:05:86:4d:34:c2   D             -    ae1.0       0          0
    vlan-100            00:05:86:4d:34:f0   D             -    ae0.0       0          0
    vlan-100            00:05:86:82:68:f0   D             -    ae0.0       0          0
    vlan-100            00:0c:29:8c:00:bc   D             -    ae0.0       0          0
```

Now that we see what MAC addresses are learned inside VLAN 100, we check the ARP entry for the default gateway on vMX6:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> show arp vpn DC2-S1-VL100

MAC Address       Address         Name                      Interface               Flags
00:05:86:82:68:f0 192.168.100.253 192.168.100.253           ae3.100                 none
00:00:5e:00:01:64 192.168.100.254 192.168.100.254           ae3.100                 none
```

The interfaces that the traffic traverses are the following:

- vQFX1 receives traffic on xe-0/0/1 and forwards it across ae0
- vMX1 receives traffic on ae1 and forwards it across xe-0/0/6.46
- vMX5 receives traffic on xe-0/0/6.46 and forwards it across xe-0/0/6.48
- vMX4 receives traffic on xe-0/0/6.48 and forwards it out interface ae1
- vQFX3 receives traffic on ae0 and forwards traffic out interface ae1

Checking IPv6:

We start out on vMX6 and issue a ping:

```
jncie@vMX6> ping routing-instance DC2-S3-VL100 2001:100:0:100::1 rapid

PING6(56=40+8+8 bytes) 2001:100:100:100::3 --> 2001:100:0:100::1
!!!!!
--- 2001:100:0:100::1 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 44.904/54.380/78.775/12.767 ms
```

We check the link-layer address that the default gateway is using. We are sending out traffic from routing-instance DC2-S3-VL100.

This means we have to examine IPv6 neighbor entries for interface xe-0/0/6.100, which is the interface that is assigned to that routing-instance:

```
jncie@vMX6> show ipv6 neighbors | match xe-0/0/6.100

2001:100:100:100::252       00:05:86:4d:34:f0  stale       96  yes no      xe-0/0/6.100
2001:100:100:100::254       00:00:5e:00:02:64  stale       1142 yes no     xe-0/0/6.100
fe80::100:1                 00:00:5e:00:02:64  stale       1099 no no      xe-0/0/6.100
```

The gateway address in use is the 2001:100:100:100::254 address. The link-layer address that the gateway is using is 00:00:5e:00:02:64. Last thing is to check what link-layer address the ping is sourced from:

```
jncie@vMX6> show interfaces xe-0/0/6 | match Current

  Current address: 00:0c:29:8c:00:bc, Hardware address: 00:0c:29:17:77:17
```

The ping is sourced from the current address, or 00:0c:29:8c:00:bc.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Next, we move to vQFX5. Here we consult the MAC table and look for the source and destination MAC address that is used for the ping:

```
jncie@vQFX5> show ethernet-switching table 00:0c:29:8c:00:bc

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 : 8 entries, 8 learned
Routing instance : default-switch
  Vlan          MAC          MAC     Age  Logical          NH    RTR
  name          address        flags       interface        Index   ID
  vlan-100      00:0c:29:8c:00:bc  D       -  xe-0/0/1.0       0     0

jncie@vQFX5> show ethernet-switching table 00:00:5e:00:02:64

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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags             interface              Index
ID
    vlan-100            00:00:5e:00:02:64   D                 -   ae0.0                  0
0
```

Traffic enters the vQFX5 on interface xe-0/0/1 and exits the switch via interface ae0. The ae0 interface connects to the MC-LAG we created for a previous task.

Here we can see what device is at the other end of the link:

```
jncie@vQFX5> show lacp interfaces ae0 extensive | match Partner

      xe-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/2     Partner         127  03:04:03:04:00:03       127   32771     1
      xe-0/0/4     Partner         127  03:04:03:04:00:03       127       5     1
```

During the next step , we check the MAC table on vMX3 and vMX4:

```
jncie@vMX3> show bridge mac-table 00:0c:29:8c:00:bc

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-100, VLAN : 100
   MAC                MAC      Logical          NH     RTR
   addresssss         flags    interface        Index  ID
   00:0c:29:8c:00:bc  DR       ae3.0
```

```
jncie@vMX4> show bridge mac-table 00:0c:29:8c:00:bc

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : bd-100, VLAN : 100
   MAC                MAC      Logical          NH     RTR
   addresssss         flags    interface        Index  ID
   00:0c:29:8c:00:bc  DL       ae3.0
```

On vMX3, we can see that the MAC address was learned remotely, indicating that the first packet went through vMX4.

We can find a neighbor entry for the host we sourced the ping from on both vMX routers:

```
jncie@vMX3> show ipv6 neighbors | match 00:0c:29:8c:00:bc

2001:100:100:100::3          00:0c:29:8c:00:bc  stale       1054 yes no     irb.100 [ae3.0]

jncie@vMX4> show ipv6 neighbors | match 00:0c:29:8c:00:bc

2001:100:100:100::3          00:0c:29:8c:00:bc  stale       1072 yes no     irb.100 [ae3.0]
```

vMX4 is acting as the VRRP master:

```
jncie@vMX4> show vrrp
..
irb.100        up            100   master   Active      A  0.508 lcl    2001:100:100:100::253
                                                                 vip    fe80::100:1
                                                                 vip    2001:100:100:100::254
```

When this vMX receives traffic destined to the gateway address, it will do a routing lookup and forward the packet based on the destination IP address:

```
jncie@vMX4> show route 2001:100:0:100::1

inet6.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100:0:100::/64*[OSPF3/10] 00:16:52, metric 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
> to fe80::20c:2900:30b9:e97d via xe-0/0/6.48
```

We now have the incoming interface on vMX4 (ae3) as well as the outgoing interface (xe-0/0/6.48).

To check what the next router in the path is, we check what OSPFv3 neighbor is learned across that interface:

```
jncie@vMX4> show ospf3 neighbor interface xe-0/0/6.48

ID              Interface             State     Pri   Dead
10.0.0.5        xe-0/0/6.48           Full      128     37
  Neighbor-address fe80::20c:2900:30b9:e97d
```

CORE-1 is a logical system that is configured on the vMX5 router. We move over to the vMX5 router and check what the OSPFv3 neighbors are as well as where the vMX5 router forwards the packets to:

```
jncie@vMX5> show route logical-system CORE-1 2001:100:0:100::1

inet6.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100:0:100::/64*[OSPF3/10] 02:32:27, metric 2
                    > to fe80::20c:2900:2e68:6e2e via xe-0/0/6.46

jncie@vMX5> show ospf3 neighbor logical-system CORE-1

ID              Interface             State     Pri   Dead
10.0.0.6        lt-0/0/0.0            Full      128     33
  Neighbor-address fe80::20c:2900:0:159
10.0.0.1        xe-0/0/6.46           Full      128     34
  Neighbor-address fe80::20c:2900:2e68:6e2e
10.0.0.4        xe-0/0/6.48           Full      128     39
  Neighbor-address fe80::20c:2900:307d:2034
```

Packets send from the vMX4 router are received across the xe-0/0/6.48 interface.

The routing lookup indicates that packets will be forwarded through interface xe-0/0/6.46. The OSPFv3 neighbor is vMX1, so this is our next stop.

On vMX1, we first check to see on what interface we find the vMX5 router as an OSPFv3 neighbor:

```
jncie@vMX1> show ospf3 neighbor 10.0.0.5

ID              Interface             State     Pri   Dead
10.0.0.5        xe-0/0/6.46           Full      128     34
  Neighbor-address fe80::20c:2900:2eb9:e97d
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Then we examine the routing table:

```
jncie@vMX1> show route 2001:100:0:100::1

inet6.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100:0:100::/64*[Direct/0] 02:59:57
                    > via irb.100
```

The destination host resides within a locally configured subnet. We check the IPv6 neighbors:

```
jncie@vMX1> show ipv6 neighbors | match 2001:100:0:100::1

2001:100:0:100::1            00:0c:29:8c:00:80  reachable   31  yes no      irb.100 [ae1.0]
```

We now have both incoming as well as the outgoing interface on vMX1. Before moving on, we take note of the IRB MAC address that the vMX is using:

```
jncie@vMX1> show interfaces irb.100 | match MAC

    MAC: 00:00:00:00:11:22
```

Now we check vQFX1 and vQFX2 for the host MAC as well as the gateway MAC:

```
jncie@vQFX1> show ethernet-switching table 00:00:00:00:11:22

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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                 NH
RTR
    name                address             flags              interface               Index
ID
    vlan-100            00:00:00:00:11:22   DL              -   ae0.0                   0
0

jncie@vQFX1> show ethernet-switching table 00:0c:29:8c:00:80

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)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Ethernet switching table : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags               interface              Index
ID
    vlan-100            00:0c:29:8c:00:80   DL              -   xe-0/0/1.0             0
0
```

```
jncie@vQFX2> show ethernet-switching table 00:00:00:00:11:22

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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags               interface              Index
ID
    vlan-100            00:00:00:00:11:22   DR              -   ae0.0                  0
0
```

```
jncie@vQFX2> show ethernet-switching table 00:0c:29:8c:00:80

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 : 8 entries, 8 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags               interface              Index
ID
    vlan-100            00:0c:29:8c:00:80   DR              -   ae1.0                  0
0
```

We can see that vQFX1 learned both gateway as well as host MAC locally. The gateway is located behind interface ae0 and the host is located behind interface xe-0/0/1.

The interfaces that the traffic traverses are the following:

- vQFX5 receives traffic on xe-0/0/1 and forwards it across ae0
- vMX4 receives traffic on ae3 and forwards it across xe-0/0/6.48

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- vMX5 receives traffic on xe-0/0/6.48 and forwards it across xe-0/0/6.46
- vMX1 receives traffic on xe-0/0/6.46 and forwards it across ae0
- vQFX1 receives traffic on ae0 and forwards it across interface xe-0/0/1

Part 3: VCF

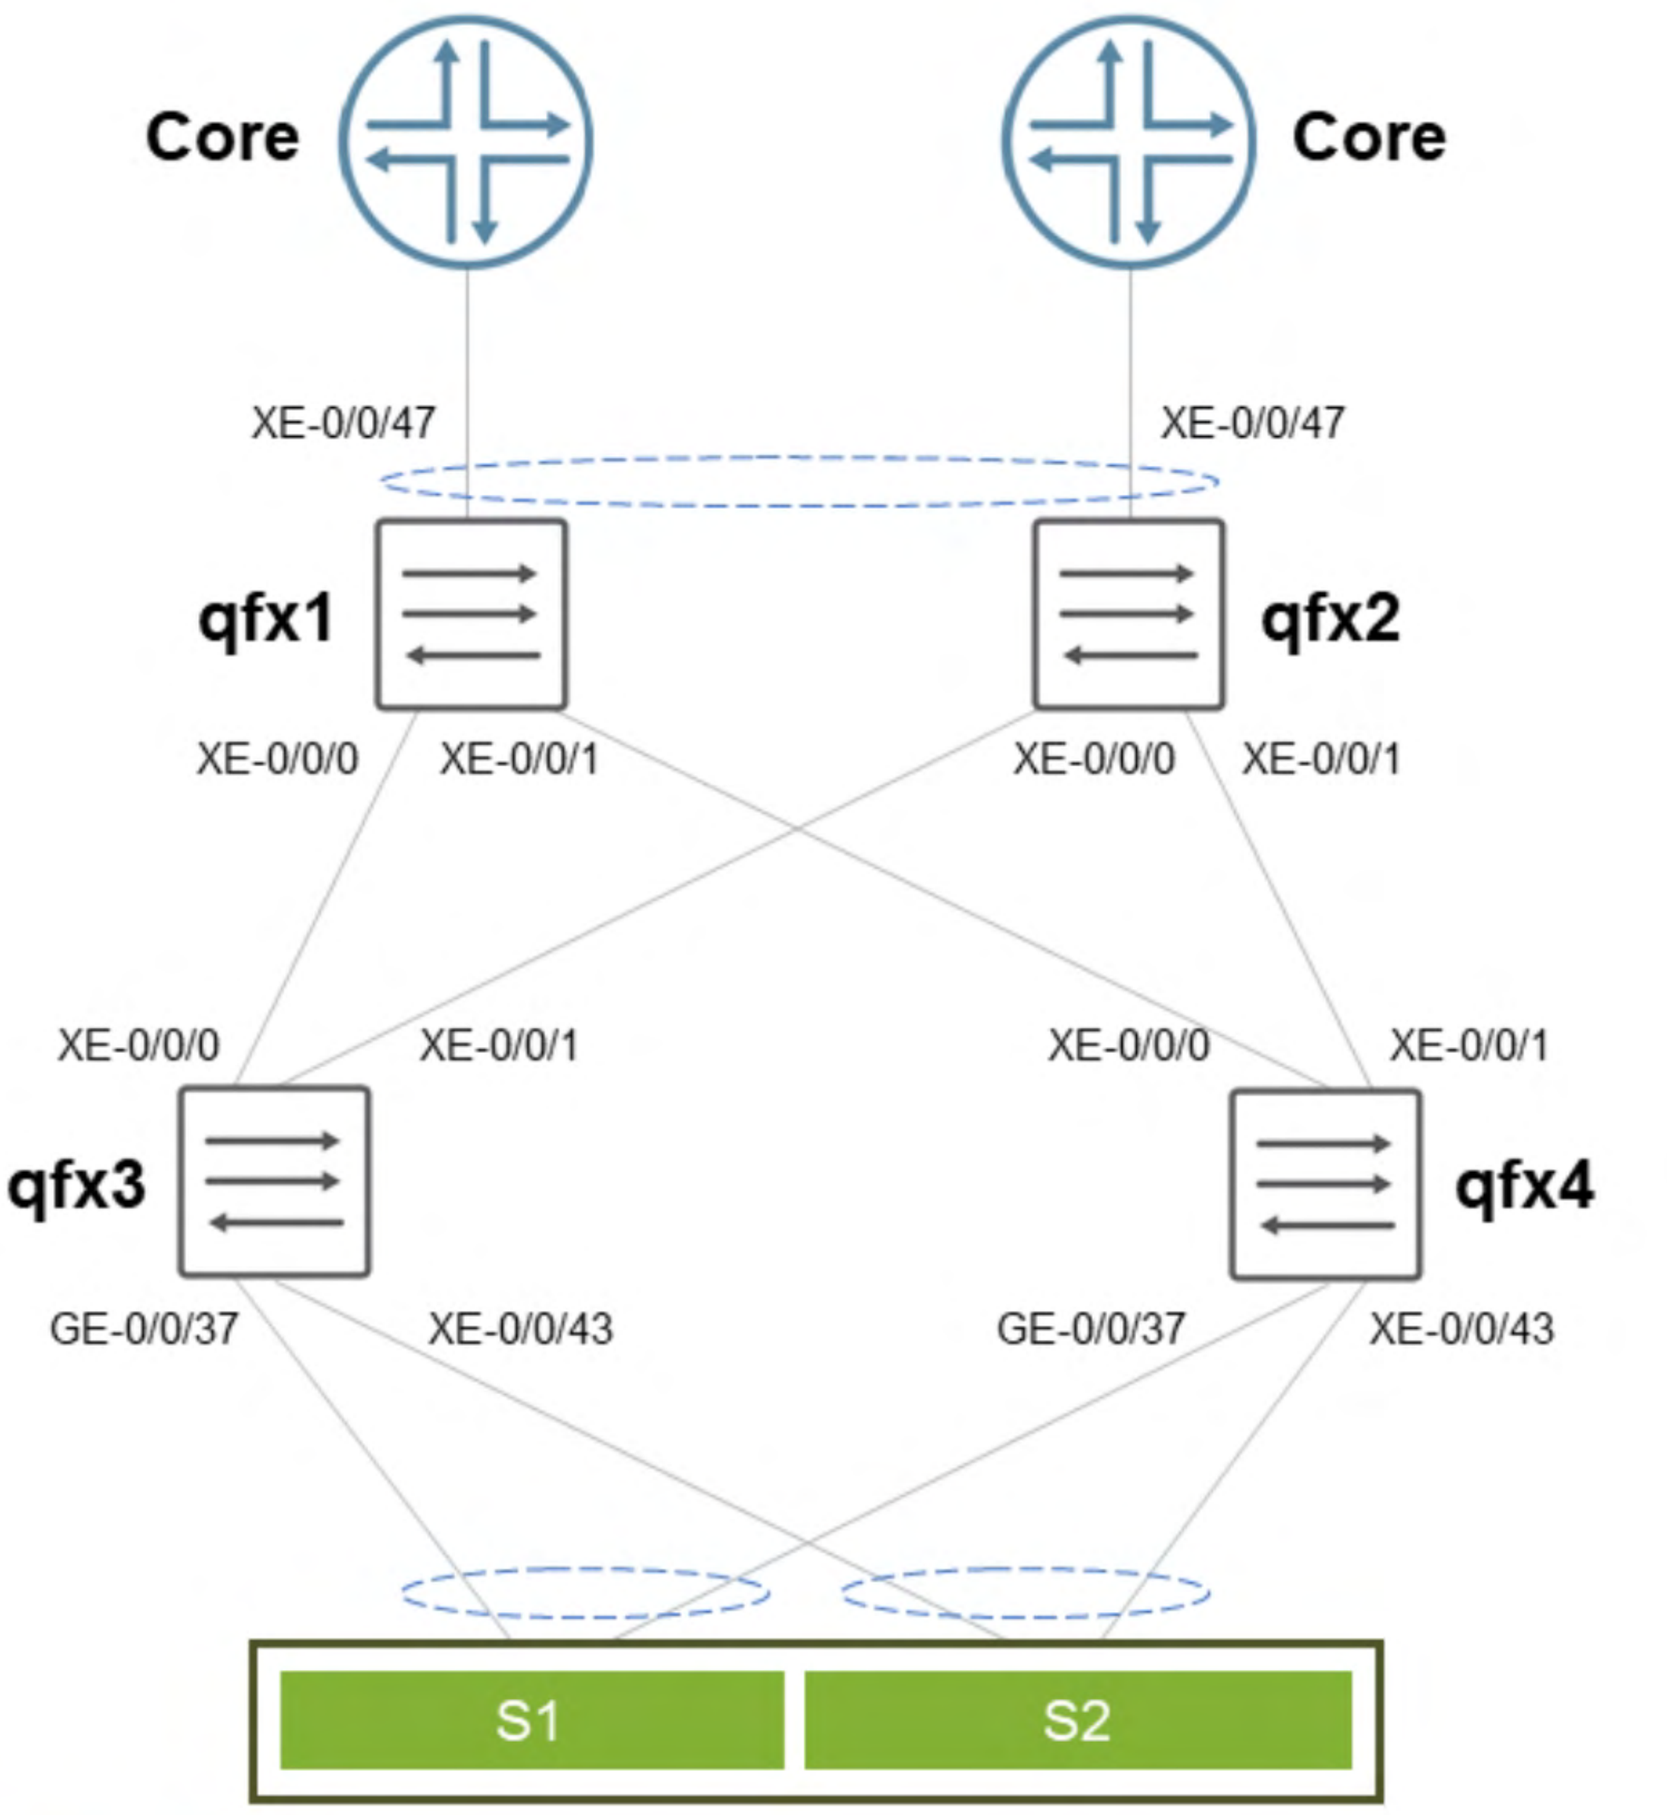

Figure 16 - Physical topology

Notes: This assignment is not included in the Juniper Networks rack rental topology. The rack rental only includes vMX and vQFX devices and unfortunately, vQFX does not support VCF.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Console access to every switch is required for these assignments.

Revert the QFX switches to factory default before you start this lab.

Task 1.14: VCF Configuration

- Configure the QFX switches to form a single VCF with the hostname 'jncie-vcf'.
- QFX1 should be the active master switch and QFX2 should be the backup switch.
- Configuration must determine the role of every device.

Solution

Since the configuration should determine the role of every device, we can conclude that this task requires us to configure a preprovisioned VCF.

Before we start, we check the virtual-chassis mode on every device:

qfx1:

```
Amnesiac (ttyd0)

login: root

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@:RE:0%
root@:RE:0% cli

{master:0}
root> show virtual-chassis

Virtual Chassis ID: 84f0.8d18.8285
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    VF3715490286 qfx5100-48s-6q 128   Master*      N  VC

Member ID for next new member: 1 (FPC 1)
```

qfx2:

```
root> show virtual-chassis

Virtual Chassis ID: b79d.5920.c1f8
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    VB3715280174 qfx5100-96s-8q 128   Master*      N  VC

Member ID for next new member: 1 (FPC 1)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

qfx3:

```
root> show virtual-chassis

Virtual Chassis ID: 660a.6a5f.4421
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    VF3715510230 qfx5100-48s-6q 128   Master*      N  VC
Member ID for next new member: 1 (FPC 1)
```

qfx4:

```
root> show virtual-chassis

Virtual Chassis ID: 2d31.85ec.1ed1
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    VF3714370037 qfx5100-48s-6q 128   Master*      N  VC

Member ID for next new member: 1 (FPC 1)
```

None of the switches are in Fabric mode. Additionally, the command output gave us the serial number for every switch. We note write these down for later use:

Switch:	**Serial number:**
qfx1	VF3715490286
qfx2	VB3715280174
qfx3	VF3715510230
qfx4	VF3714370037`

Note: there are alternate ways to get the device serial number. You can check the actual physical switch, or you can issue the following command:

```
root> show chassis hardware | match Chas

Chassis                             VF3715490286      QFX5100-48S-6Q
```

Previously, the ‘show virtual-chassis’ command output also informed us about the current mode that the devices are in. The ‘Route Mode’ VC indicates that every device is ready to form a virtual-chassis.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

At this stage, no ports are acting as vc-ports:

```
root> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------

{master:0}
root>
```

The steps to turn these switches into a single VCF are as follows:

1. Configure qfx1 as a VCF.
2. Enabling interfaces between the QFX vc-ports on every switch.
3. Reboot the switches into the proper mode.
4. Verify the VCF.

First, we configure qfx1 as a VCF:

qfx1:

```
set system host-name jncie-vcf
set system root-authentication encrypted-password "$1$GhBwtb0f$oAYgCJIH8qJka8JcJmy9y0"
set virtual-chassis preprovisioned
set virtual-chassis member 0 role routing-engine
set virtual-chassis member 0 serial-number VF3715490286
set virtual-chassis member 1 role routing-engine
set virtual-chassis member 1 serial-number VB3715280174
set virtual-chassis member 2 role line-card
set virtual-chassis member 2 serial-number VF3715510230
set virtual-chassis member 3 role line-card
set virtual-chassis member 3 serial-number VF3714370037
```

Next, we enable the vc-ports interconnecting the QFX switches on every device.

every QFX:

```
request virtual-chassis vc-port set pic-slot 0 port 0
request virtual-chassis vc-port set pic-slot 0 port 1
```

After enabling the vc-ports, the VCF should form:

```
root@jncie-vcf> show virtual-chassis

Preprovisioned Virtual Chassis
Virtual Chassis ID: 84f0.8d18.8285
Virtual Chassis Mode: Enabled
                                                  Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model             prio  Role      Mode  Mode ID  Interface
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      N  VC   3  vcp-255/0/1
                                                                          2  vcp-255/0/0
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q 129   Backup       N  VC   2  vcp-255/0/0
                                                                          3  vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q   0   Linecard     N  VC   0  vcp-255/0/0
                                                                          1  vcp-255/0/1
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q   0   Linecard     N  VC   0  vcp-255/0/0
                                                                          1  vcp-255/0/1
```

From the command output, we can see that four switches have formed a preprovisioned Virtual Chassis.

To turn this Virtual Chassis into a Virtual Chassis Fabric, we need to reboot the switches into fabric mode. We can do this by issuing the following command from the master switch:

```
root@jncie-vcf> request virtual-chassis mode fabric reboot all-members

fpc1:
--------------------------------------------------------------------------------
Mode set to 'Fabric with similar devices'.  Rebooting system...

fpc2:
--------------------------------------------------------------------------------
Mode set to 'Fabric with similar devices'.  Rebooting system...

fpc3:
--------------------------------------------------------------------------------
Mode set to 'Fabric with similar devices'.  Rebooting system...

fpc0:
--------------------------------------------------------------------------------
Mode set to 'Fabric with similar devices'.  Rebooting system...

*** System shutdown message from root@jncie-vcf ***

System going down in 1 minute

root@jncie-vcf>
```

After a reboot, we can start verifying the VCF. First, we verify that the VCF is functioning properly:

```
root@jncie-vcf> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 84f0.8d18.8285
Fabric Mode: Enabled

                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model           prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      N  F    2  vcp-255/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                                                                    3  vcp-255/0/1
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q 129   Backup      N  F  2  vcp-255/0/0
                                                                    3  vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q   0   Linecard    N  F  0  vcp-255/0/0
                                                                    1  vcp-255/0/1
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q   0   Linecard    N  F  0  vcp-255/0/0
                                                                    1  vcp-255/0/1

root@jncie-vcf>
```

We now have a preprovisioned Virtual Chassis Fabric in which every device is enabled for Fabric mode.

Task 1.15: VCF Features

- When the master fails, the backup should continue forwarding layer 3 as well as layer 2 traffic immediately.
- The MDT should be rooted at the spine node.
- The root user should be able to manage the VCF on 10.10.20.120/24 using SSH.

Solution

To complete the task, we need to configure several redundancy mechanisms. These are the following:

- graceful routing-engine switchover (GRES)
- non-stop routing (NSR)
- non-stop bridging (NSB)

Enabling GRES:

```
set system commit synchronize
set chassis redundancy graceful-switchover
```

Enabling NSR:

```
set routing-options nonstop-routing
```

Enabling NSB:

```
set protocols layer2-control nonstop-bridging
```

To verify the configuration of the redundancy mechanisms, we need to issue several commands from the backup routing-engine.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Console and SSH connections will always be redirected to the master routing-engine. Since we need to verify everything from the backup routing-engine, we need to log in to the backup first:

```
root@jncie-vcf> request session member 1

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@jncie-vcf:BK:1%
root@jncie-vcf:BK:1% cli
warning: This chassis is operating in a non-master role as part of a virtual-chassis fabric
(VCF) system.
warning: Use of interactive commands should be limited to debugging and VC Port operations.
warning: Full CLI access is provided by the Virtual Chassis Fabric Master (VCF-M) chassis.
warning: The VCF-M can be identified through the show virtual-chassis status command executed
at this console.
warning: Please logout and log into the VCF-M to use CLI.
{backup:1}
root@jncie-vcf>
```

When we are connected to the backup, we start out verifying GRES:

```
root@jncie-vcf> show system switchover

fpc1:
--------------------------------------------------------------------------
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State

{backup:1}
root@jncie-vcf>
```

Even though we have not configured any VLANs or routing on the VCF, we can still verify that NSR and NSB are active. We simply verify whether or not the relevant daemons are running:

```
root@jncie-vcf> show system processes extensive | match "rpd|l2cpd"

4244 root        1   4    0 16536K  8612K kqread   0:00  0.00% l2cpd
 4245 root        1   4    0 57264K 12584K kqread   0:00  0.00% rpd
```

Without NSR, the RPD (routing protocol daemon) would not be running on the backup routing-engine. Without NSB, the l2cpd (layer 2 control protocol daemon) would not be running on the backup routing-engine.

Next is rooting the MDT at the spines:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set virtual-chassis member 0 fabric-tree-root
set virtual-chassis member 1 fabric-tree-root
```

And finally, we enable the root user to login using SSH:

```
set system services ssh root-login allow
set system services ssh connection-limit 2
set system services ssh rate-limit 10

set interfaces vme unit 0 family inet address 10.10.20.120/24
```

Task 1.16: Deconstructing the VCF

- Deconstruct the VCF.
- Return all switches to the default configuration and mode.
- No VC or VCF should be able to form.

Solution

To deconstruct the VCF, we disable the fabric mode from the master.

```
root@jncie-vcf> request virtual-chassis mode fabric disable all-members

fpc1:
--------------------------------------------------------------------------
Mode set to 'Virtual Chassis with similar devices'. (Reboot required)

fpc2:
--------------------------------------------------------------------------
Mode set to 'Virtual Chassis with similar devices'. (Reboot required)

fpc3:
--------------------------------------------------------------------------
Mode set to 'Virtual Chassis with similar devices'. (Reboot required)

fpc0:
--------------------------------------------------------------------------
Mode set to 'Virtual Chassis with similar devices'. (Reboot required)

{master:0}
root@jncie-vcf>
```

After that, we issue a ‘request system zeroize’ on every switch. If you zeroize qfx1 and qfx2 first, logging into qfx3 and qfx4 will take a little while longer. This is because switches are operating in linecard mode will first try to connect to a master. Only after a timeout will the device allow you to login locally. Because of this, the quickest way to ‘reset’ the VCF is to zeroize the leaf switches first.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

qfx4:

```
Wed Jul 12 09:38:23 UTC 2017

Amnesiac (ttyd0)

login: root

Logging to master
Password:

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0% cli

root@jncie-vcf>

root@jncie-vcf> request session member 3

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@jncie-vcf:LC:3%
root@jncie-vcf:LC:3% cli
warning: This chassis is operating in a non-master role as part of a virtual-chassis fabric
(VCF) system.
warning: Use of interactive commands should be limited to debugging and VC Port operations.
warning: Full CLI access is provided by the Virtual Chassis Fabric Master (VCF-M) chassis.
warning: The VCF-M can be identified through the show virtual-chassis status command executed
at this console.
warning: Please logout and log into the VCF-M to use CLI.

{linecard:3}
root@jncie-vcf> request system zeroize

warning: System is part of a Virtual Chassis. This operation might impact traffic on the
Virtual Chassis.
Erase all data, including configuration and log files? [yes,no] (no) yes

Terminated
```

qfx3:

```
Wed Jul 12 09:38:23 UTC 2017

Amnesiac (ttyd0)

login: root

Logging to master
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Password:

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0%
root@jncie-vcf:RE:0% cli

root@jncie-vcf>

root@jncie-vcf> request session member 2

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC

root@jncie-vcf:LC:2%

root@jncie-vcf:LC:2% cli
warning: This chassis is operating in a non-master role as part of a virtual-chassis fabric
(VCF) system.
warning: Use of interactive commands should be limited to debugging and VC Port operations.
warning: Full CLI access is provided by the Virtual Chassis Fabric Master (VCF-M) chassis.
warning: The VCF-M can be identified through the show virtual-chassis status command executed
at this console.
warning: Please logout and log into the VCF-M to use CLI.

{linecard:2}
root@jncie-vcf> request system zeroize

warning: System is part of a Virtual Chassis. This operation might impact traffic on the
Virtual Chassis.
Erase all data, including configuration and log files? [yes,no] (no) yes

Terminated
```

qfx2:

```
{master:1}
root@jncie-vcf> request system zeroized
```

qfx1:

```
{linecard:0}
root@jncie-vcf> request system zeroize
```

Note: ‘request system zeroize’ will also restore the default port-settings. This means that after the reboot, the QFX5100 will no longer list any vc-ports.

After completing every action, you should see the following on every device:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Amnesiac (ttyd0)

login: root

--- JUNOS 14.1X53-D40.8 built 2016-11-09 02:13:22 UTC
root@:RE:0% cli


root> show virtual-chassis

Virtual Chassis ID: e992.5747.52f3
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    VF3715490286 qfx5100-48s-6q 128   Master*      N  VC

Member ID for next new member: 1 (FPC 1)

{master:0}
root> show virtual-chassis vc-port
fpc0:
--------------------------------------------------------------------------

{master:0}
root> show virtual-chassis mode
fpc0:
--------------------------------------------------------------------------
Current mode : Virtual Chassis with similar devices
Future mode after reboot : Virtual Chassis with similar devices
```

Task 1.17: Autoprovisioning a VCF

- Configure an autoprovisioned VCF with hostname ‘jncie-vcf’
- The qfx1 and qfx2 devices should operate as spines.
- Enable the VCF for future additions of EX4300 switches.
- The root user should be able to connect to the device on 10.10.20.120/24 using SSH.
- Enable GRES, NSB and NSR.

Solution

We first put the QFX into the proper mode. We examine every device using the following command:

```
root> show virtual-chassis

Virtual Chassis ID: df51.ce94.ca64
Virtual Chassis Mode: Enabled
                                                    Mstr          Mixed Route Neighbor List
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 128   Master*      N  VC

Member ID for next new member: 1 (FPC 1)
```

This device is enabled for VC and has mixed mode disabled. We need to enable the switches for Fabric mode and mixed mode.

We issue the following command on every device that is not enabled for a mixed fabric:

```
request virtual-chassis mode fabric mixed reboot
```

Before continuing, every switch should display the following output when the 'show virtual-chassis' command is entered:

```
root> show virtual-chassis

Fabric ID: df51.ce94.ca64
Fabric Mode: Enabled
                                                Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 128   Master*      Y  F

Member ID for next new member: 1 (FPC 1)
```

After this, we gather the serial numbers for both spine switches. The serial numbers can be obtained by issuing the 'show virtual-chassis' command or by issuing the following command:

```
root> show chassis hardware | match Chassis

Chassis                                 VF3715490286      QFX5100-48S-6Q
```

After setting the correct mode on every device, we configure the qfx1 switch as a VCF:

qfx1:

```
set system host-name jncie-vcf
set system root-authentication plain-text-password
set system services ssh root-login allow
delete interfaces vme
set interfaces vme unit 0 family inet address 10.10.20.120/24

set system commit synchronize
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set protocols layer2-control nonstop-bridging
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set virtual-chassis auto-provisioned
set virtual-chassis member 0 role routing-engine
set virtual-chassis member 0 serial-number VF3715490286
set virtual-chassis member 1 role routing-engine
set virtual-chassis member 1 serial-number VB3715280174
```

Optionally, to remove any periodic 'clutter' from our terminal screen, we can also remove the following configuration:

```
delete interfaces irb unit 0
delete chassis auto-image-upgrade
```

After configuring qfx1, we wait until the VCF forms. This will take somewhere between 3 and 5 minutes:

```
root# commit and-quit

configuration check succeeds
se      cli     commit complete
Exiting configuration mode

{master:0}
root@jncie-vcf> set cli timestamp

Jul 12 10:04:06
CLI timestamp set to: %b %d %T

{master:0}
root@jncie-vcf> show virtual-chassis | refresh 5

Jul 12 10:04:07
---(refreshed at 2017-07-12 10:04:08 UTC)---

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed
                                                Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      Y  F
---(refreshed at 2017-07-12 10:04:13 UTC)---

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed
                                                Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      Y  F

<output omitted>
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
---(refreshed at 2017-07-12 10:07:19 UTC)---

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed

                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model           prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q  129   Master*     Y  F    2  vcp-255/0/0
                                                                          3  vcp-255/0/1
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q  128   Backup      Y  F    3  vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q    0   Linecard    Y  F    0  vcp-255/0/0
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q    0   Linecard    Y  F    0  vcp-255/0/0
---(*more)---[abort]
```

The end-result should be as follows:

```
root@jncie-vcf> show virtual-chassis

Jul 12 10:14:23

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed

                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model           prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q  129   Master*     Y  F    2  vcp-255/0/0
                                                                          3  vcp-255/0/1
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q  129   Backup      Y  F    2  vcp-255/0/0
                                                                          3  vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q    0   Linecard    Y  F    0  vcp-255/0/0
                                                                          1  vcp-255/0/1
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q    0   Linecard    Y  F    0  vcp-255/0/0
                                                                          1  vcp-255/0/1

root@jncie-vcf> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type                Trunk  Status       Speed        Neighbor
or                               ID                (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured     -1     Up           10000        2   vcp-255/0/0
0/1         Auto-Configured     -1     Up           10000        3   vcp-255/0/0

fpc1:
--------------------------------------------------------------------------
Interface   Type                Trunk  Status       Speed        Neighbor
or                               ID                (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured     -1     Up           10000        2   vcp-255/0/1
0/1         Auto-Configured     -1     Up           10000        3   vcp-255/0/1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
fpc2:
--------------------------------------------------------------------------
Interface   Type                Trunk   Status       Speed        Neighbor
or                               ID                  (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured     -1      Up           10000        0   vcp-255/0/0
0/1         Auto-Configured     -1      Up           10000        1   vcp-255/0/0

fpc3:
--------------------------------------------------------------------------
Interface   Type                Trunk   Status       Speed        Neighbor
or                               ID                  (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured     -1      Up           10000        0   vcp-255/0/1
0/1         Auto-Configured     -1      Up           10000        1   vcp-255/0/1
```

Task 1.18: VCF VLAN and Interface Configuration

- Enable all core and server-links for VLAN 20 and VLAN 30.
- VLAN 20 should be the native VLAN on every link.
- Ensure that all current as well as future VLANs are limited to 500 MAC addresses. The packets that are violating this setting should be dropped.
- None of the VLANs should not allow for MAC flapping to exceed 2 per second. Packets that are violating this setting should be dropped.
- The default VLAN should not be allowed to forward any packets.
- Use the following interface numbering:

S1	ae18
S2	ae19
core	ae20

Solution

VCF:

```
set chassis aggregated-devices ethernet device-count 21

set interfaces xe-0/0/47 ether-options 802.3ad ae20
set interfaces xe-1/0/47 ether-options 802.3ad ae20

set interfaces ge-2/0/37 ether-options 802.3ad ae18
set interfaces ge-3/0/37 ether-options 802.3ad ae18

set interfaces xe-2/0/43 ether-options 802.3ad ae19
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces xe-3/0/43 ether-options 802.3ad ae19

set interfaces ae18 native-vlan-id 20
set interfaces ae18 aggregated-ether-options lacp active
set interfaces ae18 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae18 unit 0 family ethernet-switching vlan members 20
set interfaces ae18 unit 0 family ethernet-switching vlan members 30

copy interfaces ae18 to ae19

copy interfaces ae18 to ae20

set firewall family ethernet-switching filter reject-all term 1 then discard

set vlans default forwarding-options filter input reject-all

set vlans vlan-20 vlan-id 20
set vlans vlan-30 vlan-id 30

set groups vlans vlans <*> switch-options mac-table-size 500
set groups vlans vlans <*> switch-options mac-table-size packet-action drop
set groups vlans vlans <*> switch-options interface-mac-limit 500
set groups vlans vlans <*> switch-options interface-mac-limit packet-action none
set groups vlans vlans <*> switch-options mac-move-limit 2
set groups vlans vlans <*> switch-options mac-move-limit packet-action drop
```

After completing the configuration, we start out verifying the AE interfaces:

```
root@jncie-vcf> show interfaces terse | match ae

xe-0/0/47.0             up    up   aenet    --> ae20.0
xe-1/0/47.0             up    up   aenet    --> ae20.0
ge-2/0/37.0             up    up   aenet    --> ae18.0
xe-2/0/43.0             up    up   aenet    --> ae19.0
ge-3/0/37.0             up    up   aenet    --> ae18.0
xe-3/0/43.0             up    up   aenet    --> ae19.0
...
ae18                    up    up
ae18.0                  up    up   eth-switch
ae19                    up    up
ae19.0                  up    up   eth-switch
ae20                    up    up
ae20.0                  up    up   eth-switch
```

This tells up the interfaces are up. To gather additional information on the AE interface, we issue the following command:

```
root@jncie-vcf> show lacp interfaces
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Aggregated interface: ae18
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/37      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/37    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/37      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/37    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/37               Current   Fast periodic Collecting distributing
      ge-3/0/37               Current   Fast periodic Collecting distributing

Aggregated interface: ae19
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-3/0/43      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/0/43    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-2/0/43      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-2/0/43    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-3/0/43               Current   Fast periodic Collecting distributing
      xe-2/0/43               Current   Fast periodic Collecting distributing

Aggregated interface: ae20
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/47      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/47    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/47      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/47    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/47               Current   Fast periodic Collecting distributing
      xe-1/0/47               Current   Fast periodic Collecting distributing
```

We can see all AE interfaces are up and devices on both side of the link are configured to actively send out LACP PDUs. The ‘Receive State’ is ‘Current’. This means that LACP PDUs are being received across that particular link.

On an MX, we have the ‘show lacp interface extensive’ or ‘show lacp statistics interfaces’ commands at our disposal. Even though those commands are not available on a VCF, we are able to see the same information when we issue the following command:

```
root@jncie-vcf> show interfaces ae20 extensive | no-more

Physical interface: ae20, Enabled, Physical link is Up
  Interface index: 698, SNMP ifIndex: 593, Generation: 201
  Link-level type: Ethernet, MTU: 1514, Speed: 20Gbps, BPDU 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: 80:ac:ac:a4:25:04, Hardware address: 80:ac:ac:a4:25:04
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Last flapped   : 2017-07-12 11:25:13 UTC (00:52:11 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :              1318628                 3208 bps
 Output bytes  :              1550137                 3200 bps
 Input  packets:                10709                    3 pps
 Output packets:                12380                    3 pps
 IPv6 transit statistics:
  Input  bytes  :                    0
  Output bytes  :                    0
  Input  packets:                    0
  Output packets:                    0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Bucket drops: 0, Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0, Bucket drops: 0
Egress queues: 12 supported, 5 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                                0                 2095                    0
  3                                0                    0                    0
  4                                0                    0                    0
  7                                0                 9715                    0
  8                                0                  177                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  3                   fcoe
  4                   no-loss
  7                   network-control
  8                   mcast

Logical interface ae20.0 (Index 572) (SNMP ifIndex 596) (Generation 168)
  Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Adaptive Statistics:
      Adaptive Adjusts:          0
      Adaptive Scans  :          0
      Adaptive Updates:          0
  Link:
    xe-0/0/47.0
      Input :             0          0             0            0
      Output:           114          0         39784            0
    xe-1/0/47.0
      Input :             0          0             0            0
      Output:           115          0         40132            0
  LACP info:        Role     System           System      Port    Port  Port
                             priority       identifier  priority  number   key
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  xe-0/0/47.0    Actor         127  80:ac:ac:a4:24:00       127        1    21
  xe-0/0/47.0  Partner         127  64:64:9b:dd:64:b0       127        4     2
  xe-1/0/47.0    Actor         127  80:ac:ac:a4:24:00       127        2    21
  xe-1/0/47.0  Partner         127  64:64:9b:dd:64:b0       127        3     2
LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
  xe-0/0/47.0             3131        3116            0            0
  xe-1/0/47.0             3132        3116            0            0
Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
  xe-0/0/47.0                0           0            0            0
  xe-1/0/47.0                0           0            0            0
Protocol eth-switch, MTU: 1514, Generation: 185, Route table: 6
  Flags: Trunk-Mode
```

```
root@jncie-vcf>
```

The previous output informed us about the fact that link is up and operating at an aggregate speed of 20Gbit. In addition to this, we also saw the 'LACP info' and the 'LACP statistics' section. In these sections we can see the local system identifier used in LACP packets by the Actor (local device), the system identifier used by the other device (partner), the LACP key in use and the amount of LACP PDUs send and received on the physical interface.

After we have verified the interface configuration, we can continue to verify the VLAN configuration.

We used an apply-group to ensure that both current as well as future VLANs inherit the same MAC-table and mac-movement restrictions. To verify that our apply-group is working, we issue the following command:

```
root@jncie-vcf> show configuration vlans | display inheritance | no-more

...
vlan-30 {
    vlan-id 30;
    ##
    ## 'switch-options' was inherited from group 'vlans'
    ##
    switch-options {
        ##
        ## 'mac-table-size' was inherited from group 'vlans'
        ##
        mac-table-size {
            ##
            ## '500' was inherited from group 'vlans'
            ##
            500;
            ##
            ## 'drop' was inherited from group 'vlans'
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
        ##
        packet-action drop;
    }
    ##
    ## 'interface-mac-limit' was inherited from group 'vlans'
    ##
    interface-mac-limit {
        ##
        ## '500' was inherited from group 'vlans'
        ##
        500;
        ##
        ## 'drop-and-log' was inherited from group 'vlans'
        ##
        packet-action drop-and-log;
    }
    ##
    ## 'mac-move-limit' was inherited from group 'vlans'
    ## '2' was inherited from group 'vlans'
    ## 'drop' was inherited from group 'vlans'
    ##
    mac-move-limit 2 packet-action drop;
  }
}
```

We now proceed with the verification of the configured VLANs. The following command display the VLANs that are currently active on the VCF:

```
root@jncie-vcf> show vlans

Routing instance        VLAN name               Tag          Interfaces
default-switch          default                 1
                                                             ae18.0*
                                                             ae19.0*
                                                             ae20.0*
default-switch          vlan-20                 20
                                                             ae18.0*
                                                             ae19.0*
                                                             ae20.0*
default-switch          vlan-30                 30
                                                             ae18.0*
                                                             ae19.0*
                                                             ae20.0*
```

We see that VLAN 20 and 30 are active on all three AE interfaces. Also notice the asterisk. This indicates that the interface for which the VLAN is active is ‘UP’.

To verify whether the VLAN is tagged or untagged, we can look at detailed VLAN information:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
root@jncie-vcf> show vlans 20 detail

Routing instance: default-switch
  VLAN Name: vlan-20                          State: Active
Tag: 20
Internal index: 3, Generation Index: 3, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ae18.0*,untagged,trunk
    ae19.0*,untagged,trunk
    ae20.0*,untagged,trunk
Number of interfaces: Tagged 0    , Untagged 3
Total MAC count: 0
VLAN Flags: Limit action drop

root@jncie-vcf> show vlans 30 detail

Routing instance: default-switch
  VLAN Name: vlan-30                          State: Active
Tag: 30
Internal index: 4, Generation Index: 4, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ae18.0*,tagged,trunk
    ae19.0*,tagged,trunk
    ae20.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 2
VLAN Flags: Limit action drop
```

Alternatively, we can also check the tagging parameters that are displayed when we issue the ‘show ethernet-switching interface’ command:

```
root@jncie-vcf> show ethernet-switching interface ae20.0

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down,
                         SCTL - shutdown by Storm-control,
                         MMAS - Mac-move action shutdown, AS - Autostate-exclude enabled)

Logical          Vlan          TAG     MAC          STP          Logical           Tagging
interface        members               limit        state        interface flags
ae20.0                                 294912                                      tagged
                 vlan-30       30      500          Forwarding    AD                tagged
                 vlan-20       20      500          Forwarding    AD                untagged
```

Finally, we can examine the switching table to see if any MAC addresses are being learned:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
root@jncie-vcf> show ethernet-switching table interface ae20

MAC database for interface ae20
MAC database for interface ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan-30             64:64:9b:dd:64:b8   D                -   ae20.0
```

Additional VCF Information and Details

Most commonly used VCF verification commands:

show virtual-chassis	Check the status of the VCF.
show virtual-chassis vc-port	List all ports that are VC-ports.
show virtual-chassis mode	Check the current mode and the mode after a reboot.
show virtual-chassis login	Check to see from what FPC you logged into another member.
show virtual-chassis protocol adjacency detail	Verify how long an adjacency has been Up.

Most commonly used VCF operational mode commands (possibly) required for successful VCF formation:

request virtual-chassis mode fabric	Set mode to fabric.
request virtual-chassis mode mixed	Set mode to mixed.
request virtual-chassis vc-port set pic-slot 0 port 0 member 0	Convert port into vc-port.
request virtual-chassis vc-port delete pic-slot 0 port 0 member 0	Convert port into non-vc-port.
request virtual-chassis renumber member-id 3 new-member-id 16	Renumber member 3.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Most commonly used VCF configuration commands:

set virtual-chassis preprovisioned	Set VCF mode to preprovisioned.
set virtual-chassis auto-provisioned	Set VCF mode to auto-provisioned.
set virtual-chassis member 0 role routing-engine	Configure member 0 as routing-engine.
set virtual-chassis member 2 role line-card	Configure member 2 as line-card.
set virtual-chassis member 0 serial-number XX	Configure the serial number for a member.
set virtual-chassis member 0 fabric-tree-root	Root the BUM tree at member 0.
set virtual-chassis member 0 location mars	Specify the location of a VCF member.
set virtual-chassis traceoptions	Activate traceoptions for the VCF.

The following sample output has been gathered on a VCF right after finishing task 1.18.
This means it is an autoprovisioned VCF.

Verify the status of the VCF:

This is the most common command to check the status of a VCF:

```
root@jncie-vcf> show virtual-chassis

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed
                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      Y  F    2  vcp-255/0/0
                                                                          3  vcp-255/0/1
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q 129   Backup       Y  F    2  vcp-255/0/0
                                                                          3  vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q   0   Linecard     Y  F    0  vcp-255/0/0
                                                                          1  vcp-255/0/1
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q   0   Linecard     Y  F    0  vcp-255/0/0
                                                                          1  vcp-255/0/1
```

It informs us about the following:

- We are looking at an autoprovisioned VCF
- The Fabric mode is mixed. This allows for a combination of QFX5100, QFX3500, QFX3600 and EX4300 switches.
- There are 4 member switches and each of the switches have their status, serial, model and role are listed.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

List currently active VC-ports:

In this example, the VC-ports are all auto-configured. This is due to the fact that the output was gathered on an auto-provisioned VCF:

```
root@jncie-vcf> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                 (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        2   vcp-255/0/0
0/1         Auto-Configured    -1    Up           10000        3   vcp-255/0/0
0/12        Auto-Configured    -1    Up           10000        2   vcp-255/0/12

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                 (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        2   vcp-255/0/1
0/1         Auto-Configured    -1    Up           10000        3   vcp-255/0/1

fpc2:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                 (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        0   vcp-255/0/0
0/1         Auto-Configured    -1    Up           10000        1   vcp-255/0/0
0/12        Auto-Configured    -1    Up           10000        0   vcp-255/0/12

fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                 (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        0   vcp-255/0/1
0/1         Auto-Configured    -1    Up           10000        1   vcp-255/0/1
```

If we would have issued the command on a preprovisioned VCF, or if any of the ports would have been configured using the '**request virtual-chassis vc-portvc-port set pic-slot 0 port 10 member 0**' command, the output would have looked like this:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
root@example-vcf-pre-provisioned> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type                Trunk  Status       Speed        Neighbor
or                               ID                 (mbps)       ID  Interface
PIC / Port
0/10        Configured           -1    Up           10000        5   vcp-255/2/0
```

Verify the current VCF member mode and check what mode will be active after a reboot:

```
root@jncie-vcf> show virtual-chassis mode member 0

fpc0:
--------------------------------------------------------------------------
Current mode : Fabric with mixed devices
Future mode after reboot : Fabric with mixed devices
```

Verify from what member you are logged in:

```
root@jncie-vcf> show virtual-chassis login

Current login session initiated from host re1

root@jncie-vcf> show virtual-chassis login

Current login session initiated from host fpc2
```

Checking VCF member neighbor relationship.

The VCF uses a modified version of IS-IS to build a PFE topology. To see how long nodes have been adjacent, we can issue the following command and check for the last transition:

```
root@jncie-vcf> show virtual-chassis protocol adjacency detail

fpc0:
--------------------------------------------------------------------------

80ac.acb7.ee40
  interface-name: vcp-255/0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:41:55 ago

f01c.2d45.5f00
  interface-name: vcp-255/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:41:56 ago

80ac.acb7.ee40
  interface-name: vcp-255/0/12.32768, State: Up, Expires in 57 secs
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Priority: 0, Up/Down transitions: 1, Last transition: 00:24:31 ago

fpc1:
--------------------------------------------------------------------------

80ac.acb7.ee40
  interface-name: vcp-255/0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 04:39:32 ago

f01c.2d45.5f00
  interface-name: vcp-255/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:39:33 ago

fpc2:
--------------------------------------------------------------------------

80ac.aca4.2400
  interface-name: vcp-255/0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:41:56 ago

541e.567d.5ec0
  interface-name: vcp-255/0/1.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 04:39:33 ago

80ac.aca4.2400
  interface-name: vcp-255/0/12.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:24:31 ago

fpc3:
--------------------------------------------------------------------------

80ac.aca4.2400
  interface-name: vcp-255/0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:41:57 ago

541e.567d.5ec0
  interface-name: vcp-255/0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:39:34 ago
```

Every VCF member, or FPC, lists the adjacent FPCs. To find out what identifier belongs to a particular FPC, we issue the following command:

```
root@jncie-vcf> show virtual-chassis protocol statistics | match "fpc|IS-IS stat"

fpc0:
--------------------------------------------------------------------------
IS-IS statistics for 80ac.aca4.2400:
PDU type       Received    Processed       Drops         Sent       Rexmit
fpc1:
--------------------------------------------------------------------------
IS-IS statistics for 541e.567d.5ec0:
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
PDU type       Received    Processed       Drops         Sent       Rexmit
fpc2:
--------------------------------------------------------------------------
IS-IS statistics for 80ac.acb7.ee40:
PDU type       Received    Processed       Drops         Sent       Rexmit
fpc3:
--------------------------------------------------------------------------
IS-IS statistics for f01c.2d45.5f00:
PDU type       Received    Processed       Drops         Sent       Rexmit
```

Back to our first command:

```
root@jncie-vcf> show virtual-chassis protocol adjacency detail
fpc0:
--------------------------------------------------------------------------

80ac.acb7.ee40
  interface-name: vcp-255/0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 3, Last transition: 04:41:55 ago
```

We can now conclude that FPC 0 has an adjacency with FPC2 across VCP-255/0/0. VCP-255/0/0 translates to port 0 on PIC 0. This translates to XE-0/0/0.

Deleting and configuring VC ports.

In this example, we delete an auto-configure VC-port. This action turns the port into a regular port. After this, we turn the port into a configured VC-port.

```
root@jncie-vcf> show virtual-chassis vc-port member 0

fpc0:
--------------------------------------------------------------------------
Interface   Type                Trunk   Status      Speed        Neighbor
or                               ID                 (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured     -1      Up          10000        2   vcp-255/0/0
0/1         Auto-Configured     -1      Up          10000        3   vcp-255/0/0
0/12        Auto-Configured     -1      Up          10000        2   vcp-255/0/12

root@jncie-vcf> request virtual-chassis vc-port delete pic-slot 0 port 12

vc-port successfully deleted

root@jncie-vcf> show virtual-chassis vc-port member 0

fpc0:
--------------------------------------------------------------------------
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        2   vcp-255/0/0
0/1         Auto-Configured    -1    Up           10000        3   vcp-255/0/0

root@jncie-vcf> request virtual-chassis vc-portvc-port set pic-slot 0 port 12 member 0

fpc0:
--------------------------------------------------------------------------
Port conversion initiated,  use show virtual-chassis vc-port to verify

{master:0}
root@jncie-vcf> show virtual-chassis vc-port member 0

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
0/0         Auto-Configured    -1    Up           10000        2   vcp-255/0/0
0/1         Auto-Configured    -1    Up           10000        3   vcp-255/0/0
0/12        Configured         -1    Up           10000        2   vcp-255/0/12
```

Renumbering VCF member:

In this example, we change member 3 into member 16:

```
root@jncie-vcf> show virtual-chassis

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed
                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model              prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q     129   Master*   Y     F    2   vcp-255/0/0
                                                                               3   vcp-255/0/1
                                                                               2   vcp-255/0/12
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q     129   Backup    Y     F    2   vcp-255/0/0
                                                                               3   vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q     0     Linecard  Y     F    0   vcp-255/0/0
                                                                               0   vcp-255/0/12
                                                                               1   vcp-255/0/1
3 (FPC 3)  Prsnt    VF3714370037 qfx5100-48s-6q     0     Linecard  Y     F    0   vcp-255/0/0
                                                                               1   vcp-255/0/1

{master:0}
root@jncie-vcf> request virtual-chassis renumber member-id 3 new-member-id 16

To move configuration specific to member ID 3 to member ID 16, please
use the replace command. e.g. replace pattern ge-3/ with ge-16/
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
If member-specific configuration groups are present, perform a
"commit full" to synchronize inheritance with the new member number.

Do you want to continue ? [yes,no] (no) yes
```

After 1 minute or so:

```
root@jncie-vcf> show virtual-chassis

Auto-provisioned Virtual Chassis Fabric
Fabric ID: 5a5a.f8d5.3b3e
Fabric Mode: Mixed
                                                    Mstr           Mixed Route Neighbor List
Member ID  Status   Serial No    Model          prio  Role      Mode  Mode ID  Interface
0 (FPC 0)  Prsnt    VF3715490286 qfx5100-48s-6q 129   Master*      Y  F    2  vcp-255/0/0
                                                                          2  vcp-255/0/12
                                                                          16 vcp-255/0/1
1 (FPC 1)  Prsnt    VB3715280174 qfx5100-96s-8q 129   Backup       Y  F    2  vcp-255/0/0
                                                                          16 vcp-255/0/1
2 (FPC 2)  Prsnt    VF3715510230 qfx5100-48s-6q   0   Linecard     Y  F    0  vcp-255/0/0
                                                                          0  vcp-255/0/12
                                                                          1  vcp-255/0/1
16 (FPC 16) Prsnt   VF3714370037 qfx5100-48s-6q   0   Linecard     Y  F    0  vcp-255/0/0
                                                                          1  vcp-255/0/1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Appendix – Chapter 2: Clos IP Fabric

Part 1: EBGP-based Clos IP Fabric

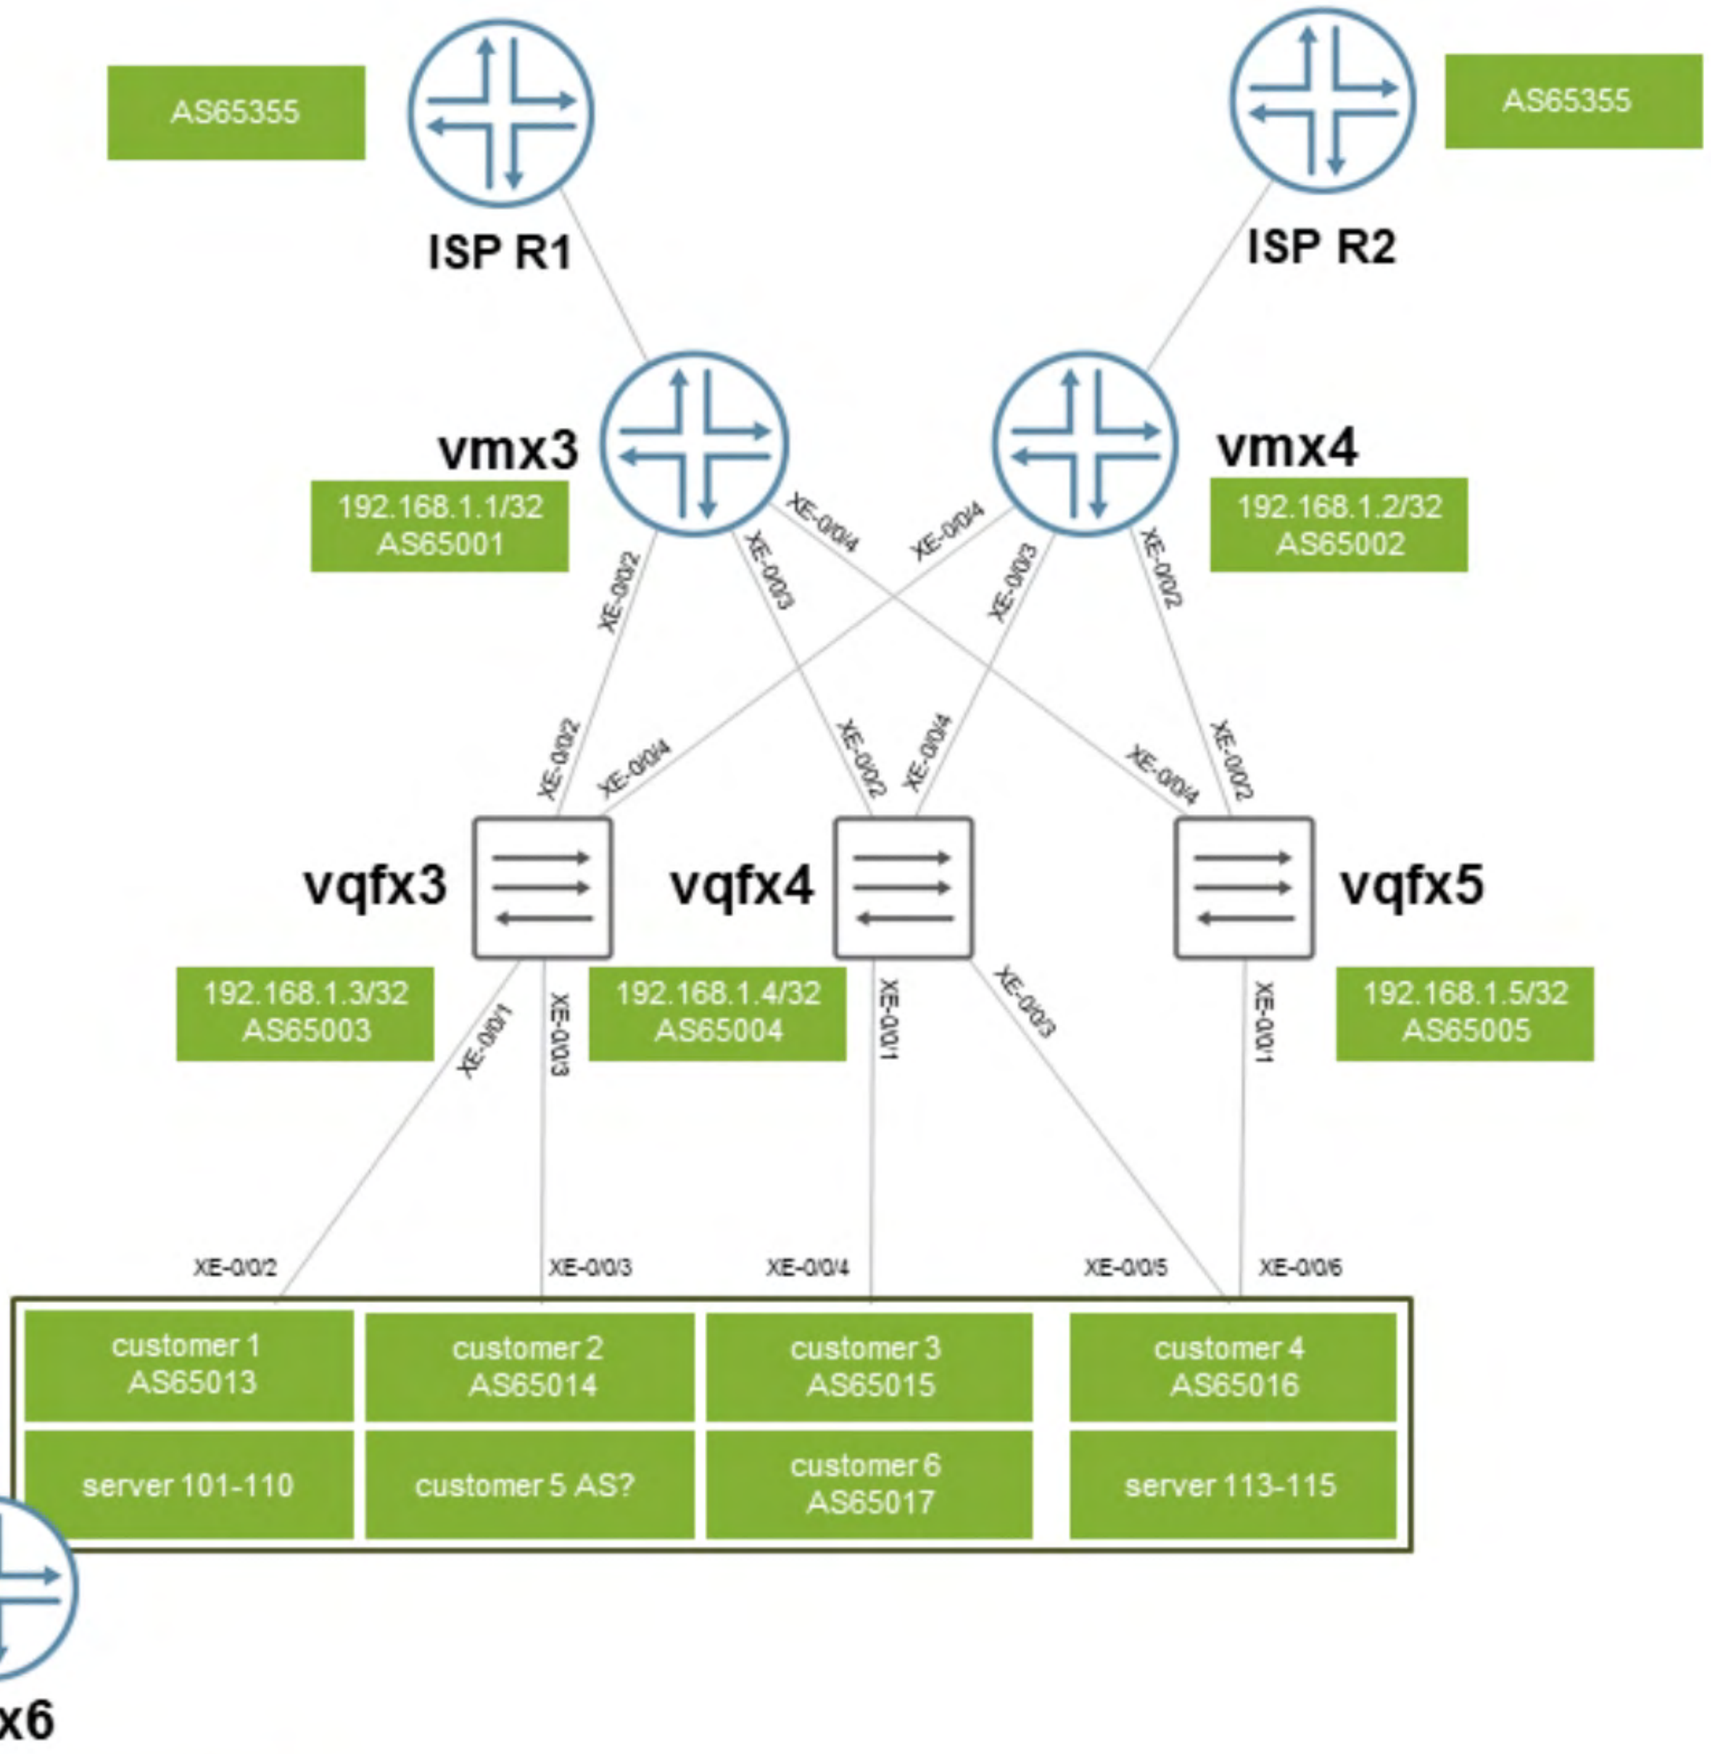

Figure 17 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

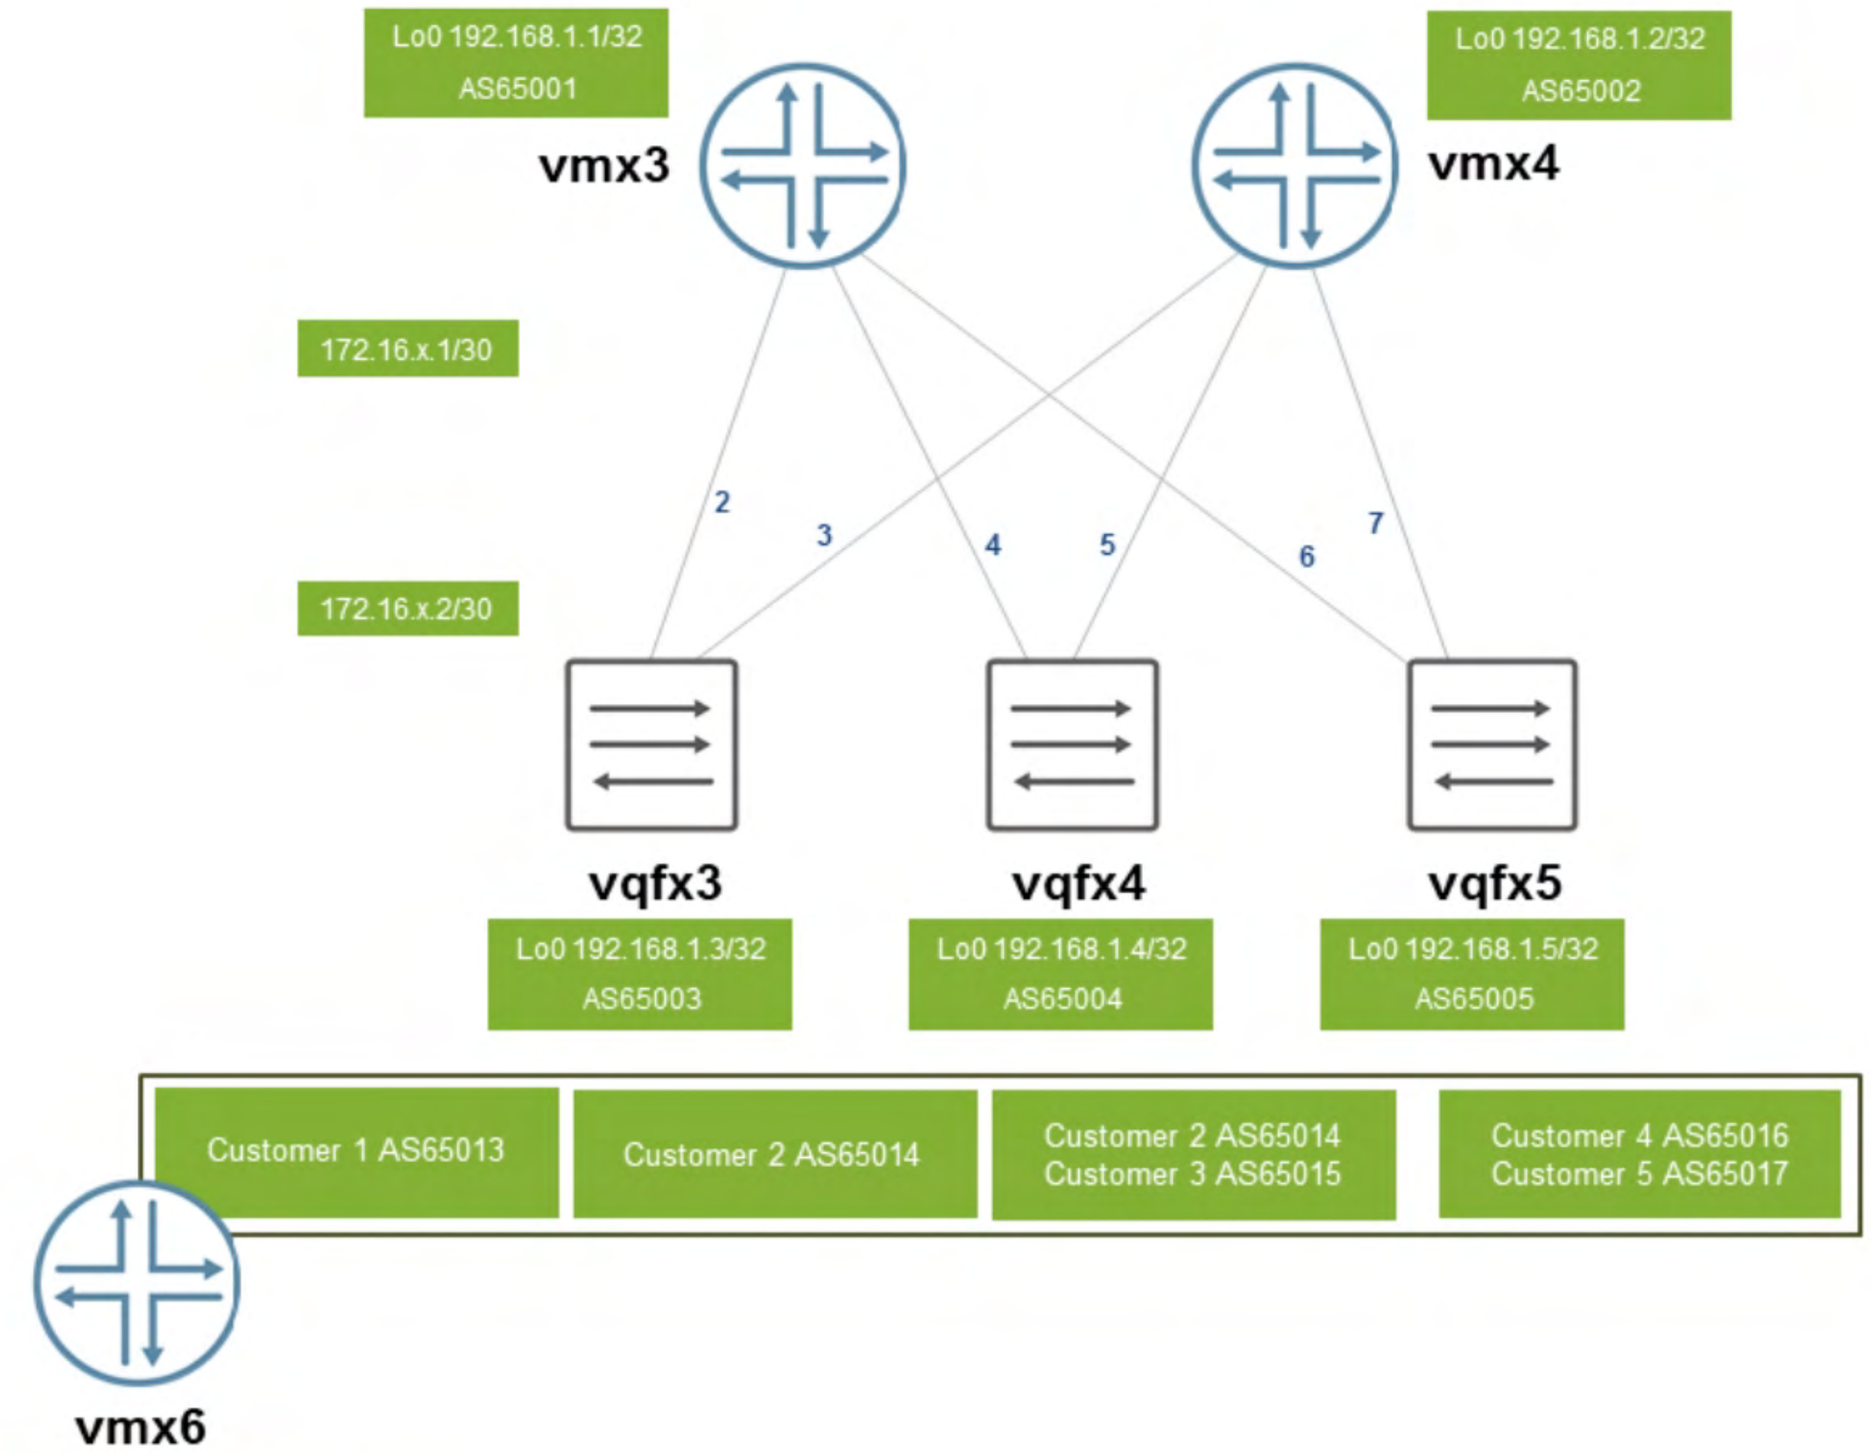

Figure 2 - Logical topology

Note: starting topology is preconfigured with IP-addressing only.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.1: Core Network BGP Configuration 1

- Configure BGP according to the information provided in the logical topology.
- Ensure IP-connectivity between all of the loopback IP addresses.
- Make sure all BGP sessions are 'Established' and that sessions state changes are logged to syslog.
- Write all BGP-related syslog messages to a file called ‘bgp.log’.

Solution

vMX3:

To write BGP related messages to a file called bgp.log:

```
set system syslog file bgp.log daemon any
set system syslog file bgp.log match bgp
```

BGP configuration:

```
set protocols bgp log-updown
set protocols bgp group jncie-fabric export export-lo0
set protocols bgp group jncie-fabric neighbor 172.16.2.2 peer-as 65003
set protocols bgp group jncie-fabric neighbor 172.16.4.2 peer-as 65004
set protocols bgp group jncie-fabric neighbor 172.16.6.2 peer-as 65005

set routing-options autonomous-system 65001
```

BGP export policy configuration:

```
set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept
```

vMX4:

```
set system syslog file bgp.log daemon any
set system syslog file bgp.log match bgp
set protocols bgp log-updown
set protocols bgp group jncie-fabric export export-lo0
set protocols bgp group jncie-fabric neighbor 172.16.3.2 peer-as 65003
set protocols bgp group jncie-fabric neighbor 172.16.5.2 peer-as 65004
set protocols bgp group jncie-fabric neighbor 172.16.7.2 peer-as 65005

set routing-options autonomous-system 65002

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set system syslog file bgp.log daemon any
set system syslog file bgp.log match bgp

set protocols bgp log-updown
set protocols bgp group jncie-fabric export export-lo0
set protocols bgp group jncie-fabric neighbor 172.16.2.1 peer-as 65001
set protocols bgp group jncie-fabric neighbor 172.16.3.1 peer-as 65002

set routing-options autonomous-system 65003

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept
```

vQFX4:

```
set system syslog file bgp.log daemon any
set system syslog file bgp.log match bgp

set protocols bgp log-updown
set protocols bgp group jncie-fabric export export-lo0
set protocols bgp group jncie-fabric neighbor 172.16.4.1 peer-as 65001
set protocols bgp group jncie-fabric neighbor 172.16.5.1 peer-as 65002

set routing-options autonomous-system 65004

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept
```

vQFX5:

```
set system syslog file bgp.log daemon any
set system syslog file bgp.log match bgp

set protocols bgp log-updown
set protocols bgp group jncie-fabric export export-lo0
set protocols bgp group jncie-fabric neighbor 172.16.6.1 peer-as 65001
set protocols bgp group jncie-fabric neighbor 172.16.7.1 peer-as 65002

set routing-options autonomous-system 65005

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

The fastest and easiest way to verify all EBGP sessions is by issuing the '**show bgp summary**' command from both of the spines:

```
jncie@vMX3> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       8          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2               65003          6          5       0       0           9 2/3/3/0
0/0/0/0
172.16.4.2               65004          6          7       0       0           7 1/2/2/0
0/0/0/0
172.16.6.2               65005          6          8       0       0           7 1/3/3/0
0/0/0/0

jncie@vMX4> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.3.2               65003          6          6       0       0          30 2/3/3/0
0/0/0/0
172.16.5.2               65004          7          8       0       0          27 1/4/4/0
0/0/0/0
172.16.7.2               65005          8          7       0       0          55 1/3/3/0
0/0/0/0
```

We also needed to ensure IP-connectivity between all of the loopback IP addresses. We start the verification of this objective by examining the routing-table:

```
jncie@vMX3> show route protocol bgp

inet.0: 13 destinations, 17 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.2/32     *[BGP/170] 00:01:37, localpref 100
                      AS path: 65003 65002 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:01:33, localpref 100
                      AS path: 65004 65002 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
                    [BGP/170] 00:01:36, localpref 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    AS path: 65005 65002 I, validation-state: unverified
                  > to 172.16.6.2 via xe-0/0/4.0
192.168.1.3/32    *[BGP/170] 00:01:38, localpref 100
                    AS path: 65003 I, validation-state: unverified
                  > to 172.16.2.2 via xe-0/0/2.0
                  [BGP/170] 00:01:36, localpref 100
                    AS path: 65005 65002 65003 I, validation-state: unverified
                  > to 172.16.6.2 via xe-0/0/4.0
192.168.1.4/32    *[BGP/170] 00:01:37, localpref 100
                    AS path: 65004 I, validation-state: unverified
                  > to 172.16.4.2 via xe-0/0/3.0
192.168.1.5/32    *[BGP/170] 00:01:36, localpref 100
                    AS path: 65005 I, validation-state: unverified
                  > to 172.16.6.2 via xe-0/0/4.0
                  [BGP/170] 00:01:37, localpref 100
                    AS path: 65003 65002 65005 I, validation-state: unverified
                  > to 172.16.2.2 via xe-0/0/2.0
```

We can see all of the loopback IP addresses are present in the routing table of vMX3. To verify that vMX3 also advertised its loopback IP to its peers, we can issue the following command:

```
jncie@vMX3> show route advertising-protocol bgp 172.16.2.2

inet.0: 13 destinations, 17 routes (13 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.1/32          Self                                    I
* 192.168.1.2/32          Self                                    65004 65002 I
* 192.168.1.4/32          Self                                    65004 I
* 192.168.1.5/32          Self                                    65005 I
```

Instead of checking the advertised routes from vMX3, we could have also chosen to examine the routing table of another device. vQFX3 for example:

```
jncie@vQFX3> show route protocol bgp

inet.0: 13 destinations, 15 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[BGP/170] 00:01:18, localpref 100
                     AS path: 65001 I, validation-state: unverified
                   > to 172.16.2.1 via xe-0/0/2.0
192.168.1.2/32     *[BGP/170] 00:01:16, localpref 100
                     AS path: 65002 I, validation-state: unverified
                   > to 172.16.3.1 via xe-0/0/4.0
192.168.1.4/32     *[BGP/170] 00:01:16, localpref 100
                     AS path: 65001 65004 I, validation-state: unverified
                   > to 172.16.2.1 via xe-0/0/2.0
                   [BGP/170] 00:01:12, localpref 100
                     AS path: 65002 65004 I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                     > to 172.16.3.1 via xe-0/0/4.0
192.168.1.5/32     *[BGP/170] 00:01:16, localpref 100
                       AS path: 65002 65005 I, validation-state: unverified
                     > to 172.16.3.1 via xe-0/0/4.0
                     [BGP/170] 00:01:16, localpref 100
                       AS path: 65001 65005 I, validation-state: unverified
                     > to 172.16.2.1 via xe-0/0/2.0

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

Next, we examine the content of the 'bgp.log' file: I

```
jncie@vMX3> show log bgp.log | last 10

Apr 1 13:31:06 vMX3 rpd[4946]: BGP_SET_INTERFACE_FAILED: bgp_set_peer_if: BGP peer 172.16.4.2
(External AS 65004) interface not found. Leaving peer idled
Apr 1 13:31:06 vMX3 rpd[4946]: BGP_SET_INTERFACE_FAILED: bgp_set_peer_if: BGP peer 172.16.6.2
(External AS 65005) interface not found. Leaving peer idled
Apr 1 13:31:06 vMX3 rpd[4946]: jtask_jthr: RUNNING scheduler main loop: bgpio-0-th
Apr 1 13:31:11 vMX3 rpd[4946]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 172.16.2.2 (External
AS 65003) changed state from OpenConfirm to Established (event RecvKeepAlive) (instance
master)
Apr 1 13:31:14 vMX3 rpd[4946]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 172.16.4.2 (External
AS 65004) changed state from OpenConfirm to Established (event RecvKeepAlive) (instance
master)
Apr 1 13:31:19 vMX3 rpd[4946]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 172.16.6.2 (External
AS 65005) changed state from OpenConfirm to Established (event RecvKeepAlive) (instance
master)
```

Instead of outputting the file to the terminal, we can also continuously monitor the log and have newly added lines send to the terminal as we continue to work on the lab. To do this, we issue the following command:

```
jncie@vMX3> monitor start bgp.log
```

To demonstrate how this works, we move to vQFX3 and issue the following command:

```
jncie@vQFX3> restart routing

Routing protocols process started, pid 1699
```

When we move back to vMX3 terminal, we can see that the following messages have appeared on screen:

```
jncie@vMX3>
*** bgp.log ***
Apr 1 13:32:27 vMX3 rpd[4946]: bgp_handle_notify:4220: NOTIFICATION received from 172.16.2.2
(External AS 65003): code 6 (Cease) subcode 3 (Peer Unconfigured)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Apr 1 13:32:27 vMX3 rpd[4946]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 172.16.2.2 (External
AS 65003) changed state from Established to Idle (event RecvNotify) (instance master)
```

This can also be a nice feat during maintenance. When you want stop monitoring the file, simply use the 'monitor stop' command.

Task 2.2: Core Network BGP Configuration 2

- Make sure that BGP can detect the loss of a neighbor within 15 seconds. Do not alter the hold time anywhere.
- Activate per-flow load-balancing between the leaf switches and the spine routers.
- Authenticate the BGP sessions using the MD5-hased password ‘jncie’

Solution

On every device:

To ensure that the loss of a neighbor is detected within 15 seconds, we use BFD:

```
set protocols bgp group jncie-fabric bfd-liveness-detection minimum-interval 3000
set protocols bgp group jncie-fabric bfd-liveness-detection multiplier 4
```

To have BGP consider equal-cost paths from different, or multiple, autonomous systems:

```
set protocols bgp group jncie-fabric multipath multiple-as
```

Device load-balancing configuration:

```
set routing-options forwarding-table export lbpp
set policy-options policy-statement lbpp term 1 then load-balance per-packet
```

To authenticate BGP sessions:

```
set protocols bgp group jncie-fabric authentication-key jncie
```

Verification

To verify BFD, we issue the following command from both spine routers:

```
jncie@vMX3> show bfd session

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.2.2               Up        xe-0/0/2.0     12.000    3.000        4
172.16.4.2               Up        xe-0/0/3.0     12.000    3.000        4
172.16.6.2               Up        xe-0/0/4.0     12.000    3.000        4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To properly verify load balancing, we need to check the forwarding table entries. When we check the forwarding table of a leaf switch, we need to check the route towards a destination that we know is eligible for load balancing. We know that the leaf switches will be able to load balance traffic to each other using the uplinks towards the spines. Therefore, we check the forwarding table entry on a leaf switch (vQFX3 in this example) towards another leaf switch:

```
jncie@vQFX3> show route forwarding-table destination 192.168.1.5

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.5/32     user     0                    ulst   131070    12
                              172.16.2.1         ucst     1713     5 xe-0/0/2.0
                              172.16.3.1         ucst     1714     6 xe-0/0/4.0
```

When we verify load balancing on one of the spines, we need to check the route towards the other spine:

```
jncie@vMX3> show route forwarding-table destination 192.168.1.2

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.2/32     user     0                    ulst  1048575     2
                              172.16.2.2         ucst      513    22 xe-0/0/2.0
                              172.16.4.2         ucst      575     9 xe-0/0/3.0
                              172.16.6.2         ucst      576     9 xe-0/0/4.0
```

It is important to realize that not all destinations will be load balanced. This can be because the routing entries are not considered equal cost or because there is only one entry in the routing table for a particular destination.

Task 2.3: ISP Uplink Configuration

- Configure EBGP between the ISP and both vMX3 as well as vMX4.
- Make sure that the keepalive is 3 seconds.
- Present the IP-Fabric to both ISP routers as AS65500.

Solution

vMX3:

To influence the keepalive value, we configure a hold-time for the BGP session:

```
set protocols bgp group isp neighbor 10.0.0.1 hold-time 9
set protocols bgp group isp neighbor 10.0.0.1 peer-as 65355
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We use the 'local-as' knob to negotiate a BGP session using an AS other than the one configured:

```
set protocols bgp group isp neighbor 10.0.0.1 local-as 65500
```

vMX4:

```
set protocols bgp group isp neighbor 10.0.0.5 hold-time 9
set protocols bgp group isp neighbor 10.0.0.5 peer-as 65355
set protocols bgp group isp neighbor 10.0.0.5 local-as 65500
```

Verification

First, we check whether or not both sessions with the ISP reach the established state:

```
jncie@vMX3> show bgp summary

Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      73         34          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1              65355        189        202       0       0        8:23 1/1/1/0
0/0/0/0
172.16.2.2            65003         23         25       0       1        4:20 19/30/29/0
0/0/0/0
172.16.4.2            65004         20         29       0       1        4:22 7/12/11/0
0/0/0/0
172.16.6.2            65005         23         20       0       1        4:24 7/30/29/0
0/0/0/0
```

```
jncie@vMX4> show bgp summary

Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      37         34          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.5              65355        212        218       0       0        9:25 1/1/1/0
0/0/0/0
172.16.3.2            65003         21         23       0       1        5:45 19/19/19/0
0/0/0/0
172.16.5.2            65004         22         22       0       1        5:45 7/10/10/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.7.2              65005        19         25       0       1        5:45 7/7/7/0
0/0/0/0
```

We then verify the session characteristics. Since we cannot configure a keepalive value directly, it is important to understand how this value is determined.

When two BGP peers start to form a peering relation, they exchange a BGP 'Open message'. This message contains a 'hold-time'. Both peers will have a received hold-time and a local hold-time. Both peers will use the lowest of the two hold time values. The hold-time value is then used to determine the keepalive frequency. The keepalive will be one third of the value of the hold-time. Therefore, we configured a hold-time of 9:

```
jncie@vMX3> show bgp neighbor 10.0.0.1

Peer: 10.0.0.1+179 AS 65355    Local: 10.0.0.2+63247 AS 65500
  Group: isp                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference HoldTime LogUpDown PeerAS LocalAS Refresh>
  Holdtime: 9 Preference: 170 Local AS: 65500 Local System AS: 65001
  Number of flaps: 0
  Peer ID: 10.0.0.1         Local ID: 192.168.1.1       Active Holdtime: 9
  Keepalive Interval: 3           Group index: 1    Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  Local Interface: xe-0/0/5.0
```

The output above shows us that the keepalive has been set to 3 seconds. Note that when you use a BGP hold-time value of 1 through 19, Juniper recommends that you also configure precision timers using the following configuration command: **set protocols bgp precision-timers**.

The previous output also told us that we are using autonomous system 65500 in the session with our ISP.

Task 2.4: IP-Fabric BGP Customer Configuration

- Configure EBGP between the IP-Fabric and all of the customers.
- The IP address customer 3 is using to set up peering is unknown. Use monitoring tools to obtain what address the BGP session is sourced with.
- The AS number of customer 5 is unknown, you need to find out what AS the customer is using.
- Customer 6 is configured to peer with the wrong AS number. Make sure the peering is established and make sure this wrong AS number will not be added to the AS-path when prefixes from that customer are advertise to other devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- All sessions need to be authenticated using the password 'jncie'.

Solution

To learn what IP address customer 3 is using to source the BGP session, we can issue the following command on any device connected to customer 3. The following steps are carried out on vQFX3 since it also has a connection to customer 5:

```
monitor traffic interface xe-0/0/1.111 matching "tcp port 179"
```

The other information required to configure the BGP session can be obtained through the use of traceoptions. The following example demonstrates how we can obtain the AS of customer 5. First, we simply configure the BGP peering with that customer and use any AS. In this example, we will use AS 1:

```
set protocols bgp group customer-5 authentication-key jncie
set protocols bgp group customer-5 peer-as 1
set protocols bgp group customer-5 neighbor 172.16.200.22
```

Then, we configure traceoptions for this peer. We start out setting the flag to general:

```
set protocols bgp group customer-5 traceoptions file bgp.trace.log
set protocols bgp group customer-5 traceoptions flag general
```

If the general flag would not be enough, we can start to consider some of the more detailed options. In this case, setting the general flag is enough. After a minute or so, we should be able to observe the following in the log:

```
jncie@vQFX3> show log bgp.trace.log

Sep 20 09:38:32 trace_on: Tracing to "/var/log/bgp.trace.log" started
...
Sep 20 09:38:44.805922 bgp_recv_open: called for peer 172.16.200.22 (External AS 1)
Sep 20 09:38:44.805976 bgp_process_open:3764: NOTIFICATION sent to 172.16.200.22 (External AS
1): code 2 (Open Message Error) subcode 2 (bad peer AS number), Reason: peer 172.16.200.22
(External AS 1) claims 65535, 1 configured
...
```

From this information, we can conclude that the peer is configured with AS 65535.

When we configure the peering with customer 6, we also use the traceoptions to figure out what AS that customer is trying to peer with. As soon as we learn what AS customer 6 is configured to peer with, we add that AS to the BGP peer configuration using the ‘local-as’ knob. Since we do not want this AS number to appear in the AS-path when we advertise routing information to other pears, we add the ‘private’ keyword.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set protocols bgp group customer-1 authentication-key jncie
set protocols bgp group customer-1 peer-as 65013
set protocols bgp group customer-1 neighbor 172.16.200.2

set protocols bgp group customer-2 authentication-key jncie
set protocols bgp group customer-2 peer-as 65014
set protocols bgp group customer-2 neighbor 172.16.200.6

set protocols bgp group customer-5 authentication-key jncie
set protocols bgp group customer-5 peer-as 65535
set protocols bgp group customer-5 neighbor 172.16.200.22
```

vQFX4:

```
set protocols bgp group customer-3 authentication-key jncie
set protocols bgp group customer-3 peer-as 65015
set protocols bgp group customer-3 neighbor 172.16.200.10

set protocols bgp group customer-4 authentication-key jncie
set protocols bgp group customer-4 peer-as 65016
set protocols bgp group customer-4 neighbor 172.16.200.14

set protocols bgp group customer-6 peer-as 65017
set protocols bgp group customer-6 local-as 65003
set protocols bgp group customer-6 local-as private
set protocols bgp group customer-6 neighbor 172.16.200.26
set protocols bgp group customer-6 authentication-key jncie
```

vQFX5:

```
set protocols bgp group customer-4 authentication-key jncie
set protocols bgp group customer-4 peer-as 65016
set protocols bgp group customer-4 neighbor 172.16.200.18
```

Verification

The verification of this task is pretty straightforward. We simply issue the 'show bgp summary' command on all of the leaf switches:

```
jncie@vQFX3> show bgp summary

Groups: 4 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      32         30          0          0          0          0
inet6.0
                       0          0          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Peer                    AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.1              65001      30         28       0       1        7:11 12/13/13/0
0/0/0/0
172.16.3.1              65002      26         24       0       1        7:17 12/13/13/0
0/0/0/0
172.16.200.2            65013     151        159       0       0       22:24 2/2/2/0
0/0/0/0
172.16.200.6            65014      51         60       0       0       22:28 2/2/2/0
0/0/0/0
172.16.200.22           65535      51         59       0       0       22:20 2/2/2/0
0/0/0/0
```

```
jncie@vQFX4> show bgp summary

Groups: 4 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      58         11          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                    AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.4.1              65001      34         25       0       1        7:17 3/27/9/0
0/0/0/0
172.16.5.1              65002      25         24       0       1        7:23 3/25/7/0
0/0/0/0
172.16.200.10           65015      51         74       0       0       22:19 1/2/2/0
0/0/0/0
172.16.200.14           65016      59         59       0       0       22:07 2/2/2/0
0/0/0/0
172.16.200.26           65017      51         59       0       0       22:15 2/2/2/0
0/0/0/0
```

```
jncie@vQFX5> show bgp summary

Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      53         50          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                    AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.6.1              65001      26         28       0       1        7:21 24/24/24/0
0/0/0/0
172.16.7.1              65002      28         21       0       1        7:23 24/26/26/0
0/0/0/0
172.16.200.18           65016      54         59       0       0       22:01 2/3/2/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To further clarify the use of the local-as private option, we delete it from the configuration:

```
jncie@vQFX4# delete protocols bgp group customer-6 local-as private
```

After this, we check a route advertisement to customer 4 on vQFX5:

Note: This may be learned in the opposite direction depending on the protocol convergence timing for your environment.

```
jncie@vQFX5> show route advertising-protocol bgp 172.16.200.18 10.0.10.0/24

inet.0: 34 destinations, 48 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.10.0/24            Self                                    65001 65004 65003 65017 I
```

The route we just checked originated from within the customer 6 network. We can see that 65003, the AS used in the BGP peering session, appears in the AS path.

Now, we re-activate the private option:

```
set protocols bgp group customer-6 local-as private
```

After this, we can observe the following:

```
jncie@vQFX5> show route advertising-protocol bgp 172.16.200.18 10.0.10.0/24

inet.0: 34 destinations, 51 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.10.0/24            Self                                    65002 65004 65017 I
```

AS 65003 no longer appears in the AS path.

Task 2.5: Server 100-110 Internet Feed.

- Server 100-110 have their gateway configured on vQFX3. Advertise the subnets in which these servers are placed.
- Make sure that vQFX3 will advertise any additional subnet that is configured for the same interface without any additional configuration.

Solution

vQFX3:

To complete this task, we configure an export policy that references a prefix-list configured to use 'apply-path'. This feature will make the prefix-list inherit its configuration from another part of the configuration stanza. We configure the apply-path to include all IPv4 addresses configured on any unit of the xe-0/0/1 interface:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group jncie-fabric export export-direct

set policy-options policy-statement export-direct term xe-0/0/1 from protocol direct
set policy-options policy-statement export-direct term xe-0/0/1 from prefix-list xe-0/0/1
set policy-options policy-statement export-direct term xe-0/0/1 then accept

set policy-options prefix-list xe-0/0/1 apply-path "interfaces xe-0/0/1 unit <*> family inet
address <*>"
```

Verification

To see what the prefix-list inherits, we issue the following command:

```
jncie@vQFX3> show configuration policy-options prefix-list xe-0/0/1 |display inheritance

##
## apply-path was expanded to:
##     10.200.100.0/24;
##     10.200.101.0/24;
##     10.200.102.0/24;
##     10.200.103.0/24;
##     10.200.104.0/24;
##     10.200.105.0/24;
##     10.200.106.0/24;
##     10.200.107.0/24;
##     10.200.108.0/24;
##     10.200.109.0/24;
##     10.200.110.0/24;
##
apply-path "interfaces xe-0/0/1 unit <*> family inet address <*>";
```

Using 'apply-path' can be useful for certain provisioning tasks. Whenever another subnet is configured, this additional range is automatically advertised without touching the BGP policy configuration.

To verify our policy, we check the advertised routes from vQFX3:

```
jncie@vQFX3> show route advertising-protocol bgp 172.16.2.1 10.200.100.0/20

inet.0: 56 destinations, 69 routes (56 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.200.100.0/24         Self                                    I
* 10.200.101.0/24         Self                                    I
* 10.200.102.0/24         Self                                    I
* 10.200.103.0/24         Self                                    I
* 10.200.104.0/24         Self                                    I
* 10.200.105.0/24         Self                                    I
* 10.200.106.0/24         Self                                    I
* 10.200.107.0/24         Self                                    I
* 10.200.108.0/24         Self                                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 10.200.109.0/24          Self                                    I
* 10.200.110.0/24          Self                                    I
```

Task 2.6: Server 113-115 Configuration

- Configure VRRP and authenticate the control messages with the 'jncie' password.
- You can only have VRRP exchange messages inside VLAN 113.
- As long as vQFX5 has a default route, vQFX4 should not be the primary router.
- The rest of the IP-Fabric should prefer vQFX5 to reach the servers.
- Make sure the interface of vQFX5 will remain down for at least 45 seconds after the device boots.

Solution

vQFX4:

When we configure VRRP for an interface, the router will start exchanging VRRP messages inside that subnet/VLAN. Since we will exchange VRRP messages inside VLAN 113, we need to configure VRRP with all of the required parameters.

Since we can only exchange VRRP messages inside VLAN 113, we need to make sure that our VRRP configuration for VLANs 114 and 115 include the 'vrrp-inherit-from active-interface' statement. When using this statement, we make sure that the VRRP group will follow the referenced interface without sending VRRP messages inside its own subnet/VLAN.

```
set interfaces xe-0/0/3 unit 113 vlan-id 113
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.252/24 vrrp-group 1 virtual-
address 10.200.113.254
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.252/24 vrrp-group 1 priority
150
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.252/24 vrrp-group 1 accept-
data
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.252/24 vrrp-group 1
authentication-type md5
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.252/24 vrrp-group 1
authentication-key jncie
set interfaces xe-0/0/3 unit 113 family inet address 10.200.113.253/24

set interfaces xe-0/0/3 unit 114 vlan-id 114
set interfaces xe-0/0/3 unit 114 family inet address 10.200.114.252/24 vrrp-group 2 virtual-
address 10.200.114.254
set interfaces xe-0/0/3 unit 114 family inet address 10.200.114.252/24 vrrp-group 2 vrrp-
inherit-from active-interface xe-0/0/3.113
set interfaces xe-0/0/3 unit 114 family inet address 10.200.114.252/24 vrrp-group 2 vrrp-
inherit-from active-group 1

set interfaces xe-0/0/3 unit 115 vlan-id 115
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces xe-0/0/3 unit 115 family inet address 10.200.115.252/24 vrrp-group 2 virtual-
address 10.200.115.254
set interfaces xe-0/0/3 unit 115 family inet address 10.200.115.252/24 vrrp-group 2 vrrp-
inherit-from active-interface xe-0/0/3.113
set interfaces xe-0/0/3 unit 115 family inet address 10.200.115.252/24 vrrp-group 2 vrrp-
inherit-from active-group 1
```

To make sure the other routers prefer vQFX5, we advertise the subnets from vQFX4 with additional AS paths prepended to the BGP path:

```
set protocols bgp group jncie-fabric export export-vrrp

set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.113.0/24 exact
set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.114.0/24 exact
set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.115.0/24 exact
set policy-options policy-statement export-vrrp term 1 then as-path-prepend "65004 65004"
set policy-options policy-statement export-vrrp term 1 then accept
```

vQFX5:

We configure vQFX5 to damp interface transitions using the following configuration command:

```
set interfaces xe-0/0/1 hold-time up 45000
set interfaces xe-0/0/1 hold-time down 0
```

This will make sure that if and when the link comes 'Up', the vQFX will damp the state for 45 seconds. Consider the situation wherein you reboot a QFX. When the QFX comes back online, interfaces are brought up immediately. This is long before the routing-protocols are started. Without damping the interface, vQFX5 would be able to become VRRP master before the routing protocols converge, effectively drawing in and blackholing traffic.

Note how on vQFX5, we configure VRRP using ' track route 0.0.0.0/0 '. Because of this, vQFX5 will continuously test this condition. As soon as the route disappears from the routing-table, the priority will decrease with 250, leaving the vQFX with a priority of 0:

```
set interfaces xe-0/0/1 unit 113 vlan-id 113
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1 virtual-
address 10.200.113.254
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1 priority
250
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1 accept-
data
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1
authentication-type md5
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1
authentication-key jncie
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces xe-0/0/1 unit 113 family inet address 10.200.113.253/24 vrrp-group 1 track
route 0.0.0.0/0 routing-instance default priority-cost 250

set interfaces xe-0/0/1 unit 114 vlan-id 114
set interfaces xe-0/0/1 unit 114 family inet address 10.200.114.253/24 vrrp-group 2 virtual-
address 10.200.114.254
set interfaces xe-0/0/1 unit 114 family inet address 10.200.114.253/24 vrrp-group 2 vrrp-
inherit-from active-interface xe-0/0/1.113
set interfaces xe-0/0/1 unit 114 family inet address 10.200.114.253/24 vrrp-group 2 vrrp-
inherit-from active-group 1

set interfaces xe-0/0/1 unit 115 vlan-id 115
set interfaces xe-0/0/1 unit 115 family inet address 10.200.115.253/24 vrrp-group 2 virtual-
address 10.200.115.254
set interfaces xe-0/0/1 unit 115 family inet address 10.200.115.253/24 vrrp-group 2 vrrp-
inherit-from active-interface xe-0/0/1.113
set interfaces xe-0/0/1 unit 115 family inet address 10.200.115.253/24 vrrp-group 2 vrrp-
inherit-from active-group 1

set protocols bgp group jncie-fabric export export-vrrp

set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.113.0/24 exact
set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.114.0/24 exact
set policy-options policy-statement export-vrrp term 1 from route-filter 10.200.115.0/24 exact
set policy-options policy-statement export-vrrp term 1 then accept
```

Note: in this case the configuration of 'track' implies that we would not necessarily have to dampen the interface state. Since it is a task requirement, both are configured.

Note: To effectively verify this task, you should wait until completing task 2.8. This will ensure that you are not learning the 0/0 route from your EBGP peer on vQFX5 as well. The following outputs were captured after completing task 2.8.

Verification

First, we verify VRRP:

```
jncie@vQFX5> show vrrp summary

Interface     State       Group   VR state       VR Mode     Type    Address
xe-0/0/1.113  up              1   master          Active     lcl     10.200.113.253
                                                             vip     10.200.113.254
xe-0/0/1.114  up              2   master          Inherit    lcl     10.200.114.253
                                                             vip     10.200.114.254
xe-0/0/1.115  up              2   master          Inherit    lcl     10.200.115.253
                                                             vip     10.200.115.254

jncie@vQFX4> show vrrp summary

Interface     State       Group   VR state       VR Mode     Type    Address
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
xe-0/0/3.113   up          1    backup       Active     lcl    10.200.113.252
                                                         vip    10.200.113.254
xe-0/0/3.114   up          2    backup       Inherit    lcl    10.200.114.252
                                                         vip    10.200.114.254
xe-0/0/3.115   up          2    backup       Inherit    lcl    10.200.115.252
                                                         vip    10.200.115.254
```

As expected, vQFX5 is acting as the master for the virtual router and interfaces inside VLAN 114 and VLAN 115 follow the VRRP state of the interface inside VLAN 113.

```
jncie@vQFX5> show vrrp interface xe-0/0/1.113
Interface: xe-0/0/1.113, Interface index :567, Groups: 1, Active :1
  Interface VRRP PDU statistics
    Advertisement sent                        :1790
    Advertisement received                    :134
    Packets received                          :134
    No group match received                   :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received           :0
    Invalid VRRP TTL value received           :0
    Invalid VRRP version received             :0
    Invalid VRRP PDU type received            :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received            :0
    Invalid VRRP checksum received            :0

Physical interface: xe-0/0/1, Unit: 113, Vlan-id: 113, Address: 10.200.113.253/24
  Index: 567, SNMP ifIndex: 528, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 1, State: master, VRRP Mode: Active
  Priority: 250, Advertisement interval: 1, Authentication type: md5
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: yes, VIP count: 1, VIP: 10.200.113.254
  Advertisement Timer: 0.375s, Master router: 10.200.113.253
  Virtual router uptime: 00:28:04, Master router uptime: 00:13:25
  Virtual Mac: 00:00:5e:00:01:01
  Tracking: enabled
    Current priority: 250, Configured priority: 250
    Priority hold time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route               VRF name     Route state      Priority cost
      0.0.0.0/0           default      up                        250
  Group VRRP PDU statistics
    Advertisement sent                        :1790
    Advertisement received                    :134
  Group VRRP PDU error statistics
    Bad authentication Type received          :0
    Bad password received                     :0
    Bad MD5 digest received                   :0
    Bad advertisement timer received          :0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Bad VIP count received                  :0
    Bad VIPADDR received                    :0
  Group state transition statistics
    Idle to master transitions              :0
    Idle to backup transitions              :4
    Backup to master transitions            :4
    Master to backup transitions            :0

jncie@vQFX5> show vrrp interface xe-0/0/1.114

Physical interface: xe-0/0/1, Unit: 114, Vlan-id: 114, Address: 10.200.114.253/24
  Interface state: up, Group: 2, State: master
  VRRP Mode: Inherit, Active Node : xe-0/0/1.113, Active Group: 1
  Accept-data mode: no, VIP count: 1, VIP: 10.200.114.254
  Active router: 10.200.113.253
  Virtual Mac: 00:00:5e:00:01:02

jncie@vQFX4> show route advertising-protocol bgp 172.16.4.1 10.200.113.0/22 detail

inet.0: 49 destinations, 81 routes (31 active, 0 holddown, 36 hidden)
* 10.200.113.0/24 (4 entries, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
     AS path: 65004 65004 [65004] I

* 10.200.114.0/24 (3 entries, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
     AS path: 65004 65004 [65004] I

* 10.200.115.0/24 (3 entries, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
     AS path: 65004 65004 [65004] I

jncie@vQFX5> show route advertising-protocol bgp 172.16.6.1 10.200.113.0/22 detail

inet.0: 45 destinations, 71 routes (45 active, 0 holddown, 1 hidden)
* 10.200.113.0/24 (1 entry, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
     AS path: [65005] I

* 10.200.114.0/24 (1 entry, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
     AS path: [65005] I

* 10.200.115.0/24 (1 entry, 1 announced)
 BGP group jncie-fabric type External
     Nexthop: Self
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
     AS path: [65005] I

jncie@vMX3> show route 10.200.113.0/22 active-path

inet.0: 42 destinations, 85 routes (42 active, 0 holddown, 3 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.113.0/24    *[BGP/170] 00:16:11, localpref 100
                      AS path: 65005 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
10.200.114.0/24    *[BGP/170] 00:16:11, localpref 100
                      AS path: 65005 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
10.200.115.0/24    *[BGP/170] 00:16:11, localpref 100
                      AS path: 65005 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0

jncie@vMX3> show route 10.200.113.0/22 inactive-path

inet.0: 42 destinations, 85 routes (42 active, 0 holddown, 3 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.113.0/24     [BGP/170] 00:16:17, localpref 100
                      AS path: 65003 65002 65005 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:16:19, localpref 100
                      AS path: 65004 65004 65004 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
10.200.114.0/24     [BGP/170] 00:16:17, localpref 100
                      AS path: 65003 65002 65005 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:16:19, localpref 100
                      AS path: 65004 65004 65004 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
10.200.115.0/24     [BGP/170] 00:16:17, localpref 100
                      AS path: 65003 65002 65005 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:16:19, localpref 100
                      AS path: 65004 65004 65004 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0

jncie@vQFX5> show interfaces xe-0/0/1 extensive

...
  Hold-times     : Up 45000 ms, Down 0 ms
...
```

To verify that vQFX5 will release its mastership over the VRRP virtual-router, we simulate the loss of EBGP peers on the vQFX by deactivating the 'jncie-fabric' BGP group:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX5> configure

Entering configuration mode

{master:0}[edit]
jncie@vQFX5# deactivate protocols bgp group jncie-fabric

{master:0}[edit]
jncie@vQFX5# commit and-quit

configuration check succeeds
commit complete
Exiting configuration mode

{master:0}
jncie@vQFX5> show route 0/0 exact

inet.0: 20 destinations, 20 routes (19 active, 0 holddown, 1 hidden)

{master:0}
```

The vQFX no longer has any default route. Now we check VRRP:

```
jncie@vQFX5> show vrrp summary

Interface     State       Group   VR state       VR Mode    Type   Address
xe-0/0/1.113  up              1   idle            Active    lcl    10.200.113.253
                                                            vip    10.200.113.254
xe-0/0/1.114  up              2   idle            Inherit   lcl    10.200.114.253
                                                            vip    10.200.114.254
xe-0/0/1.115  up              2   idle            Inherit   lcl    10.200.115.253
                                                            vip    10.200.115.254
```

The priority has dropped. On vQFX4, we will be able to witness the following:

```
jncie@vQFX4> show vrrp summary

Interface     State       Group   VR state       VR Mode    Type   Address
xe-0/0/3.113  up              1   master          Active    lcl    10.200.113.252
                                                            vip    10.200.113.254
xe-0/0/3.114  up              2   master          Inherit   lcl    10.200.114.252
                                                            vip    10.200.114.254
xe-0/0/3.115  up              2   master          Inherit   lcl    10.200.115.252
                                                            vip    10.200.115.254
```

Before moving on to the next-task, we restore the situation wherein vQFX5 is VRRP master again:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX5> configure
Entering configuration mode

{master:0}[edit]
jncie@vQFX5# rollback 1
load complete

{master:0}[edit]
jncie@vQFX5# commit and-quit
configuration check succeeds
commit complete
Exiting configuration mode
```

```
{master:0}
jncie@vQFX5> show route 0/0 exact

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:00:05, localpref 100
                      AS path: 65001 65500 65355 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:00:02, localpref 100
                      AS path: 65002 65500 65355 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0

{master:0}
jncie@vQFX5> show vrrp summary

Interface     State       Group   VR state       VR Mode     Type   Address
xe-0/0/1.113  up              1   master          Active     lcl    10.200.113.253
                                                             vip    10.200.113.254
xe-0/0/1.114  up              2   master          Inherit    lcl    10.200.114.253
                                                             vip    10.200.114.254
xe-0/0/1.115  up              2   master          Inherit    lcl    10.200.115.253
                                                             vip    10.200.115.254
```

Task 2.7: IP-Fabric Customer Policy Configuration

- Make sure that customer 1 has a keepalive interval of 10 seconds.
- The session with customer 1 should be torn down for 5 minutes when more than 10 prefixes are announced.
- Make sure that flapping routes from customer-2 are dampened.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Solution

vQFX3:

```
set protocols bgp group customer-1 hold-time 30
set protocols bgp group customer-1 family inet unicast prefix-limit maximum 5
set protocols bgp group customer-1 family inet unicast prefix-limit teardown 80
set protocols bgp group customer-1 family inet unicast prefix-limit teardown idle-timeout 5

set protocols bgp group customer-2 damping
```

Verification

```
jncie@vQFX3> show bgp neighbor 172.16.200.2 | match "Options|Keepalive"

  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Options: <Preference HoldTime AuthKey LogUpDown AddressFamily PeerAS PrefixLimit Refresh>
  Keepalive Interval: 10          Group index: 5    Peer index: 0


jncie@vQFX3> show bgp neighbor 172.16.200.6 | match op

  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Options: <Preference AuthKey LogUpDown Damping PeerAS Refresh>
```

Task 2.8: IP-Fabric Customer Policy Configuration

- The customers should only receive a default route.
- Do not accept any default route.
- Prefixes advertised by the customers should be identifiable by a unique community.

Solution

vQFX3:

```
set protocols bgp group customer-1 import default
set protocols bgp group customer-1 import customer-1-import
set protocols bgp group customer-1 export default-only
set protocols bgp group customer-2 import default
set protocols bgp group customer-2 import customer-2-import
set protocols bgp group customer-2 export default-only
set protocols bgp group customer-5 import default
set protocols bgp group customer-5 import customer-5-import
set protocols bgp group customer-5 export default-only

set policy-options policy-statement default term 1 from route-filter 0.0.0.0/0 through
0.0.0.0/32
set policy-options policy-statement default term 1 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options policy-statement default-only term 1 from protocol bgp
set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options policy-statement customer-1-import term community then community add
customer-1
set policy-options policy-statement customer-2-import term community then community add
customer-2
set policy-options policy-statement customer-5-import term community then community add
customer-5

set policy-options community customer-1 members 65013:1
set policy-options community customer-2 members 65014:1
set policy-options community customer-3 members 65015:1
set policy-options community customer-4 members 65016:1
set policy-options community customer-5 members 65535:1
set policy-options community customer-6 members 65017:1
```

vQFX4:

```
set protocols bgp group customer-3 import default
set protocols bgp group customer-3 import customer-3-import
set protocols bgp group customer-3 export default-only
set protocols bgp group customer-4 import default
set protocols bgp group customer-4 import customer-4-import
set protocols bgp group customer-4 export default-only
set protocols bgp group customer-6 import customer-6-import
set protocols bgp group customer-6 export default-only

set policy-options policy-statement default term 1 from route-filter 0.0.0.0/0 through
0.0.0.0/32
set policy-options policy-statement default term 1 then reject

set policy-options policy-statement default-only term 1 from protocol bgp
set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options policy-statement customer-3-import term community then community add
customer-3
set policy-options policy-statement customer-4-import term 1 then community add customer-4
set policy-options policy-statement customer-6-import term 1 then community add customer-6

set policy-options community customer-1 members 65013:1
set policy-options community customer-2 members 65014:1
set policy-options community customer-3 members 65015:1
set policy-options community customer-4 members 65016:1
set policy-options community customer-5 members 65535:1
set policy-options community customer-6 members 65017:1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX5:

```
set protocols bgp group customer-4 import default
set protocols bgp group customer-4 import customer-4-import
set protocols bgp group customer-4 export default-only

set policy-options policy-statement default term 1 from route-filter 0.0.0.0/0 through
0.0.0.0/32
set policy-options policy-statement default term 1 then reject

set policy-options policy-statement default-only term 1 from protocol bgp
set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options policy-statement customer-4-import term 1 then community add customer-4

set policy-options community customer-1 members 65013:1
set policy-options community customer-2 members 65014:1
set policy-options community customer-3 members 65015:1
set policy-options community customer-4 members 65016:1
set policy-options community customer-5 members 65535:1
set policy-options community customer-6 members 65017:1
```

Verification

To verify the advertisement of a default route, we check the routes we advertised to all of the customers. The following output displays how this is done on vQFX3:

```
jncie@vQFX3> show route advertising-protocol bgp 172.16.200.2

inet.0: 56 destinations, 69 routes (56 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    65001 65500 65355 I

{master:0}
jncie@vQFX3> show route advertising-protocol bgp 172.16.200.6

inet.0: 56 destinations, 69 routes (56 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    65001 65500 65355 I

{master:0}
jncie@vQFX3> show route advertising-protocol bgp 172.16.200.22

inet.0: 56 destinations, 69 routes (56 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    65001 65500 65355 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To examine whether or not we are receiving a default route from any of the customers, we simply check the routing table. When we check the routing table for a default route, we know there should only be two routes present. Both of those routes should be originated from AS 65355:

```
jncie@vQFX3> show route 0/0 exact

inet.0: 56 destinations, 69 routes (56 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:25:52, localpref 100
                      AS path: 65001 65500 65355 I, validation-state: unverified
                    > to 172.16.2.1 via xe-0/0/2.0
                      to 172.16.3.1 via xe-0/0/4.0
                    [BGP/170] 00:14:31, localpref 100
                      AS path: 65002 65500 65355 I, validation-state: unverified
                    > to 172.16.3.1 via xe-0/0/4.0

jncie@vQFX4> show route 0/0 exact

inet.0: 49 destinations, 77 routes (31 active, 0 holddown, 36 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:26:09, localpref 100, from 172.16.4.1
                      AS path: 65001 65500 65355 I, validation-state: unverified
                      to 172.16.4.1 via xe-0/0/2.0
                    > to 172.16.5.1 via xe-0/0/4.0
                    [BGP/170] 00:14:49, localpref 100
                      AS path: 65002 65500 65355 I, validation-state: unverified
                    > to 172.16.5.1 via xe-0/0/4.0


jncie@vQFX5> show route 0/0 exact

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:26:28, localpref 100, from 172.16.6.1
                      AS path: 65001 65500 65355 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
                      to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:15:07, localpref 100
                      AS path: 65002 65500 65355 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
```

Both vQFX4 and vQFX5 show a number of hidden routes. Routes can be hidden for all sorts of reasons. For example, an illegal AS path, a next-hop that cannot be resolved, a damping policy or an import policy.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

On vQFX5, we can see that the 1 hidden route is the result of the import policy that rejects the default route:

```
jncie@vQFX5> show route 0/0 exact hidden detail

inet.0: 43 destinations, 69 routes (43 active, 0 holddown, 1 hidden)
0.0.0.0/0 (3 entries, 1 announced)
         BGP                  /-101
                Next hop type: Router, Next hop index: 1693
                Address: 0x9db3550
                Next-hop reference count: 5
                Source: 172.16.200.18
                Next hop: 172.16.200.18 via xe-0/0/1.111, selected
                Session Id: 0x0
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 65005 Peer AS: 65016
                Age: 3:45
                Validation State: unverified
                Task: BGP_65016.172.16.200.18
                AS path: 65016 65004 65002 65500 65355 I
                Localpref: 100
                Router ID: 172.16.200.14
                Hidden reason: rejected by import policy
```

Notice that checking the received routes will not display the route that an import policy has rejected. When we check the received route, we do see the number of hidden routes:

```
jncie@vQFX5> show route receive-protocol bgp 172.16.200.18

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.6.0/24             172.16.200.18                           65016 I
* 10.0.7.0/24             172.16.200.18                           65016 I
```

To see what type of route this is, we need to issue the same command with the keyword hidden:

```
jncie@vQFX5> show route receive-protocol bgp 172.16.200.18 hidden

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  0.0.0.0/0               172.16.200.18                           65016 65004 65002 65500
65355 I
```

The final objective in this task was to make sure we can identify the prefixes from every customer via a unique community. To verify our work, we examine the routing table with the newly configured communities like this:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX5> show route community-name customer-1

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/24         *[BGP/170] 00:24:47, localpref 100, from 172.16.7.1
                      AS path: 65002 65003 65013 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                     [BGP/170] 00:18:15, localpref 100
                      AS path: 65001 65003 65013 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0
10.0.2.0/24         *[BGP/170] 00:24:47, localpref 100
                      AS path: 65002 65003 65013 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
                      to 172.16.6.1 via xe-0/0/4.0
                     [BGP/170] 00:18:15, localpref 100
                      AS path: 65001 65003 65013 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0

jncie@vQFX5> show route community-name customer-2

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.3.0/24         *[BGP/170] 00:24:48, localpref 100, from 172.16.7.1
                      AS path: 65002 65003 65014 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                     [BGP/170] 00:18:16, localpref 100
                      AS path: 65001 65003 65014 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0
10.0.4.0/24         *[BGP/170] 00:24:48, localpref 100, from 172.16.7.1
                      AS path: 65002 65003 65014 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                     [BGP/170] 00:18:16, localpref 100
                      AS path: 65001 65003 65014 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0

jncie@vQFX5> show route community-name customer-3

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.5.0/24         *[BGP/170] 00:24:52, localpref 100, from 172.16.6.1
                      AS path: 65001 65004 65015 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
                      to 172.16.6.1 via xe-0/0/4.0
                     [BGP/170] 00:24:49, localpref 100
                      AS path: 65002 65004 65015 I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    > to 172.16.7.1 via xe-0/0/2.0

jncie@vQFX5> show route community-name customer-4

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.6.0/24         *[BGP/170] 01:00:05, localpref 250
                      AS path: 65016 I, validation-state: unverified
                    > to 172.16.200.18 via xe-0/0/1.111
                    [BGP/170] 00:24:53, localpref 100
                      AS path: 65001 65004 65016 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:24:50, localpref 100
                      AS path: 65002 65004 65016 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
10.0.7.0/24         *[BGP/170] 01:00:05, localpref 250
                      AS path: 65016 I, validation-state: unverified
                    > to 172.16.200.18 via xe-0/0/1.111
                    [BGP/170] 00:24:53, localpref 100
                      AS path: 65001 65004 65016 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:24:50, localpref 100
                      AS path: 65002 65004 65016 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0

jncie@vQFX5> show route community-name customer-5

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.8.0/24         *[BGP/170] 00:24:52, localpref 100, from 172.16.7.1
                      AS path: 65002 65003 65535 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:18:20, localpref 100
                      AS path: 65001 65003 65535 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0
10.0.9.0/24         *[BGP/170] 00:24:52, localpref 100
                      AS path: 65002 65003 65535 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
                      to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:18:20, localpref 100
                      AS path: 65001 65003 65535 I, validation-state: unverified
                    > to 172.16.6.1 via xe-0/0/4.0

jncie@vQFX5> show route community-name customer-6

inet.0: 45 destinations, 75 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
10.0.10.0/24       *[BGP/170] 00:04:40, localpref 100
                      AS path: 65001 65004 65017 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:04:40, localpref 100
                      AS path: 65002 65004 65017 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
10.0.11.0/24       *[BGP/170] 00:04:40, localpref 100
                      AS path: 65001 65004 65017 I, validation-state: unverified
                      to 172.16.7.1 via xe-0/0/2.0
                    > to 172.16.6.1 via xe-0/0/4.0
                    [BGP/170] 00:04:40, localpref 100
                      AS path: 65002 65004 65017 I, validation-state: unverified
                    > to 172.16.7.1 via xe-0/0/2.0
```

Task 2.9: ISP Policy Configuration

- Advertise a single summary IPv4 route (10.0.0.0/8) that aggregates the datacenter range to the ISPs.
- The ISP only allows you to use your own AS.
- Do not accept any other route then the default route from the ISP.

Solution

vMX3:

```
set routing-options aggregate route 10.0.0.0/8 policy contributor

set protocols bgp group isp export aggregate-only
set protocols bgp group isp remove-private
set protocols bgp group isp import default-only

set policy-options policy-statement aggregate-only term 1 from protocol aggregate
set policy-options policy-statement aggregate-only term 1 from route-filter 10.0.0.0/8 exact
set policy-options policy-statement aggregate-only term 1 then accept
set policy-options policy-statement aggregate-only term 2 then reject

set policy-options policy-statement default-only term 1 from protocol bgp
set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options policy-statement contributor term direct from protocol direct
set policy-options policy-statement contributor term direct then accept
set policy-options policy-statement contributor term reject then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set routing-options aggregate route 10.0.0.0/8

set protocols bgp group isp export aggregate-only
set protocols bgp group isp remove-private
set protocols bgp group isp import default-only

set policy-options policy-statement aggregate-only term 1 from protocol aggregate
set policy-options policy-statement aggregate-only term 1 from route-filter 10.0.0.0/8 exact
set policy-options policy-statement aggregate-only term 1 then accept
set policy-options policy-statement aggregate-only term 2 then reject
```

```
set policy-options policy-statement default-only term 1 from protocol bgp
set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options policy-statement contributor term direct from protocol direct
set policy-options policy-statement contributor term direct then accept
set policy-options policy-statement contributor term reject then reject
```

Verification

We can use either a static route or an aggregate route in this task. Keep in mind that on a JNCIE exam, you are not allowed to use a static route anywhere unless this is specifically mentioned.

Since the use of a static can cause a loss of points, we turn to the use of an aggregate route in this scenario. The ISP in this scenario only allows you to use AS65001. When we configure and advertise an aggregate, we can observe the following:

Note: The output was taken prior to configuration changes.

```
jncie@vMX3> show route advertising-protocol bgp 10.0.0.1

inet.0: 42 destinations, 79 routes (42 active, 0 holddown, 2 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.0.0/8              Self                                    65001 {65003 65004 65005
65013 65014 65015 65016 65017 65535} I
```

We do not want to send all the other autonomous systems we use when we advertise the route to our ISP. When we configure an aggregate, so-called contributor routes not only determine whether or not the aggregate is active. The contributor routes can also add to the properties of the aggregate. To see the contributors for an aggregate route, issue the following command:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX4# show route protocol aggregate extensive

inet.0: 42 destinations, 55 routes (42 active, 0 holddown, 0 hidden)
10.0.0.0/8 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.0.0.0/8 -> {}
Page 0 idx 1, (group isp type External) Type 1 val 0xb22189c (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: 65002 {65003 65004 65005 65013 65014 65015 65016 65017 65535} I (LocalAgg)
     Communities:
Path 10.0.0.0 Vector len 4.  Val: 1
        *Aggregate Preference: 130
         Next hop type: Reject, Next hop index: 0
         Address: 0x9f48f84
         Next-hop reference count: 2
         State: <Active Int Ext>
         Local AS: 65002
         Age: 15:19
         Validation State: unverified
         Task: Aggregate
         Announcement bits (2): 0-KRT 1-BGP_RT_Background
         AS path: {65003 65004 65005 65013 65014 65015 65016 65017 65535} I (LocalAgg)
         Flags:                  Depth: 0        Active
         AS path list:
         AS path: 65005 65016 I Refcount: 2
         AS path: 65003 65014 I Refcount: 2
         AS path: 65005 I Refcount: 3
         AS path: 65003 65013 I Refcount: 2
         AS path: I Refcount: 2
         AS path: 65004 65015 I Refcount: 1
         AS path: 65003 I Refcount: 11
         AS path: 65003 65535 I Refcount: 2
         AS path: 65004 65017 I Refcount: 2
         Contributing Routes (27):
         10.0.0.4/30 proto Direct
         10.10.20.0/24 proto Direct
         10.0.1.0/24 proto BGP
         10.0.2.0/24 proto BGP
         10.0.3.0/24 proto BGP
         10.0.4.0/24 proto BGP
         10.0.5.0/24 proto BGP
         10.0.6.0/24 proto BGP
         10.0.7.0/24 proto BGP
         10.0.8.0/24 proto BGP
         10.0.9.0/24 proto BGP
         10.0.10.0/24 proto BGP
         10.0.11.0/24 proto BGP
         10.200.100.0/24 proto BGP
         10.200.101.0/24 proto BGP
         10.200.102.0/24 proto BGP
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
10.200.103.0/24 proto BGP
10.200.104.0/24 proto BGP
10.200.105.0/24 proto BGP
10.200.106.0/24 proto BGP
10.200.107.0/24 proto BGP
10.200.108.0/24 proto BGP
10.200.109.0/24 proto BGP
10.200.110.0/24 proto BGP
10.200.113.0/24 proto BGP
10.200.114.0/24 proto BGP
10.200.115.0/24 proto BGP
```

To be able to advertise the route without adding any information from the contributors to the AS-path, we configure a policy that rejects the BGP routes from being able to contribute. After applying the policy, we can see the following:

```
jncie@vMX4# show route protocol aggregate extensive

inet.0: 42 destinations, 55 routes (42 active, 0 holddown, 0 hidden)
10.0.0.0/8 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.0.0.0/8 -> {}
Page 0 idx 1, (group isp type External) Type 1 val 0xb22189c (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [65500] I (LocalAgg)
     Communities:
Path 10.0.0.0 Vector len 4.  Val: 1
        *Aggregate Preference: 130
                Next hop type: Reject, Next hop index: 0
                Address: 0x9f48f84
                Next-hop reference count: 2
                State: <Active Int Ext>
                Local AS: 65002
                Age: 17
                Validation State: unverified
                Task: Aggregate
                Announcement bits (2): 0-KRT 1-BGP_RT_Background
                AS path: I (LocalAgg)
                Flags:                  Depth: 0        Active
                AS path list:
                AS path: I Refcount: 2
                Contributing Routes (2):
                        10.0.0.4/30 proto Direct
                        10.10.20.0/24 proto Direct
```

Only the direct routes are adding to the aggregate. When we check the advertised routes now, we can see the following:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show route advertising-protocol bgp 10.0.0.1

inet.0: 42 destinations, 78 routes (42 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.0.0/8              Self                                    I

jncie@vMX4> show route advertising-protocol bgp 10.0.0.5

inet.0: 42 destinations, 54 routes (42 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.0.0/8              Self                                    I
```

To verify our import policy:

```
jncie@vMX3> show route receive-protocol bgp 10.0.0.1

inet.0: 42 destinations, 80 routes (42 active, 0 holddown, 3 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               10.0.0.1                                65355 I

jncie@vMX4> show route receive-protocol bgp 10.0.0.5

inet.0: 42 destinations, 56 routes (42 active, 0 holddown, 2 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               10.0.0.5                                65355 I
```

Note: as an alternative to the configuration of a policy for the aggregate, we could also consider the 'remove-private' knob. Without the 'contributor' policy applied to the aggregate, we could observe the following:

Note: The output was taken prior to configuration changes being applied:

```
jncie@vMX3> show route advertising-protocol bgp 10.0.0.1
inet.0: 42 destinations, 79 routes (42 active, 0 holddown, 2 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.0.0/8              Self                                    65001 {65003 65004 65005
65013 65014 65015 65016 65017 65535} I
```

To make sure that the AS-path only shows the locally configured AS, we can use the 'remove-private' knob as follows:

```
jncie@vMX3> configure
Entering configuration mode

[edit]
jncie@vMX3# set protocols bgp group isp remove-private
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
[edit]
jncie@vMX3# commit and-quit
commit complete
Exiting configuration mode

jncie@vMX3> show route advertising-protocol bgp 10.0.0.1

inet.0: 42 destinations, 79 routes (42 active, 0 holddown, 2 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.0.0/8              Self                                    I
```

Part 2: IBGP-based Clos IP Fabric

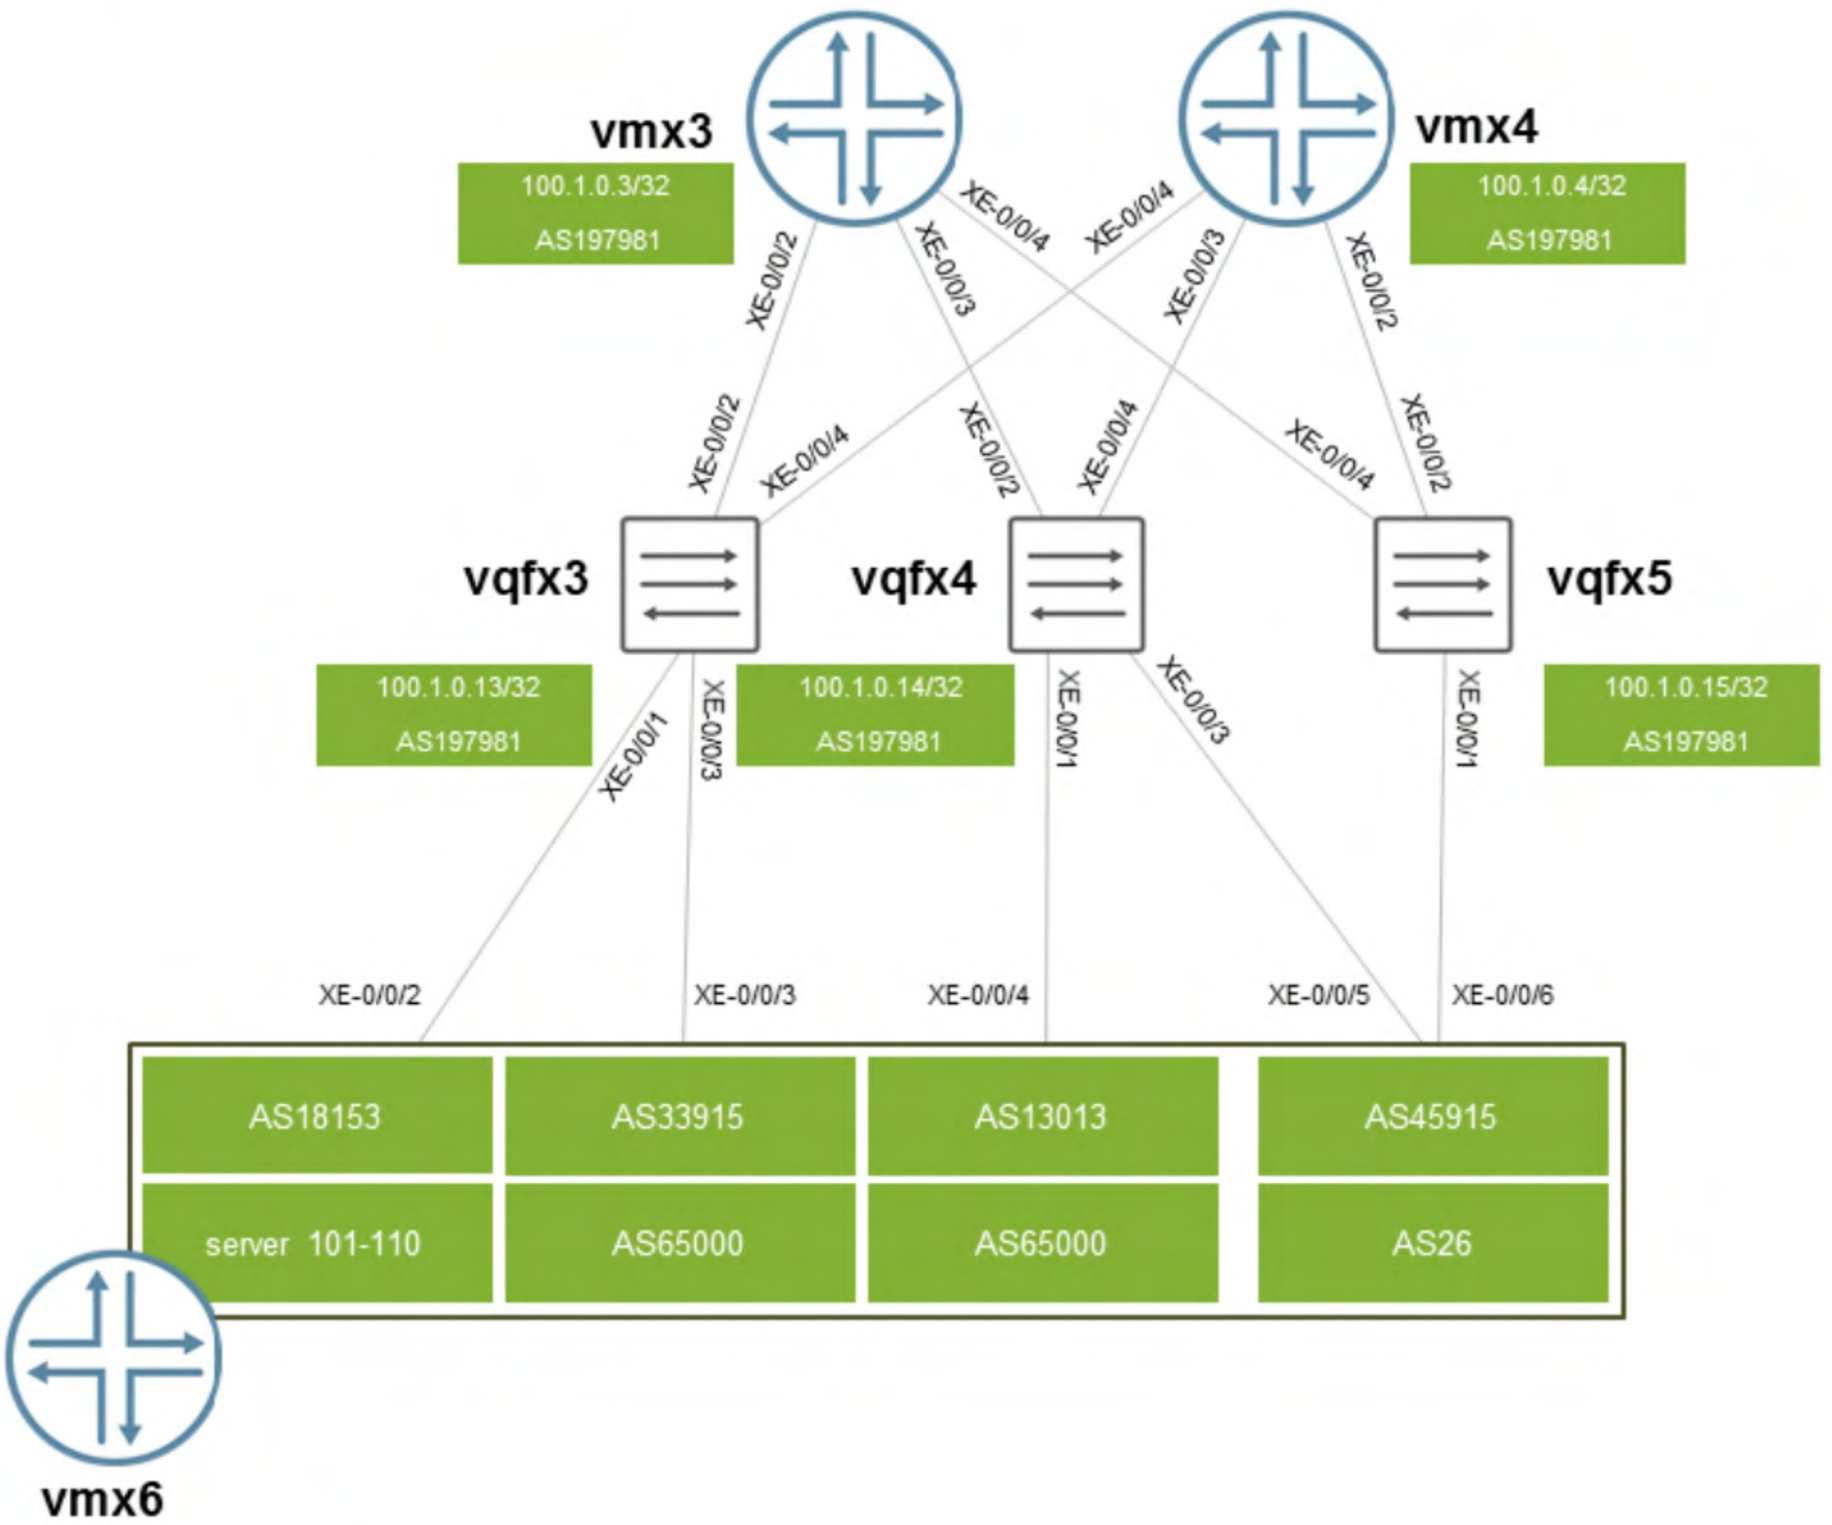

Figure 18 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

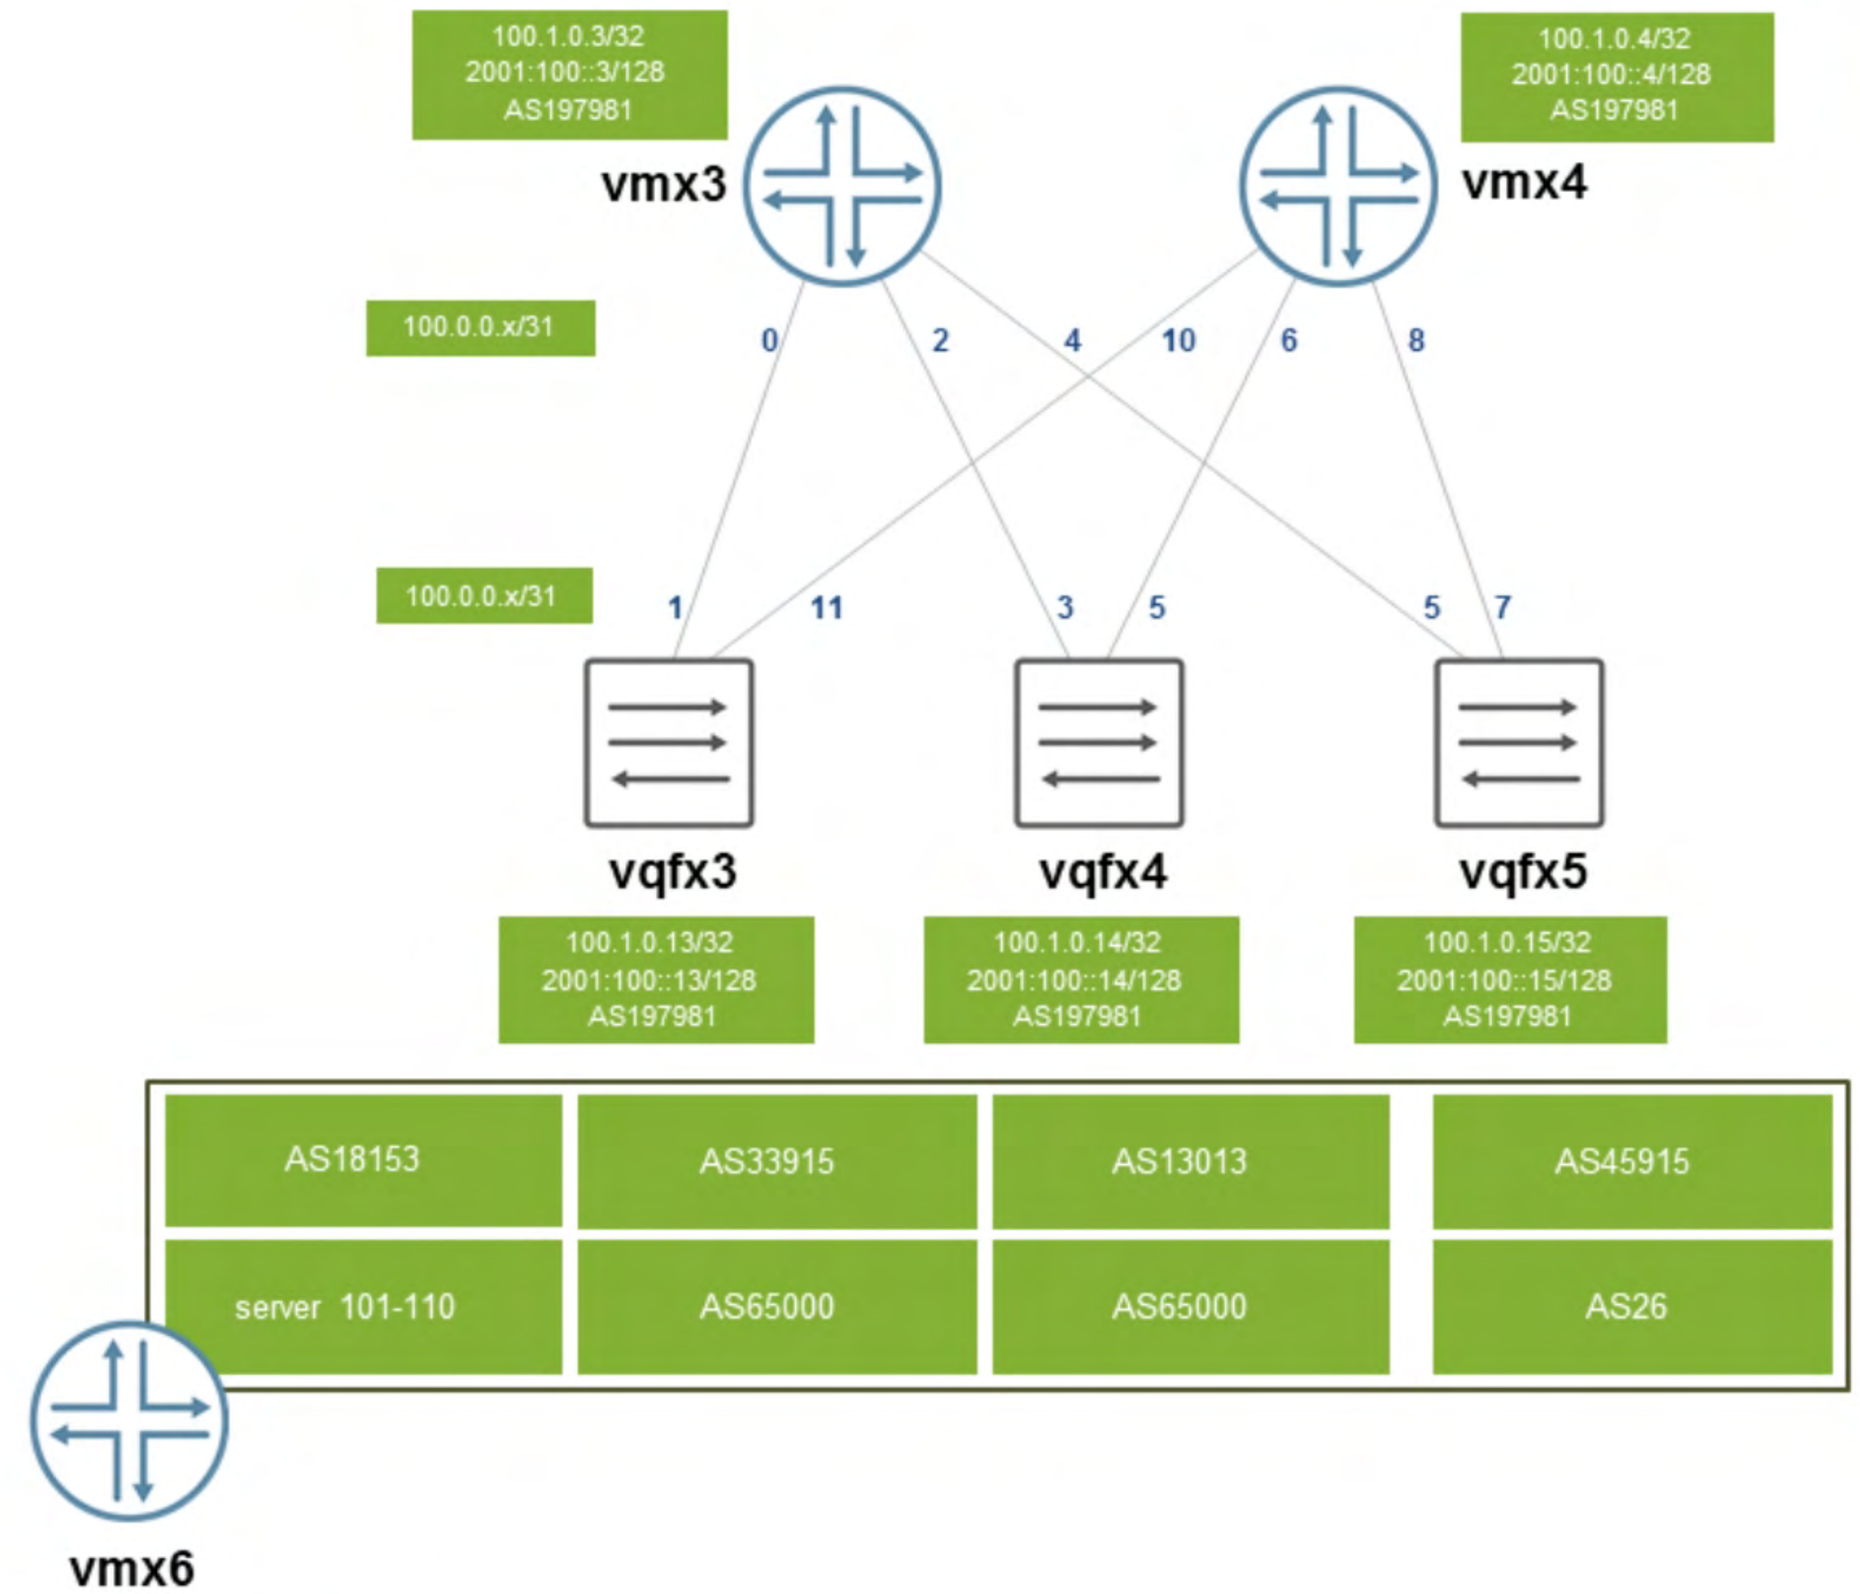

Figure 2 - Logical topology

Note: starting topology is configured with IP-addressing only.

Task 2.10: IGP Configuration

- Place the nodes that make up the IP-Fabric into a single area OSPFv3.

Solution

vMX3:

```
set protocols ospf3 area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/3.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set protocols ospf3 area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/3.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

vQFX3:

```
set protocols ospf3 area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

vQFX4:

```
set protocols ospf3 area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

vQFX5:

```
set protocols ospf3 area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
```

Verification

The easiest way to verify the OSPFv3 neighbor relationship establishment is by examining them from the spine nodes. First, we check OSPFv3 from vMX3:

```
jncie@vMX3> show ospf3 neighbor

ID              Interface              State     Pri   Dead
100.1.0.13      xe-0/0/2.0             Full      128     39
  Neighbor-address fe80::205:86ff:fe71:9f0b
100.1.0.14      xe-0/0/3.0             Full      128     38
  Neighbor-address fe80::205:86ff:fe71:3f0b
100.1.0.15      xe-0/0/4.0             Full      128     38
  Neighbor-address fe80::205:86ff:fe71:813
```

Another useful verification command is the following:

```
jncie@vMX3> show ospf3 interface detail

Interface          State   Area            DR ID           BDR ID          Nbrs
lo0.0              DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Address fe80::20c:290f:fcb1:7e18, Prefix-length 128
  OSPF3-Intf-index 1, Type LAN, MTU 65535, Cost 0
  Adj count: 0, Router LSA ID: -, Passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None
xe-0/0/2.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0          1
  Address fe80::20c:29ff:fecf:e172, Prefix-length 64
  OSPF3-Intf-index 2, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None
xe-0/0/3.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0          1
  Address fe80::20c:29ff:fecf:e17c, Prefix-length 64
  OSPF3-Intf-index 3, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None
xe-0/0/4.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0          1
  Address fe80::20c:29ff:fecf:e186, Prefix-length 64
  OSPF3-Intf-index 4, Type P2P, MTU 1500, Cost 1
  Adj count: 1, Router LSA ID: 0
  Hello 10, Dead 40, ReXmit 5, Not Stub
  Protection type: None
```

The command informs us on all the configuration parameters applied to the interface. In addition to this, we can also see if an OSPF neighbor adjacency has been established across the interface.

The following command can be used to examine the LSAs stored inside the OSPF database:

```
jncie@vMX3> show ospf3 database

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          100.1.0.3       0x80000003   27  0x679f  72
Router      0.0.0.0          100.1.0.4       0x80000003   34  0x43c1  72
Router      0.0.0.0          100.1.0.13      0x80000002   40  0xf4a8  56
Router      0.0.0.0          100.1.0.14      0x80000002   35  0xdebd  56
Router      0.0.0.0          100.1.0.15      0x80000002   28  0xd8c2  56
IntraArPfx *0.0.0.1          100.1.0.3       0x80000001   62  0xa2e2  52
IntraArPfx  0.0.0.1          100.1.0.4       0x80000001   51  0xccb5  52
IntraArPfx  0.0.0.1          100.1.0.13      0x80000001   45  0x154c  52
IntraArPfx  0.0.0.1          100.1.0.14      0x80000001   42  0x3f1f  52
IntraArPfx  0.0.0.1          100.1.0.15      0x80000001   35  0x69f1  52

    OSPF3 Link-Local database, interface xe-0/0/2.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          100.1.0.3       0x80000001   62  0x80d1  44
Link        0.0.0.2          100.1.0.13      0x80000001   45  0x55a4  44

    OSPF3 Link-Local database, interface xe-0/0/3.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Link        *0.0.0.3          100.1.0.3         0x80000001    62  0x53f3  44
Link         0.0.0.2          100.1.0.14        0x80000001    40  0x67f1  44

    OSPF3 Link-Local database, interface xe-0/0/4.0 Area 0.0.0.0
 Type        ID               Adv Rtr           Seq         Age  Cksum  Len
Link        *0.0.0.4          100.1.0.3         0x80000001    62  0x2616  44
Link         0.0.0.3          100.1.0.15        0x80000001    33  0x8006  44
```

Notice the router LSAs. We can see that every device in our topology have advertised a router LSA.

We then move to the other spine and check the adjacencies:

```
jncie@vMX4> show ospf3 neighbor
ID              Interface             State     Pri   Dead
100.1.0.15      xe-0/0/2.0            Full      128     34
  Neighbor-address fe80::205:86ff:fe71:80b
100.1.0.14      xe-0/0/3.0            Full      128     33
  Neighbor-address fe80::205:86ff:fe71:3f13
100.1.0.13      xe-0/0/4.0            Full      128     34
  Neighbor-address fe80::205:86ff:fe71:9f13
```

The last thing we need to check is whether or not the loopback IPv6 addresses are advertised using OSPFv3:

```
jncie@vMX3> show route table inet6.0 protocol ospf3

inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::4/128    *[OSPF3/10] 00:02:55, metric 2
                    > to fe80::205:86ff:fe71:9f0b via xe-0/0/2.0
                      to fe80::205:86ff:fe71:3f0b via xe-0/0/3.0
                      to fe80::205:86ff:fe71:813 via xe-0/0/4.0
2001:100::13/128   *[OSPF3/10] 00:03:50, metric 1
                    > to fe80::205:86ff:fe71:9f0b via xe-0/0/2.0
2001:100::14/128   *[OSPF3/10] 00:03:45, metric 1
                    > to fe80::205:86ff:fe71:3f0b via xe-0/0/3.0
2001:100::15/128   *[OSPF3/10] 00:03:40, metric 1
                    > to fe80::205:86ff:fe71:813 via xe-0/0/4.0
ff02::5/128        *[OSPF3/10] 00:04:17, metric 1
                      MultiRecv
```

To see if vMX3 advertised the loopback IPv6 address, we can check this from the vMX3 router itself:

```
jncie@vMX3> show ospf3 database advertising-router self detail

    OSPF3 database, Area 0.0.0.0
 Type        ID               Adv Rtr           Seq         Age  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Router     *0.0.0.0          100.1.0.3        0x80000003   280  0x679f  72
  bits 0x0, Options 0x33
  Type PointToPoint (1), Metric 1
      Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 100.1.0.13
  Type PointToPoint (1), Metric 1
      Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 100.1.0.14
  Type PointToPoint (1), Metric 1
      Loc-If-Id 4, Nbr-If-Id 3, Nbr-Rtr-Id 100.1.0.15
  Type: PointToPoint, Node ID: 100.1.0.15, Metric: 1, Bidirectional
  Type: PointToPoint, Node ID: 100.1.0.14, Metric: 1, Bidirectional
  Type: PointToPoint, Node ID: 100.1.0.13, Metric: 1, Bidirectional
IntraArPfx *0.0.0.1          100.1.0.3        0x80000001   315  0xa2e2  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 100.1.0.3
  Prefix-count 1
  Prefix 2001:100::3/128
    Prefix-options 0x2, Metric 0

    OSPF3 Link-Local database, interface xe-0/0/2.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          100.1.0.3        0x80000001   315  0x80d1  44
  fe80::20c:29ff:fecf:e172
  Options 0x33, Priority 128
  Prefix-count 0

    OSPF3 Link-Local database, interface xe-0/0/3.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.3          100.1.0.3        0x80000001   315  0x53f3  44
  fe80::20c:29ff:fecf:e17c
  Options 0x33, Priority 128
  Prefix-count 0

    OSPF3 Link-Local database, interface xe-0/0/4.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.4          100.1.0.3        0x80000001   315  0x2616  44
  fe80::20c:29ff:fecf:e186
  Options 0x33, Priority 128
  Prefix-count 0
```

Alternatively, we can simply examine the routing table on another router:

```
jncie@vMX4> show route protocol ospf3

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

inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::3/128    *[OSPF3/10] 00:04:22, metric 2
                    > to fe80::205:86ff:fe71:80b via xe-0/0/2.0
                      to fe80::205:86ff:fe71:3f13 via xe-0/0/3.0
                      to fe80::205:86ff:fe71:9f13 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
2001:100::13/128   *[OSPF3/10] 00:05:20, metric 1
                    > to fe80::205:86ff:fe71:9f13 via xe-0/0/4.0
2001:100::14/128   *[OSPF3/10] 00:05:15, metric 1
                    > to fe80::205:86ff:fe71:3f13 via xe-0/0/3.0
2001:100::15/128   *[OSPF3/10] 00:04:22, metric 1
                    > to fe80::205:86ff:fe71:80b via xe-0/0/2.0
ff02::5/128        *[OSPF3/10] 00:05:30, metric 1
                      MultiRecv
```

Task 2.11: IGP Configuration

- Links actively running OSPFv3 should be active for IPv4 as well.
- Have OSPFv3 advertise the IPv4 loopback address on every node.

Solution

vMX3:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/3.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
```

vMX4:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/3.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
```

vQFX3:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
```

vQFX4:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
```

vQFX5:

```
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface lo0.0 passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

Verifying the OSPFv3 neighbor establishment for the IPv4-unicast realm:

```
jncie@vMX3> show ospf3 neighbor realm ipv4-unicast

ID              Interface             State     Pri   Dead
100.1.0.13      xe-0/0/2.0            Full      128     37
  Neighbor-address fe80::205:86ff:fe71:9f0b
100.1.0.14      xe-0/0/3.0            Full      128     32
  Neighbor-address fe80::205:86ff:fe71:3f0b
100.1.0.15      xe-0/0/4.0            Full      128     39
  Neighbor-address fe80::205:86ff:fe71:813
```

Verifying the routing-table:

```
jncie@vMX3> show route table inet.0

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.20.0/24      *[Direct/0] 5d 18:04:43
                    > via fxp0.0
10.10.20.3/32      *[Local/0] 5d 18:04:43
                       Local via fxp0.0
100.0.0.0/31       *[Direct/0] 00:35:54
                    > via xe-0/0/2.0
100.0.0.0/32       *[Local/0] 00:35:54
                       Local via xe-0/0/2.0
100.0.0.2/31       *[Direct/0] 01:04:52
                    > via xe-0/0/3.0
100.0.0.2/32       *[Local/0] 01:04:52
                       Local via xe-0/0/3.0
100.0.0.4/31       *[Direct/0] 00:36:06
                    > via xe-0/0/4.0
100.0.0.4/32       *[Local/0] 00:36:06
                       Local via xe-0/0/4.0
100.0.0.6/31       *[OSPF3/10] 00:01:26, metric 2
                    > to 100.0.0.5 via xe-0/0/4.0
100.0.0.8/31       *[OSPF3/10] 00:01:31, metric 2
                    > to 100.0.0.3 via xe-0/0/3.0
100.0.0.10/31      *[OSPF3/10] 00:01:41, metric 2
                    > to 100.0.0.1 via xe-0/0/2.0
100.1.0.3/32       *[Direct/0] 4d 02:47:54
                    > via lo0.0
100.1.0.4/32       *[OSPF3/10] 00:01:21, metric 2
                    > to 100.0.0.1 via xe-0/0/2.0
                      to 100.0.0.3 via xe-0/0/3.0
                      to 100.0.0.5 via xe-0/0/4.0
100.1.0.13/32      *[OSPF3/10] 00:01:41, metric 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    > to 100.0.0.1 via xe-0/0/2.0
100.1.0.14/32       *[OSPF3/10] 00:01:31, metric 1
                    > to 100.0.0.3 via xe-0/0/3.0
100.1.0.15/32       *[OSPF3/10] 00:01:26, metric 1
                    > to 100.0.0.5 via xe-0/0/4.0

jncie@vMX4> show ospf3 neighbor realm ipv4-unicast

ID              Interface             State     Pri   Dead
100.1.0.15      xe-0/0/2.0            Full      128     34
  Neighbor-address fe80::205:86ff:fe71:80b
100.1.0.14      xe-0/0/3.0            Full      128     34
  Neighbor-address fe80::205:86ff:fe71:3f13
100.1.0.13      xe-0/0/4.0            Full      128     32
  Neighbor-address fe80::205:86ff:fe71:9f13

jncie@vMX4> show route table inet.0

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.20.0/24       *[Direct/0] 4d 02:08:09
                     > via fxp0.0
10.10.20.4/32       *[Local/0] 4d 02:08:09
                       Local via fxp0.0
100.0.0.0/31        *[OSPF3/10] 00:01:57, metric 2
                     > to 100.0.0.11 via xe-0/0/4.0
100.0.0.2/31        *[OSPF3/10] 00:01:47, metric 2
                     > to 100.0.0.9 via xe-0/0/3.0
100.0.0.4/31        *[OSPF3/10] 00:01:52, metric 2
                     > to 100.0.0.7 via xe-0/0/2.0
100.0.0.6/31        *[Direct/0] 00:11:03
                     > via xe-0/0/2.0
100.0.0.6/32        *[Local/0] 00:11:03
                       Local via xe-0/0/2.0
100.0.0.8/31        *[Direct/0] 00:57:20
                     > via xe-0/0/3.0
100.0.0.8/32        *[Local/0] 00:57:20
                       Local via xe-0/0/3.0
100.0.0.10/31       *[Direct/0] 00:42:43
                     > via xe-0/0/4.0
100.0.0.10/32       *[Local/0] 00:42:43
                       Local via xe-0/0/4.0
100.1.0.3/32        *[OSPF3/10] 00:01:47, metric 2
                       to 100.0.0.7 via xe-0/0/2.0
                     > to 100.0.0.9 via xe-0/0/3.0
                       to 100.0.0.11 via xe-0/0/4.0
100.1.0.4/32        *[Direct/0] 4d 02:08:09
                     > via lo0.0
100.1.0.13/32       *[OSPF3/10] 00:01:57, metric 1
                     > to 100.0.0.11 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
100.1.0.14/32      *[OSPF3/10] 00:01:47, metric 1
                    > to 100.0.0.9 via xe-0/0/3.0
100.1.0.15/32      *[OSPF3/10] 00:01:52, metric 1
                    > to 100.0.0.7 via xe-0/0/2.0
```

Task 2.12: IGP Configuration

- Ensure the utilization of equal-cost links.
- Make sure OSPFv3 calculates path cost using a reference bandwidth of 100Gbit.

Solution

On every device:

```
set protocols ospf3 realm ipv4-unicast reference-bandwidth 100g
set protocols ospf3 reference-bandwidth 100g

set routing-options forwarding-table export lbpp
set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

Verification

Before activating load-balancing:

```
jncie@vQFX3> show route forwarding-table destination 100.1.0.15

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
100.1.0.15/32      user     0 100.0.0.10         ucst     1703     7 xe-0/0/4.0

jncie@vQFX3> show route forwarding-table destination 2001:100::15

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
2001:100::15/128   user     0 fe80::20c:29ff:fecf:e172
                                                 ucst     1715     6 xe-0/0/2.0
```

After enabling load-balancing:

```
jncie@vQFX3> show route forwarding-table destination 100.1.0.15

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
100.1.0.15/32      user     0                    ulst   131071     3
                              100.0.0.0          ucst     1692     6 xe-0/0/2.0
                              100.0.0.10         ucst     1703     6 xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show route forwarding-table destination 2001:100::15

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
2001:100::15/128   user     0                    ulst   131070     3
                            fe80::20c:29ff:fecf:e172
                                                 ucst     1715     4 xe-0/0/2.0
                            fe80::20c:29ff:fe7d:2020
                                                 ucst     1716     4 xe-0/0/4.0
```

Note that it is important to understand on what node you are verifying the configuration. For instance, the spine routers only have an equal-cost path to each other. Therefore, if you are verifying the load-balancing configuration from either vMX3 or vXM4, you would need to check the destination of the other spine instead of one of the leaf devices. Observe the following:

```
jncie@vMX3> show route forwarding-table destination 100.1.0.13

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
100.1.0.13/32      user     0 100.0.0.1          ucst      513     5 xe-0/0/2.0
```

Here you can see that vMX3 only has one route towards the leaf. Now observe the following:

```
jncie@vMX3> show route forwarding-table destination 100.1.0.4

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
100.1.0.4/32       user     0                    ulst  1048574     2
                            100.0.0.1            ucst      513     5 xe-0/0/2.0
                            100.0.0.3            ucst      601     5 xe-0/0/3.0
                            100.0.0.5            ucst      575     5 xe-0/0/4.0
```

Here you can see that vMX3 has two equal cost paths towards vMX4.

To verify the reference bandwidth, we can check the metric that is assigned for certain routes before altering the reference bandwidth:

```
jncie@vQFX3> show route protocol ospf3 100.1.0.15

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.1.0.15/32      *[OSPF3/10] 00:04:15, metric 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      to 100.0.0.0 via xe-0/0/2.0
                    > to 100.0.0.10 via xe-0/0/4.0

jncie@vQFX3> show route protocol ospf3 2001:100::15

inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::15/128   *[OSPF3/10] 00:13:21, metric 2
                    > to fe80::20c:29ff:fecf:e172 via xe-0/0/2.0
                      to fe80::20c:29ff:fe7d:2020 via xe-0/0/4.0
```

And afterwards:

```
jncie@vQFX3> show route protocol ospf3 100.1.0.15

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.1.0.15/32      *[OSPF3/10] 00:04:54, metric 20
                      to 100.0.0.0 via xe-0/0/2.0
                    > to 100.0.0.10 via xe-0/0/4.0

jncie@vQFX3> show route protocol ospf3 2001:100::15

inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::15/128   *[OSPF3/10] 00:05:05, metric 20
                    > to fe80::20c:29ff:fecf:e172 via xe-0/0/2.0
                      to fe80::20c:29ff:fe7d:2020 via xe-0/0/4.0
```

Task 2.13: IPv4 IBGP Configuration

- Configure an IBGP session between the vMX routers and the vQFX switches.
- vQFX switches are not allowed to peer with each other.
- Authenticate all IBGP sessions using password ‘jncie’.

Solution

Every vQFX will have IBGP sessions with both of the vMX devices. By default, IBGP learned routes are not advertised across IBGP sessions. In order to make sure that the vQFX devices learn all routes, we configure the vMX devices as route reflectors. By doing so, we enable the vMX routers to advertise the IBGP learned routes across the IBGP sessions with the vQFX devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set protocols bgp group jncie type internal
set protocols bgp group jncie local-address 100.1.0.3
set protocols bgp group jncie authentication-key jncie
set protocols bgp group jncie cluster 100.1.0.3
set protocols bgp group jncie neighbor 100.1.0.4
set protocols bgp group jncie neighbor 100.1.0.13
set protocols bgp group jncie neighbor 100.1.0.14
set protocols bgp group jncie neighbor 100.1.0.15

set routing-options autonomous-system 197981
```

vMX4:

```
set protocols bgp group jncie type internal
set protocols bgp group jncie local-address 100.1.0.4
set protocols bgp group jncie authentication-key jncie
set protocols bgp group jncie cluster 100.1.0.4
set protocols bgp group jncie neighbor 100.1.0.3
set protocols bgp group jncie neighbor 100.1.0.13
set protocols bgp group jncie neighbor 100.1.0.14
set protocols bgp group jncie neighbor 100.1.0.15

set routing-options autonomous-system 197981
```

vQFX3:

```
set protocols bgp group jncie type internal
set protocols bgp group jncie local-address 100.1.0.13
set protocols bgp group jncie authentication-key jncie
set protocols bgp group jncie neighbor 100.1.0.3
set protocols bgp group jncie neighbor 100.1.0.4
set routing-options autonomous-system 197981
```

vQFX4:

```
set protocols bgp group jncie type internal
set protocols bgp group jncie local-address 100.1.0.14
set protocols bgp group jncie authentication-key jncie
set protocols bgp group jncie neighbor 100.1.0.3
set protocols bgp group jncie neighbor 100.1.0.4

set routing-options autonomous-system 197981
```

vQFX5:

```
set protocols bgp group jncie type internal
set protocols bgp group jncie local-address 100.1.0.15
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group jncie authentication-key jncie
set protocols bgp group jncie neighbor 100.1.0.3
set protocols bgp group jncie neighbor 100.1.0.4

set routing-options autonomous-system 197981
```

Verification

First, we verify the session establishment:

```
jncie@vMX3> show bgp summary | refresh 2
---(refreshed at 2017-04-08 07:56:43 UTC)---
...
---(refreshed at 2017-04-08 07:56:45 UTC)---
...
---(refreshed at 2017-04-08 07:56:47 UTC)---
Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
100.1.0.4            197981          3          2       0       0           0 0/0/0/0
0/0/0/0
100.1.0.13           197981          6          3       0       0        1:07 0/0/0/0
0/0/0/0
100.1.0.14           197981          6          4       0       0        1:07 0/0/0/0
0/0/0/0
100.1.0.15           197981          6          3       0       0        1:05 0/0/0/0
0/0/0/0
---(*more 100%)---[abort]

jncie@vMX3>
```

```
jncie@vMX4> show bgp summary

Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
100.1.0.3            197981          2          2       0       0          25 0/0/0/0
0/0/0/0
100.1.0.13           197981          5          4       0       0        1:29 0/0/0/0
0/0/0/0
100.1.0.14           197981          6          4       0       0        1:28 0/0/0/0
0/0/0/0
100.1.0.15           197981          6          4       0       0        1:27 0/0/0/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To verify that vMX3 and vMX4 are functioning as a route-reflector:

```
jncie@vMX3> show bgp neighbor 100.1.0.13

Peer: 100.1.0.13+64380 AS 197981 Local: 100.1.0.3+179 AS 197981
  Group: jncie                Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
...
```

To see if we authenticated all BGP sessions, we can verify the options on both spines:

```
jncie@vMX3> show bgp neighbor | match Options

  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>

jncie@vMX4> show bgp neighbor | match Options

  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
```

Task 2.14: IPv6 IBGP Configuration

- Configure an IBGP session between the vMX routers and the vQFX switches.
- vQFX switches are not allowed to peer with each other.
- Authenticate all IBGP sessions using password ‘jncie’.

Solution

vMX3:

```
set protocols bgp group jncie-v6 type internal
set protocols bgp group jncie-v6 local-address 2001:100::3
set protocols bgp group jncie-v6 family inet6 unicast
set protocols bgp group jncie-v6 authentication-key jncie
set protocols bgp group jncie-v6 cluster 100.1.0.3
set protocols bgp group jncie-v6 neighbor 2001:100::4
set protocols bgp group jncie-v6 neighbor 2001:100::13
set protocols bgp group jncie-v6 neighbor 2001:100::14
set protocols bgp group jncie-v6 neighbor 2001:100::15
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set protocols bgp group jncie-v6 type internal
set protocols bgp group jncie-v6 local-address 2001:100::4
set protocols bgp group jncie-v6 family inet6 unicast
set protocols bgp group jncie-v6 authentication-key jncie
set protocols bgp group jncie-v6 cluster 100.1.0.4
set protocols bgp group jncie-v6 neighbor 2001:100::3
set protocols bgp group jncie-v6 neighbor 2001:100::13
set protocols bgp group jncie-v6 neighbor 2001:100::14
set protocols bgp group jncie-v6 neighbor 2001:100::15
```

vQFX3:

```
set protocols bgp group jncie-v6 type internal
set protocols bgp group jncie-v6 local-address 2001:100::13
set protocols bgp group jncie-v6 family inet6 unicast
set protocols bgp group jncie-v6 authentication-key jncie
set protocols bgp group jncie-v6 neighbor 2001:100::3
set protocols bgp group jncie-v6 neighbor 2001:100::4
```

vQFX4:

```
set protocols bgp group jncie-v6 type internal
set protocols bgp group jncie-v6 local-address 2001:100::14
set protocols bgp group jncie-v6 family inet6 unicast
set protocols bgp group jncie-v6 authentication-key jncie
set protocols bgp group jncie-v6 neighbor 2001:100::3
set protocols bgp group jncie-v6 neighbor 2001:100::4
```

vQFX5:

```
set protocols bgp group jncie-v6 type internal
set protocols bgp group jncie-v6 local-address 2001:100::15
set protocols bgp group jncie-v6 family inet6 unicast
set protocols bgp group jncie-v6 authentication-key jncie
set protocols bgp group jncie-v6 neighbor 2001:100::3
set protocols bgp group jncie-v6 neighbor 2001:100::4
```

Verification

The verification for the IPv6 BGP sessions is no different from the verification of the IPv4 BGP sessions. The most efficient way to check the configuration is by examining the BGP sessions from both spines and by checking the options that are active for the BGP sessions:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show bgp summary group jncie-v6

Groups: 2 Peers: 8 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
inet6.0
                         0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:100::4          197981          6          5       0       0        1:27 Establ
  inet6.0: 0/0/0/0
2001:100::13         197981          6          4       0       0        1:05 Establ
  inet6.0: 0/0/0/0
2001:100::14         197981          5          2       0       0          56 Establ
  inet6.0: 0/0/0/0
2001:100::15         197981          5          3       0       0          59 Establ
  inet6.0: 0/0/0/0

jncie@vMX3> show bgp neighbor | match Options

  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
```

Notice that the first four sessions are the previously configured IPv4 sessions, the last four are the newly configured IPv6 BGP sessions.

After checking the BGP sessions from the vMX3 spine, we move over to the vMX4 spine and issue the same command:

```
jncie@vMX4> show bgp summary group jncie-v6

Groups: 2 Peers: 8 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
inet6.0
                         0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:100::3          197981          5          4       0       0        1:42 Establ
  inet6.0: 0/0/0/0
2001:100::13         197981          5          4       0       0        1:16 Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  inet6.0: 0/0/0/0
2001:100::14            197981         5          3       0       0        1:07 Establ
  inet6.0: 0/0/0/0
2001:100::15            197981         5          3       0       0        1:02 Establ
  inet6.0: 0/0/0/0

jncie@vMX4> show bgp neighbor | match Options

  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Refresh>
```

Task 2.15: Server 101-110

- Advertise the prefixes for server 101-110 into BGP
- Make sure you can identify the routes to these servers anywhere inside the IP-fabric.

Solution

vQFX3:

```
set protocols bgp group jncie export servers

set policy-options policy-statement servers term 1 from protocol direct
set policy-options policy-statement servers term 1 from route-filter 10.200.100.0/20 prefix-
length-range /24-/24
set policy-options policy-statement servers term 1 then community add servers
set policy-options policy-statement servers term 1 then accept

set policy-options community servers members 1:1
```

Verification

To verify whether or not vQFX3 is advertising the routes, we issue the following command on vQFX3:

```
jncie@vQFX3> show route advertising-protocol bgp 100.1.0.3

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.200.100.0/24         Self                         100        I
* 10.200.101.0/24         Self                         100        I
* 10.200.102.0/24         Self                         100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 10.200.103.0/24          Self                    100        I
* 10.200.104.0/24          Self                    100        I
* 10.200.105.0/24          Self                    100        I
* 10.200.106.0/24          Self                    100        I
* 10.200.107.0/24          Self                    100        I
* 10.200.108.0/24          Self                    100        I
* 10.200.109.0/24          Self                    100        I
* 10.200.110.0/24          Self                    100        I
```

To see whether or not a community is attached to the route, we need to use the keyword 'detail' or 'extensive':

```
jncie@vQFX3> show route advertising-protocol bgp 100.1.0.3 10.200.100.0/24 extensive

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
* 10.200.100.0/24 (1 entry, 1 announced)
 BGP group jncie type Internal
     Nexthop: Self
     Localpref: 100
     AS path: [197981] I
     Communities: 1:1
```

To check and see if the routes were propagated throughout the IP-Fabric, we issue the following command on vQFX5:

```
jncie@vQFX5> show route community 1:1

inet.0: 32 destinations, 43 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.100.0/24    *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.101.0/24    *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.102.0/24    *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.103.0/24    *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.104.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.105.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.106.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.107.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.108.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.109.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
10.200.110.0/24     *[BGP/170] 00:05:23, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                    [BGP/170] 00:05:23, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
```

You should make sure that you know all verification methods possible so that you know how to verify without actually generating actual traffic. However, since in this lab we can, we might as well verify IP-connectivity by sending some PINGs to one of the hosts:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX5> ping 10.200.100.1 rapid

PING 10.200.100.1 (10.200.100.1): 56 data bytes
!!!!!
--- 10.200.100.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 33.949/35.308/36.327/1.070 ms
```

Task 2.16: EBGP Configuration

- Configure EBGP sessions with the autonomous systems displayed in the topology.
- Make sure that the received routes are correctly distributed throughout the IP-Fabric.

AS	Neighboring IPv4	Neighboring IPv6
18153	55.0.0.1	
33915	33.0.0.1	
65000	100.2.0.2	2001:100:2::1
65000	100.2.0.6	2001:100:2::3
13013	13.0.0.2	
45915	8.0.0.1	
26	26.0.0.2	2001:26::1
45915	8.0.0.5	
26	26.0.0.6	2001:26::3

Solution

vQFX3:

```
set protocols bgp group transit-as18153 neighbor 55.0.0.1 peer-as 18153
set protocols bgp group as65000-v6 neighbor 2001:100:2::1 peer-as 65000
set protocols bgp group as33915 neighbor 33.0.0.1 peer-as 33915
set protocols bgp group as65000 neighbor 100.2.0.2 peer-as 65000

set protocols bgp group jncie export nhs
set protocols bgp group jncie-v6 export nhs
set policy-options policy-statement nhs term nhs then next-hop self
```

vQFX4:

```
set protocols bgp group transit-as26 neighbor 26.0.0.2 peer-as 26
set protocols bgp group transit-as26-v6 neighbor 2001:26::1 peer-as 26
set protocols bgp group as65000-v6 neighbor 2001:100:2::3 peer-as 65000
set protocols bgp group as45915 neighbor 8.0.0.1 peer-as 45915
set protocols bgp group as13013 neighbor 13.0.0.2 peer-as 13013
set protocols bgp group as65000 neighbor 100.2.0.6 peer-as 65000
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group jncie export nhs
set protocols bgp group jncie-v6 export nhs
set policy-options policy-statement nhs term nhs then next-hop self
```

vQFX5:

```
set protocols bgp group transit-as26 neighbor 26.0.0.6 peer-as 26
set protocols bgp group transit-as26-v6 neighbor 2001:26::3 peer-as 26
set protocols bgp group as45915 neighbor 8.0.0.5 peer-as 45915

set protocols bgp group jncie export nhs
set protocols bgp group jncie-v6 export nhs
set policy-options policy-statement nhs term nhs then next-hop self
```

Verification

First, we verify the state of the EBGP sessions:

```
jncie@vQFX3> show bgp summary

Groups: 6 Peers: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          6          0          0          0          0
inet6.0
                       5          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
33.0.0.1              33915          5         12       0       2        1:47 0/0/0/0
0/0/0/0
55.0.0.1              18153          7         12       0       2        1:45 1/1/1/0
0/0/0/0
100.1.0.3            197981       3057       3100       0       0    23:20:50 4/4/4/0
0/0/0/0
100.1.0.4            197981       3058       3098       0       0    23:20:46 0/4/4/0
0/0/0/0
100.2.0.2             65000          6         10       0       2        1:47 1/1/1/0
0/0/0/0
2001:100::3          197981       3038       3077       0       0    23:13:58 Establ
  inet6.0: 2/2/2/0
2001:100::4          197981       3038       3075       0       0    23:13:54 Establ
  inet6.0: 0/2/2/0
2001:100:2::1         65000          6          7       0       2        1:25 Establ
  inet6.0: 1/1/1/0

{master:0}
jncie@vQFX3>
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4> show bgp summary

Groups: 8 Peers: 10 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      31         17          0          0          0          0
inet6.0
                       4          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
8.0.0.1                45915          5         12       0       2        1:52 1/1/1/0
0/0/0/0
13.0.0.2               13013          6         12       0       2        1:53 2/2/2/0
0/0/0/0
26.0.0.2                  26          5         10       0       2        1:50 1/1/1/0
0/0/0/0
100.1.0.3              197981       3060       3104       0       0    23:20:54 12/13/13/0
0/0/0/0
100.1.0.4              197981       3059       3103       0       0    23:20:51 0/13/13/0
0/0/0/0
100.2.0.6               65000          7         10       0       2        1:51 1/1/1/0
0/0/0/0
2001:26::1                26          6          8       0       2        1:43 Establ
  inet6.0: 1/1/1/0
2001:100::3            197981       3037       3079       0       0    23:13:54 Establ
  inet6.0: 1/1/1/0
2001:100::4            197981       3037       3078       0       0    23:13:50 Establ
  inet6.0: 0/1/1/0
2001:100:2::3           65000          6          7       0       2        1:34 Establ
  inet6.0: 1/1/1/0

{master:0}
jncie@vQFX4>

jncie@vQFX5> show bgp summary

Groups: 5 Peers: 7 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      36         17          0          0          0          0
inet6.0
                       7          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
8.0.0.5                45915          8         13       0       2        2:40 1/1/1/0
0/0/0/0
26.0.0.6                  26          8         13       0       2        2:39 1/1/1/0
0/0/0/0
100.1.0.3              197981       3062       3095       0       0    23:21:40 15/17/17/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
100.1.0.4              197981       3061       3094       0          0       23:21:36 0/17/17/0
0/0/0/0
2001:26::3                 26          7         10       0          2            2:21 Establ
  inet6.0: 1/1/1/0
2001:100::3            197981       3042       3075       0          0       23:14:44 Establ
  inet6.0: 2/3/3/0
2001:100::4            197981       3040       3074       0          0       23:14:32 Establ
  inet6.0: 0/3/3/0

{master:0}
jncie@vQFX5>
```

Then we verify whether or not the routes have been correctly distributed throughout the IP-Fabric.

If we only configured the EBGP sessions, the route-advertisement from the vQFX3 point of view would have looked like this:

```
jncie@vQFX3> show route advertising-protocol bgp 100.1.0.3

inet.0: 47 destinations, 47 routes (47 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               55.0.0.1                     100        18153 I
* 10.200.100.0/24         Self                         100        I
* 10.200.101.0/24         Self                         100        I
* 10.200.102.0/24         Self                         100        I
* 10.200.103.0/24         Self                         100        I
* 10.200.104.0/24         Self                         100        I
* 10.200.105.0/24         Self                         100        I
* 10.200.106.0/24         Self                         100        I
* 10.200.107.0/24         Self                         100        I
* 10.200.108.0/24         Self                         100        I
* 10.200.109.0/24         Self                         100        I
* 10.200.110.0/24         Self                         100        I
* 81.0.0.0/23             100.2.0.2                    100        65000 I

jncie@vQFX3> show route advertising-protocol bgp 2001:100::3

inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:81::/32            2001:100:2::1                100        65000 I
```

If we were to examine the routing table in that state, the only BGP routes we would see would have been the locally learned EBGP routes:

```
jncie@vQFX3> show route protocol bgp

inet.0: 47 destinations, 47 routes (47 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
0.0.0.0/0          *[BGP/170] 00:05:59, localpref 100
                      AS path: 18153 I, validation-state: unverified
                    > to 55.0.0.1 via xe-0/0/1.3560
81.0.0.0/23        *[BGP/170] 00:05:55, localpref 100
                      AS path: 65000 I, validation-state: unverified
                    > to 100.2.0.2 via xe-0/0/3.115

inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:81::/32       *[BGP/170] 00:05:42, localpref 100
                      AS path: 65000 I, validation-state: unverified
                    > to 2001:100:2::1 via xe-0/0/3.115
```

We did not inject the links towards the other autonomous systems into the routing table. Because of this, the route-reflectors cannot resolve these routes:

```
jncie@vMX3> show route resolution unresolved
Tree Index 1
0.0.0.0/0
        Protocol Nexthop: 26.0.0.2
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0

81.0.0.0/23
        Protocol Nexthop: 100.2.0.2
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0

8.0.0.0/8
        Protocol Nexthop: 8.0.0.1
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0
...
Tree Index 2
2001:81::/32
        Protocol Nexthop: 2001:100:2::1
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0

2001:12::/32
        Protocol Nexthop: 2001:100:2::3
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0

::/0
        Protocol Nexthop: 2001:26::3
        Indirect nexthop: 0x0 - INH Session ID: 0x0      Path forwarding nexthops link: 0x0
         Path inh link: 0x0

...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

On vQFX5, we would have seen something similar. The only routes in the BGP routing-table would have been the locally learned EBGP routes and direct routes that vQFX3 advertised:

```
jncie@vQFX5> show route protocol bgp

inet.0: 34 destinations, 45 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:06:11, localpref 100
                      AS path: 26 I, validation-state: unverified
                    > to 26.0.0.6 via xe-0/0/1.4000
8.0.0.0/8          *[BGP/170] 00:06:09, localpref 100
                      AS path: 45915 I, validation-state: unverified
                    > to 8.0.0.5 via xe-0/0/1.116
10.200.100.0/24    *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.101.0/24    *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.102.0/24    *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.103.0/24    *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.104.0/24    *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.105.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.106.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.107.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.108.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.109.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.110.0/24     *[BGP/170] 23:04:27, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    [BGP/170] 23:04:27, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0

inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                *[BGP/170] 00:06:03, localpref 100
                      AS path: 26 I, validation-state: unverified
                    > to 2001:26::3 via xe-0/0/1.4000
```

We can solve this by injecting the links towards the EBGP peers into the OSPF table, or by having the devices rewrite the next-hop.

We chose to solve this by having the leaf devices rewrite the next-hop. After applying the 'next-hop-self' policies, we can see the following:

```
jncie@vQFX3> show route advertising-protocol bgp 100.1.0.3

inet.0: 51 destinations, 55 routes (51 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                         100        18153 I
* 10.200.100.0/24         Self                         100        I
* 10.200.101.0/24         Self                         100        I
* 10.200.102.0/24         Self                         100        I
* 10.200.103.0/24         Self                         100        I
* 10.200.104.0/24         Self                         100        I
* 10.200.105.0/24         Self                         100        I
* 10.200.106.0/24         Self                         100        I
* 10.200.107.0/24         Self                         100        I
* 10.200.108.0/24         Self                         100        I
* 10.200.109.0/24         Self                         100        I
* 10.200.110.0/24         Self                         100        I
* 81.0.0.0/23             Self                         100        65000 I

jncie@vQFX3> show route advertising-protocol bgp 2001:100::3

inet6.0: 15 destinations, 17 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:81::/32            Self                         100        65000 I
```

After setting the next-hop to self, the route reflector and the other nodes will have no problem resolving the next-hop that is associated with the BGP route. To verify this, we move to the route reflector again:

```
jncie@vMX3> show route resolution unresolved
Tree Index 1
Tree Index 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3>
```

On vQFX5, after rewriting the next-hop on all of the leaf switches, we now see EBGP routes that were reflected by the route reflectors (100.1.0.3 and 100.1.0.4:

```
jncie@vQFX5> show route protocol bgp

inet.0: 38 destinations, 57 routes (38 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:01:15, localpref 100
                      AS path: 26 I, validation-state: unverified
                    > to 26.0.0.6 via xe-0/0/1.4000
                    [BGP/170] 00:01:13, localpref 100, from 100.1.0.3
                      AS path: 18153 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 00:01:13, localpref 100, from 100.1.0.4
                      AS path: 18153 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
8.0.0.0/8          *[BGP/170] 00:01:16, localpref 100
                      AS path: 45915 I, validation-state: unverified
                    > to 8.0.0.5 via xe-0/0/1.116
                    [BGP/170] 00:01:15, localpref 100, from 100.1.0.3
                      AS path: 45915 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 00:01:15, localpref 100, from 100.1.0.4
                      AS path: 45915 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.100.0/24    *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.101.0/24    *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.102.0/24    *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.103.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                     > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                     > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.104.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
10.200.105.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
10.200.106.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
10.200.107.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
10.200.108.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                     > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.109.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
10.200.110.0/24     *[BGP/170] 23:07:43, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 23:07:43, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
12.0.0.0/17         *[BGP/170] 00:01:13, localpref 100, from 100.1.0.3
                      AS path: 65000 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 00:01:13, localpref 100, from 100.1.0.4
                      AS path: 65000 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
13.0.0.0/26         *[BGP/170] 00:01:15, localpref 100, from 100.1.0.3
                      AS path: 13013 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 00:01:15, localpref 100, from 100.1.0.4
                      AS path: 13013 I, validation-state: unverified
                      to 100.0.0.6 via xe-0/0/2.0
                    > to 100.0.0.4 via xe-0/0/4.0
21.0.0.0/21         *[BGP/170] 00:01:15, localpref 100, from 100.1.0.3
                      AS path: 13013 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 00:01:15, localpref 100, from 100.1.0.4
                      AS path: 13013 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
81.0.0.0/23         *[BGP/170] 00:01:15, localpref 100, from 100.1.0.3
                      AS path: 65000 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                     [BGP/170] 00:01:15, localpref 100, from 100.1.0.4
                      AS path: 65000 I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
inet6.0: 15 destinations, 19 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0                    *[BGP/170] 00:00:57, localpref 100
                          AS path: 26 I, validation-state: unverified
                        > to 2001:26::3 via xe-0/0/1.4000
                        [BGP/170] 00:01:06, localpref 100, from 2001:100::3
                          AS path: 26 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
                        [BGP/170] 00:01:06, localpref 100, from 2001:100::4
                          AS path: 26 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
2001:12::/32            *[BGP/170] 00:00:57, localpref 100, from 2001:100::3
                          AS path: 65000 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
                        [BGP/170] 00:00:57, localpref 100, from 2001:100::4
                          AS path: 65000 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
2001:81::/32            *[BGP/170] 00:00:53, localpref 100, from 2001:100::3
                          AS path: 65000 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
                        [BGP/170] 00:00:53, localpref 100, from 2001:100::4
                          AS path: 65000 I, validation-state: unverified
                        > to fe80::20c:29ff:fe7d:200c via xe-0/0/2.0
                          to fe80::20c:29ff:fecf:e186 via xe-0/0/4.0
```

Task 2.17: Local AS Advertisement

- Advertise the IP ranges for your AS to all autonomous systems.
  - For IPV4, advertise 100.0.0/8
  - For IPv6, advertise 2001:100::/32
- Make sure that leaf switches will not advertise any of these ranges when the connections to the spines are lost.

Solution

To make sure that the leaf switches will only advertise the routes when they have a connection to either of the spines, we have the spines advertise the IP range to the leaf devices.

On both vMX3 and vMX4, we configure an aggregate route for IPv4 as well as IPv6. After this, we write a routing policy to advertise the aggregates to the leaf switches.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set protocols bgp group jncie export aggregate
set protocols bgp group jncie-v6 export aggregate-v6

set routing-options rib inet6.0 aggregate route 2001:100::/32
set routing-options aggregate route 100.0.0.0/8

set policy-options policy-statement aggregate term aggregate from protocol aggregate
set policy-options policy-statement aggregate term aggregate from route-filter 100.0.0.0/8
exact
set policy-options policy-statement aggregate term aggregate then accept

set policy-options policy-statement aggregate-v6 term aggregate from protocol aggregate
set policy-options policy-statement aggregate-v6 term aggregate from route-filter
2001:100::/32 exact
set policy-options policy-statement aggregate-v6 term aggregate then accept
```

vMX4:

```
set protocols bgp group jncie export aggregate
set protocols bgp group jncie-v6 export aggregate-v6

set routing-options rib inet6.0 aggregate route 2001:100::/32
set routing-options aggregate route 100.0.0.0/8

set policy-options policy-statement aggregate term aggregate from protocol aggregate
set policy-options policy-statement aggregate term aggregate from route-filter 100.0.0.0/8
exact
set policy-options policy-statement aggregate term aggregate then accept

set policy-options policy-statement aggregate-v6 term aggregate from protocol aggregate
set policy-options policy-statement aggregate-v6 term aggregate from route-filter
2001:100::/32 exact
set policy-options policy-statement aggregate-v6 term aggregate then accept
```

Verification

First, we verify the state of the aggregate route on vMX3 as well as vMX4:

```
jncie@vMX3> show route protocol aggregate

inet.0: 34 destinations, 55 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.0.0.0/8         *[Aggregate/130] 00:00:14
                       Reject

inet6.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
2001:100::/32      *[Aggregate/130] 00:00:14
                      Reject

jncie@vMX4> show route protocol aggregate

inet.0: 34 destinations, 55 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.0.0.0/8        *[Aggregate/130] 00:00:26
                      Reject

inet6.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:100::/32      *[Aggregate/130] 00:00:26
                      Reject
```

The aggregates will always be in this up state. This is because among the contributing routes, we find the loopback interface:

```
jncie@vMX3> show route protocol aggregate detail

inet.0: 34 destinations, 55 routes (34 active, 0 holddown, 0 hidden)
100.0.0.0/8 (2 entries, 1 announced)
        *Aggregate Preference: 130
                Next hop type: Reject, Next hop index: 0
                Address: 0x9f48f84
                Next-hop reference count: 4
                State: <Active Int Ext>
                Local AS: 197981
                Age: 2:18
                Validation State: unverified
                Task: Aggregate
                Announcement bits (3): 0-KRT 4-BGP_RT_Background 5-Resolve tree 1
                AS path: I (LocalAgg)
                Flags:                  Depth: 0        Active
                AS path list:
                AS path: I Refcount: 11
                Contributing Routes (11):
                        100.0.0.0/31 proto Direct
                        100.0.0.2/31 proto Direct
                        100.0.0.4/31 proto Direct
                        100.1.0.3/32 proto Direct
                        100.0.0.6/31 proto OSPF3
                        100.0.0.8/31 proto OSPF3
                        100.0.0.10/31 proto OSPF3
                        100.1.0.4/32 proto OSPF3
                        100.1.0.13/32 proto OSPF3
                        100.1.0.14/32 proto OSPF3
                        100.1.0.15/32 proto OSPF3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
inet6.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)

2001:100::/32 (2 entries, 1 announced)
        *Aggregate Preference: 130
                Next hop type: Reject, Next hop index: 0
                Address: 0x9f48f84
                Next-hop reference count: 4
                State: <Active Int Ext>
                Local AS: 197981
                Age: 2:18
                Validation State: unverified
                Task: Aggregate
                Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 2
                AS path: I (LocalAgg)
                Flags:                  Depth: 0        Active
                AS path list:
                AS path: I Refcount: 5
                Contributing Routes (5):
                        2001:100::3/128 proto Direct
                        2001:100::4/128 proto OSPF3
                        2001:100::13/128 proto OSPF3
                        2001:100::14/128 proto OSPF3
                        2001:100::15/128 proto OSPF3
```

Another thing we notice from the preceding output is that fact that the vMX is advertising the route using BGP. Notice how the output would look if we had not applied an export policy on the vMX

```
jncie@vMX3> show route protocol aggregate detail | match Announ
100.0.0.0/8 (1 entry, 1 announced)
                Announcement bits (2): 0-KRT 5-Resolve tree 1
2001:100::/32 (1 entry, 1 announced)
                Announcement bits (2): 0-KRT 3-Resolve tree 2
```

To verify that the route is send to the leaf, we can also check the advertised routes like this:

```
jncie@vMX3> show route advertising-protocol bgp 100.1.0.13 detail

inet.0: 34 destinations, 55 routes (34 active, 0 holddown, 0 hidden)
* 8.0.0.0/8 (3 entries, 1 announced)
 BGP group jncie type Internal
     Nexthop: 100.1.0.14
     Localpref: 100
     AS path: [197981] 45915 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 12.0.0.0/17 (2 entries, 1 announced)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 BGP group jncie type Internal
     Nexthop: 100.1.0.14
     Localpref: 100
     AS path: [197981] 65000 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 13.0.0.0/26 (2 entries, 1 announced)
 BGP group jncie type Internal
     Nexthop: 100.1.0.14
     Localpref: 100
     AS path: [197981] 13013 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 21.0.0.0/21 (2 entries, 1 announced)
 BGP group jncie type Internal
     Nexthop: 100.1.0.14
     Localpref: 100
     AS path: [197981] 13013 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 100.0.0.0/8 (2 entries, 1 announced)
 BGP group jncie type Internal
     Nexthop: Self
     Localpref: 100
     AS path: [197981] I (LocalAgg)

jncie@vMX3> show route advertising-protocol bgp 2001:100::13 detail

inet6.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
* ::/0 (3 entries, 1 announced)
 BGP group jncie-v6 type Internal
     Nexthop: 2001:100::14
     Localpref: 100
     AS path: [197981] 26 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 2001:12::/32 (2 entries, 1 announced)
 BGP group jncie-v6 type Internal
     Nexthop: 2001:100::14
     Localpref: 100
     AS path: [197981] 65000 I
     Cluster ID: 100.1.0.3
     Originator ID: 100.1.0.14

* 2001:100::/32 (2 entries, 1 announced)
 BGP group jncie-v6 type Internal
     Nexthop: Self
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Localpref: 100
    AS path: [197981] I (LocalAgg)
```

After verifying that the route-reflectors advertise the aggregate for IPv4 as well as IPv6, we move to the leaf switches. Since vQFX4 has IPv4 as well as IPv6 peers, we check the advertisement from that leaf:

```
jncie@vQFX4> show route advertising-protocol bgp 26.0.0.2

inet.0: 43 destinations, 58 routes (43 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               Self                                    45915 I
* 10.200.100.0/24         Self                                    I
* 10.200.101.0/24         Self                                    I
* 10.200.102.0/24         Self                                    I
* 10.200.103.0/24         Self                                    I
* 10.200.104.0/24         Self                                    I
* 10.200.105.0/24         Self                                    I
* 10.200.106.0/24         Self                                    I
* 10.200.107.0/24         Self                                    I
* 10.200.108.0/24         Self                                    I
* 10.200.109.0/24         Self                                    I
* 10.200.110.0/24         Self                                    I
* 12.0.0.0/17             Self                                    65000 I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 81.0.0.0/23             Self                                    65000 I
* 100.0.0.0/8             Self                                    I

{master:0}
jncie@vQFX4> show route advertising-protocol bgp 2001:26::1

inet6.0: 19 destinations, 21 routes (19 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:12::/32            Self                                    65000 I
* 2001:81::/32            Self                                    65000 I
* 2001:100::/32           Self                                    I
```

Task 2.18: BGP RIB Configuration

- Make sure that redundant address information from multihomed autonomous systems is available in the RIB as well as the FIB on every device.

Solution

On every device:

```
set protocols bgp group jncie multipath
set protocols bgp group jncie family inet unicast add-path receive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group jncie family inet unicast add-path send path-count 3

set protocols bgp group jncie-v6 multipath
set protocols bgp group jncie-v6 family inet6 unicast add-path receive
set protocols bgp group jncie-v6 family inet6 unicast add-path send path-count 3
```

Verification

We start out verifying the add-path configuration:

```
jncie@vMX3> show bgp neighbor 100.1.0.13

Peer: 100.1.0.13+179 AS 197981 Local: 100.1.0.3+55642 AS 197981
  Group: jncie                 Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ aggregate ]
  Options: <Preference LocalAddress AuthKey Cluster AddressFamily Multipath Refresh>
  Authentication key is configured
  Address families configured: inet-unicast
  Local Address: 100.1.0.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 100.1.0.13      Local ID: 100.1.0.3         Active Holdtime: 90
  Keepalive Interval: 30         Group index: 2    Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 197981)
  NLRI's for which peer can send and receive multiple paths: inet-unicast
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              13
    Received prefixes:            13
    Accepted prefixes:            13
    Suppressed due to damping:    0
    Advertised prefixes:          21
  Last traffic (seconds): Received 585002 Sent 464 Checked 585002
  Input messages:  Total 23     Updates 4       Refreshes 0     Octets 636
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Output messages: Total 42     Updates 24      Refreshes 0     Octets 2111
  Output Queue[1]: 0            (inet.0, inet-unicast)


jncie@vMX3> show bgp neighbor 2001:100::13

Peer: 2001:100::13+179 AS 197981 Local: 2001:100::3+64448 AS 197981
  Group: jncie-v6           Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ aggregate-v6 ]
Options: <Preference LocalAddress AuthKey Cluster AddressFamily Multipath Refresh>
Authentication key is configured
Address families configured: inet6-unicast
Local Address: 2001:100::3 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 100.1.0.13      Local ID: 100.1.0.3         Active Holdtime: 90
Keepalive Interval: 30         Group index: 3    Peer index: 1
I/O Session Thread: bgpio-0 State: Enabled
BFD: disabled, down
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 197981)
NLRI's for which peer can send and receive multiple paths: inet6-unicast
Table inet6.0 Bit: 30000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              1
  Received prefixes:            1
  Accepted prefixes:            1
  Suppressed due to damping:    0
  Advertised prefixes:          5
Last traffic (seconds): Received 585006 Sent 461 Checked 585006
Input messages:  Total 21     Updates 2       Refreshes 0     Octets 468
Output messages: Total 26     Updates 8       Refreshes 0     Octets 1068
Output Queue[2]: 0            (inet6.0, inet6-unicast)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Before our configuration changes:

```
jncie@vMX3> show route forwarding-table destination 8.0.0.0/8

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
8.0.0.0/8          user     0                    indr  1048578     5
                              100.0.0.3           ucst      601     7 xe-0/0/3.0


jncie@vMX3> show route advertising-protocol bgp 100.1.0.13

inet.0: 34 destinations, 55 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               100.1.0.14                   100        45915 I
* 12.0.0.0/17             100.1.0.14                   100        65000 I
* 13.0.0.0/26             100.1.0.14                   100        13013 I
* 21.0.0.0/21             100.1.0.14                   100        13013 I
* 100.0.0.0/8             Self                         100        I
```

After our configuration changes:

```
jncie@vMX3> show route forwarding-table destination 8.0.0.0/8

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
8.0.0.0/8          user     0                    ulst  1048579     2
                                                 indr  1048578     2
                              100.0.0.5           ucst      575     7 xe-0/0/4.0
                                                 indr  1048576     5
                              100.0.0.3           ucst      601     7 xe-0/0/3.0


jncie@vMX3> show route advertising-protocol bgp 100.1.0.13

inet.0: 34 destinations, 58 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               100.1.0.13                   100        18153 I
                          100.1.0.14                   100        26 I
                          100.1.0.15                   100        26 I
* 8.0.0.0/8               100.1.0.14                   100        45915 I
                          100.1.0.15                   100        45915 I
* 10.200.100.0/24         100.1.0.13                   100        I
* 10.200.101.0/24         100.1.0.13                   100        I
* 10.200.102.0/24         100.1.0.13                   100        I
* 10.200.103.0/24         100.1.0.13                   100        I
* 10.200.104.0/24         100.1.0.13                   100        I
* 10.200.105.0/24         100.1.0.13                   100        I
* 10.200.106.0/24         100.1.0.13                   100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 10.200.107.0/24          100.1.0.13                100          I
* 10.200.108.0/24          100.1.0.13                100          I
* 10.200.109.0/24          100.1.0.13                100          I
* 10.200.110.0/24          100.1.0.13                100          I
* 12.0.0.0/17              100.1.0.14                100          65000 I
* 13.0.0.0/26              100.1.0.14                100          13013 I
* 21.0.0.0/21              100.1.0.14                100          13013 I
* 81.0.0.0/23              100.1.0.13                100          65000 I
* 100.0.0.0/8              Self                      100          I
```

Task 2.19: AS65000 Customers

- Make sure that the two different customers that are using AS65000 receive all routing information.
- When prefixes that originate from AS65000 are advertised to other public ASes, AS65000 should not appear in the AS-path.

Solution

vQFX3:

```
set protocols bgp group as65000-v6 neighbor 2001:100:2::1 as-override
set protocols bgp group as65000 neighbor 100.2.0.2 as-override

set protocols bgp group transit-as18153 remove-private
set protocols bgp group as33915 remove-private
```

vQFX4:

```
set protocols bgp group as65000-v6 neighbor 2001:100:2::3 as-override
set protocols bgp group as65000 neighbor 100.2.0.6 as-override

set protocols bgp group transit-as26 remove-private
set protocols bgp group transit-as26-v6 remove-private
set protocols bgp group as45915 remove-private
set protocols bgp group as13013 remove-private
```

vQFX5:

```
set protocols bgp group transit-as26 remove-private
set protocols bgp group transit-as26-v6 remove-private
set protocols bgp group as45915 remove-private
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

On vQFX4, we receive the following routes from AS65000:

```
jncie@vQFX4> show route receive-protocol bgp 100.2.0.6

inet.0: 43 destinations, 62 routes (43 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 12.0.0.0/17             100.2.0.6                               65000 I

inet6.0: 19 destinations, 23 routes (19 active, 0 holddown, 0 hidden)
```

```
jncie@vQFX4> show route receive-protocol bgp 2001:100:2::3

inet.0: 43 destinations, 62 routes (43 active, 0 holddown, 0 hidden)

inet6.0: 19 destinations, 23 routes (19 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:12::/32            2001:100:2::3                           65000 I
```

On vQFX3, we do not advertise these routes out to the BGP peer in AS65000 due to the loop prevention mechanism in BGP:

```
jncie@vQFX3> show route advertising-protocol bgp 100.2.0.2

inet.0: 52 destinations, 63 routes (52 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    18153 I
* 8.0.0.0/8               Self                                    45915 I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 100.0.0.0/8             Self                                    I
```

```
jncie@vQFX3> show route advertising-protocol bgp 2001:100:2::1

inet6.0: 16 destinations, 21 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    26 I
* 2001:100::/32           Self                                    I
```

Once we configure AS-override, we can see the effect this has had by checking the advertised routes on vQFX3:

```
jncie@vQFX3> show route advertising-protocol bgp 100.2.0.2

inet.0: 52 destinations, 63 routes (52 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    18153 I
* 8.0.0.0/8               Self                                    45915 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 12.0.0.0/17              Self                                    197981 I
* 13.0.0.0/26              Self                                    13013 I
* 21.0.0.0/21              Self                                    13013 I
* 100.0.0.0/8              Self                                    I
```

```
jncie@vQFX3> show route advertising-protocol bgp 2001:100:2::1

inet6.0: 16 destinations, 21 routes (16 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* ::/0                    Self                                    26 I
* 2001:12::/32            Self                                    197981 I
* 2001:100::/32           Self                                    I
```

For the other part of the assignment, we need to make sure that the routes we advertise do not contain any private autonomous systems in the AS-path. Before our configuration changes, we can see the following:

```
jncie@vQFX3> show route advertising-protocol bgp 55.0.0.1

inet.0: 52 destinations, 63 routes (52 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               Self                                    45915 I
* 12.0.0.0/17             Self                                    65000 I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 81.0.0.0/23             Self                                    65000 I
* 100.0.0.0/8             Self                                    I
```

After our configuration changes, we can observe the following:

```
jncie@vQFX3> show route advertising-protocol bgp 55.0.0.1

inet.0: 52 destinations, 63 routes (52 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               Self                                    45915 I
* 12.0.0.0/17             Self                                    I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 81.0.0.0/23             Self                                    I
* 100.0.0.0/8             Self                                    I
```

The 'remove-private' knob has removed all private autonomous systems from the AS-path.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.20: IP-Transit Policy Configuration

- Prefer the IP-transit that AS197981 is multihomed to for inbound as well as outbound traffic.
- Only advertise a default route to all non-transit networks.
- Do not alter the local-preference anywhere and only apply policies on vQFX3.
- AS33915 should only receive a default route and the route for your AS.

Solution

vQFX3:

```
set protocols bgp group transit-as18153 export as-prepend-out

set policy-options policy-statement as-prepend-out term 1 then as-path-prepend "197981 197981
197981"

set protocols bgp group transit-as18153 import as-prepend-in

set policy-options policy-statement as-prepend-in term 1 then as-path-prepend "18153 18153
18153 18153"

set protocols bgp group as33915 export default-only

set policy-options policy-statement default-only term 1 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-only term 1 then accept
set policy-options policy-statement default-only term 2 then reject

set policy-options community no-export members no-export
set policy-options policy-statement servers term 1 then community add no-export
```

Verification

Before applying the configuration to vQFX3, that leaf will use the locally learned default route:

```
jncie@vQFX3> show route 0/0 exact active-path

inet.0: 52 destinations, 63 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 01:18:51, localpref 100
                      AS path: 18153 I, validation-state: unverified
                    > to 55.0.0.1 via xe-0/0/1.3560
```

Since all the BGP attributes associated with the default route from AS 18153 as well as AS26 are identical, all of the leaf switches will prefer the route that is learned locally via EBGP.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

After we increased the AS-path length for the route received from AS 18153, we will not see this reflected in the 'show route receive-protocol' verification command:

```
jncie@vQFX3> show route receive-protocol bgp 55.0.0.1

inet.0: 52 destinations, 63 routes (52 active, 1 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  0.0.0.0/0               55.0.0.1                                18153 I
```

The vQFX will apply the policies on the BGP RIB after the routes are received. To verify the result of our policy, we can issue the following command:

```
jncie@vQFX3> show route 0.0.0.0/0 exact

inet.0: 52 destinations, 63 routes (52 active, 1 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:35:00, localpref 100, from 100.1.0.3
                      AS path: 26 I, validation-state: unverified
                    > to 100.0.0.0 via xe-0/0/2.0
                      to 100.0.0.10 via xe-0/0/4.0
                    > to 100.0.0.0 via xe-0/0/2.0
                      to 100.0.0.10 via xe-0/0/4.0
                    [BGP/170] 00:35:08, localpref 100, from 100.1.0.4
                      AS path: 26 I, validation-state: unverified
                    > to 100.0.0.0 via xe-0/0/2.0
                      to 100.0.0.10 via xe-0/0/4.0
                    [BGP/170] 00:34:57, localpref 100, from 100.1.0.3
                      AS path: 26 I, validation-state: unverified
                    > to 100.0.0.0 via xe-0/0/2.0
                      to 100.0.0.10 via xe-0/0/4.0
                    [BGP/170] 00:35:05, localpref 100, from 100.1.0.4
                      AS path: 26 I, validation-state: unverified
                    > to 100.0.0.0 via xe-0/0/2.0
                      to 100.0.0.10 via xe-0/0/4.0
                    [BGP/170] 01:14:10, localpref 100
                      AS path: 18153 18153 18153 18153 18153 I, validation-state: unverified
                    > to 55.0.0.1 via xe-0/0/1.3560
```

The export policy can be verified by examining the advertised routes:

```
jncie@vQFX3> show route advertising-protocol bgp 55.0.0.1

inet.0: 52 destinations, 63 routes (52 active, 1 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    197981 197981 197981
[197981] 26 I
* 8.0.0.0/8               Self                                    197981 197981 197981
[197981] 45915 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 12.0.0.0/17              Self                                197981 197981 197981
[197981] I
* 13.0.0.0/26              Self                                197981 197981 197981
[197981] 13013 I
* 21.0.0.0/21              Self                                197981 197981 197981
[197981] 13013 I
* 81.0.0.0/23              Self                                197981 197981 197981
[197981] I
* 100.0.0.0/8              Self                                197981 197981 197981
[197981] I
```

By default, BGP will allow the export of all BGP learned routes. After introducing a 'then reject' in the export policy towards AS 33915, we can observe the following:

```
jncie@vQFX3> show route advertising-protocol bgp 33.0.0.1

inet.0: 52 destinations, 63 routes (52 active, 1 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    26 I
```

The routes towards the servers was injected into the BGP RIB by vQFX3. By default, all of these addresses will be advertised to other autonomous systems:

```
jncie@vQFX5> show route advertising-protocol bgp 26.0.0.6

inet.0: 39 destinations, 61 routes (39 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               Self                                    45915 I
* 10.200.100.0/24         Self                                    I
* 10.200.101.0/24         Self                                    I
* 10.200.102.0/24         Self                                    I
* 10.200.103.0/24         Self                                    I
* 10.200.104.0/24         Self                                    I
* 10.200.105.0/24         Self                                    I
* 10.200.106.0/24         Self                                    I
* 10.200.107.0/24         Self                                    I
* 10.200.108.0/24         Self                                    I
* 10.200.109.0/24         Self                                    I
* 10.200.110.0/24         Self                                    I
* 12.0.0.0/17             Self                                    I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 81.0.0.0/23             Self                                    I
* 100.0.0.0/8             Self                                    I
```

One way of making sure that these routes will not be advertised to other peers would have been by configuring an export policy for every EBGP peer configured on vQFX4 and vQFX5. We could have denied the prefixes through a route-filter or by referencing the server community that we had added to the routes earlier. However, since we can only

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

configure policies on the vQFX3 leaf, we are left with no other option then the use of the no-export community. We do not need to touch vQFX4 or vQFX5 because the 'no-export' community is a well-known community.

To verify our configuration, we first check the routes advertised to the route-reflector:

```
jncie@vQFX3> show route advertising-protocol bgp 100.1.0.3 extensive

inet.0: 52 destinations, 63 routes (52 active, 1 holddown, 0 hidden)
...

* 10.200.100.0/24 (1 entry, 1 announced)
 BGP group jncie type Internal
     Nexthop: Self
     Localpref: 100
     AS path: [197981] I
     Communities: 1:1 no-export
     Addpath Path ID: 1

...
```

Observe that vQFX5 has the routes to the servers, but it does not advertise the routes to any of the peers:

```
jncie@vQFX5> show route protocol bgp community 1:1

inet.0: 39 destinations, 61 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.100.0/24    *[BGP/170] 00:12:30, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 00:12:30, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.101.0/24    *[BGP/170] 00:12:30, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 00:12:30, localpref 100, from 100.1.0.4
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
10.200.102.0/24    *[BGP/170] 00:12:30, localpref 100, from 100.1.0.3
                      AS path: I, validation-state: unverified
                    > to 100.0.0.6 via xe-0/0/2.0
                      to 100.0.0.4 via xe-0/0/4.0
                    [BGP/170] 00:12:30, localpref 100, from 100.1.0.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    AS path: I, validation-state: unverified
                  > to 100.0.0.6 via xe-0/0/2.0
                    to 100.0.0.4 via xe-0/0/4.0
...

jncie@vQFX5> show route advertising-protocol bgp 26.0.0.6

inet.0: 39 destinations, 61 routes (39 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 8.0.0.0/8               Self                                    45915 I
* 12.0.0.0/17             Self                                    I
* 13.0.0.0/26             Self                                    13013 I
* 21.0.0.0/21             Self                                    13013 I
* 81.0.0.0/23             Self                                    I
* 100.0.0.0/8             Self                                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Appendix – Chapter 3: Controllerless Overlay

Part 1: A basic EBGP-based Overlay Network

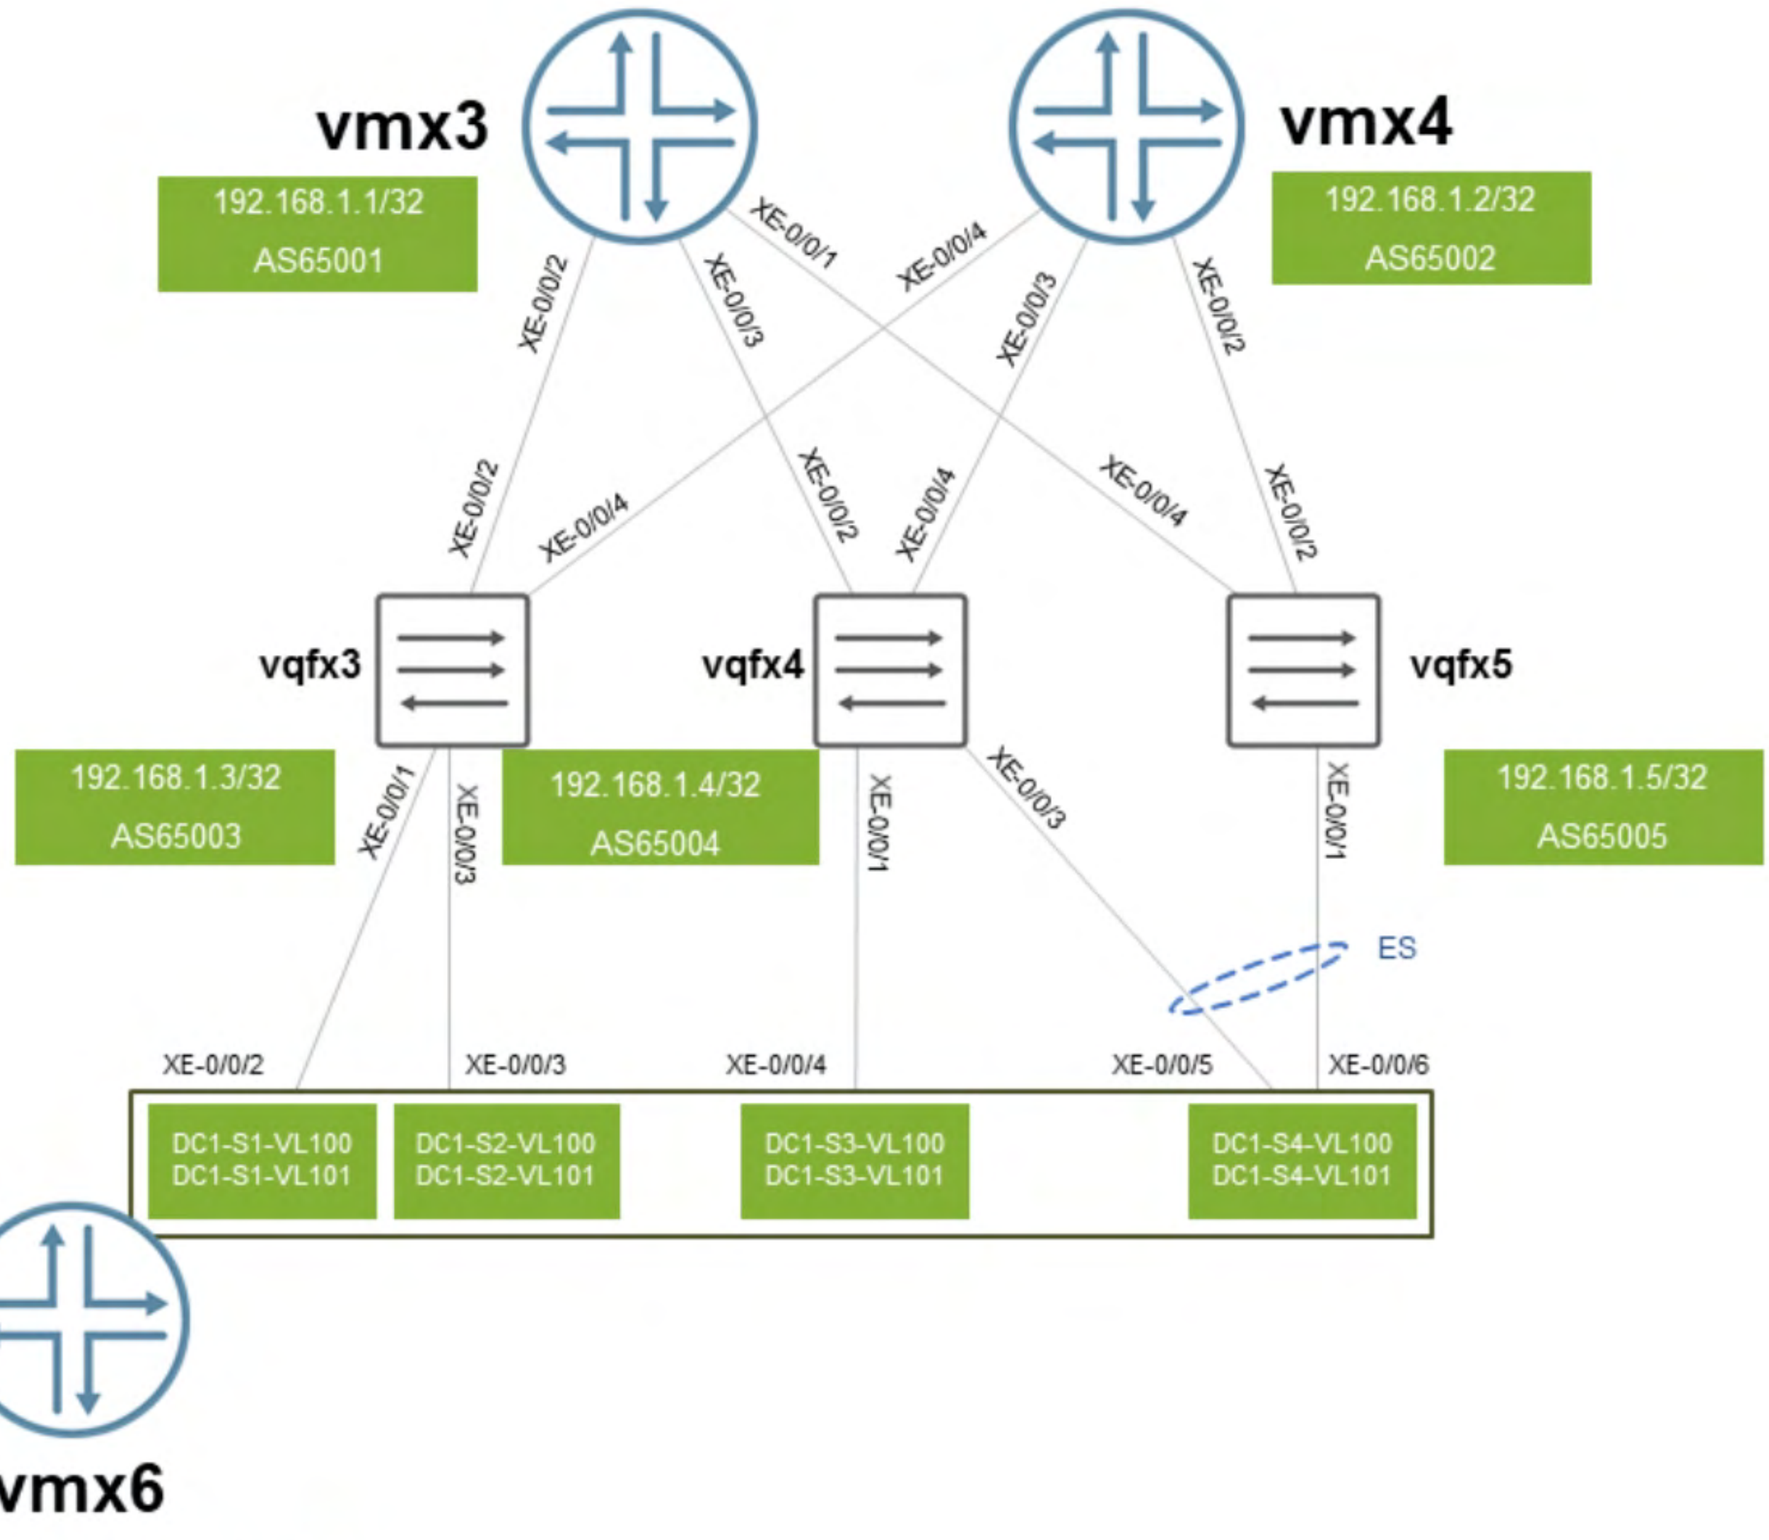

Figure 19 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

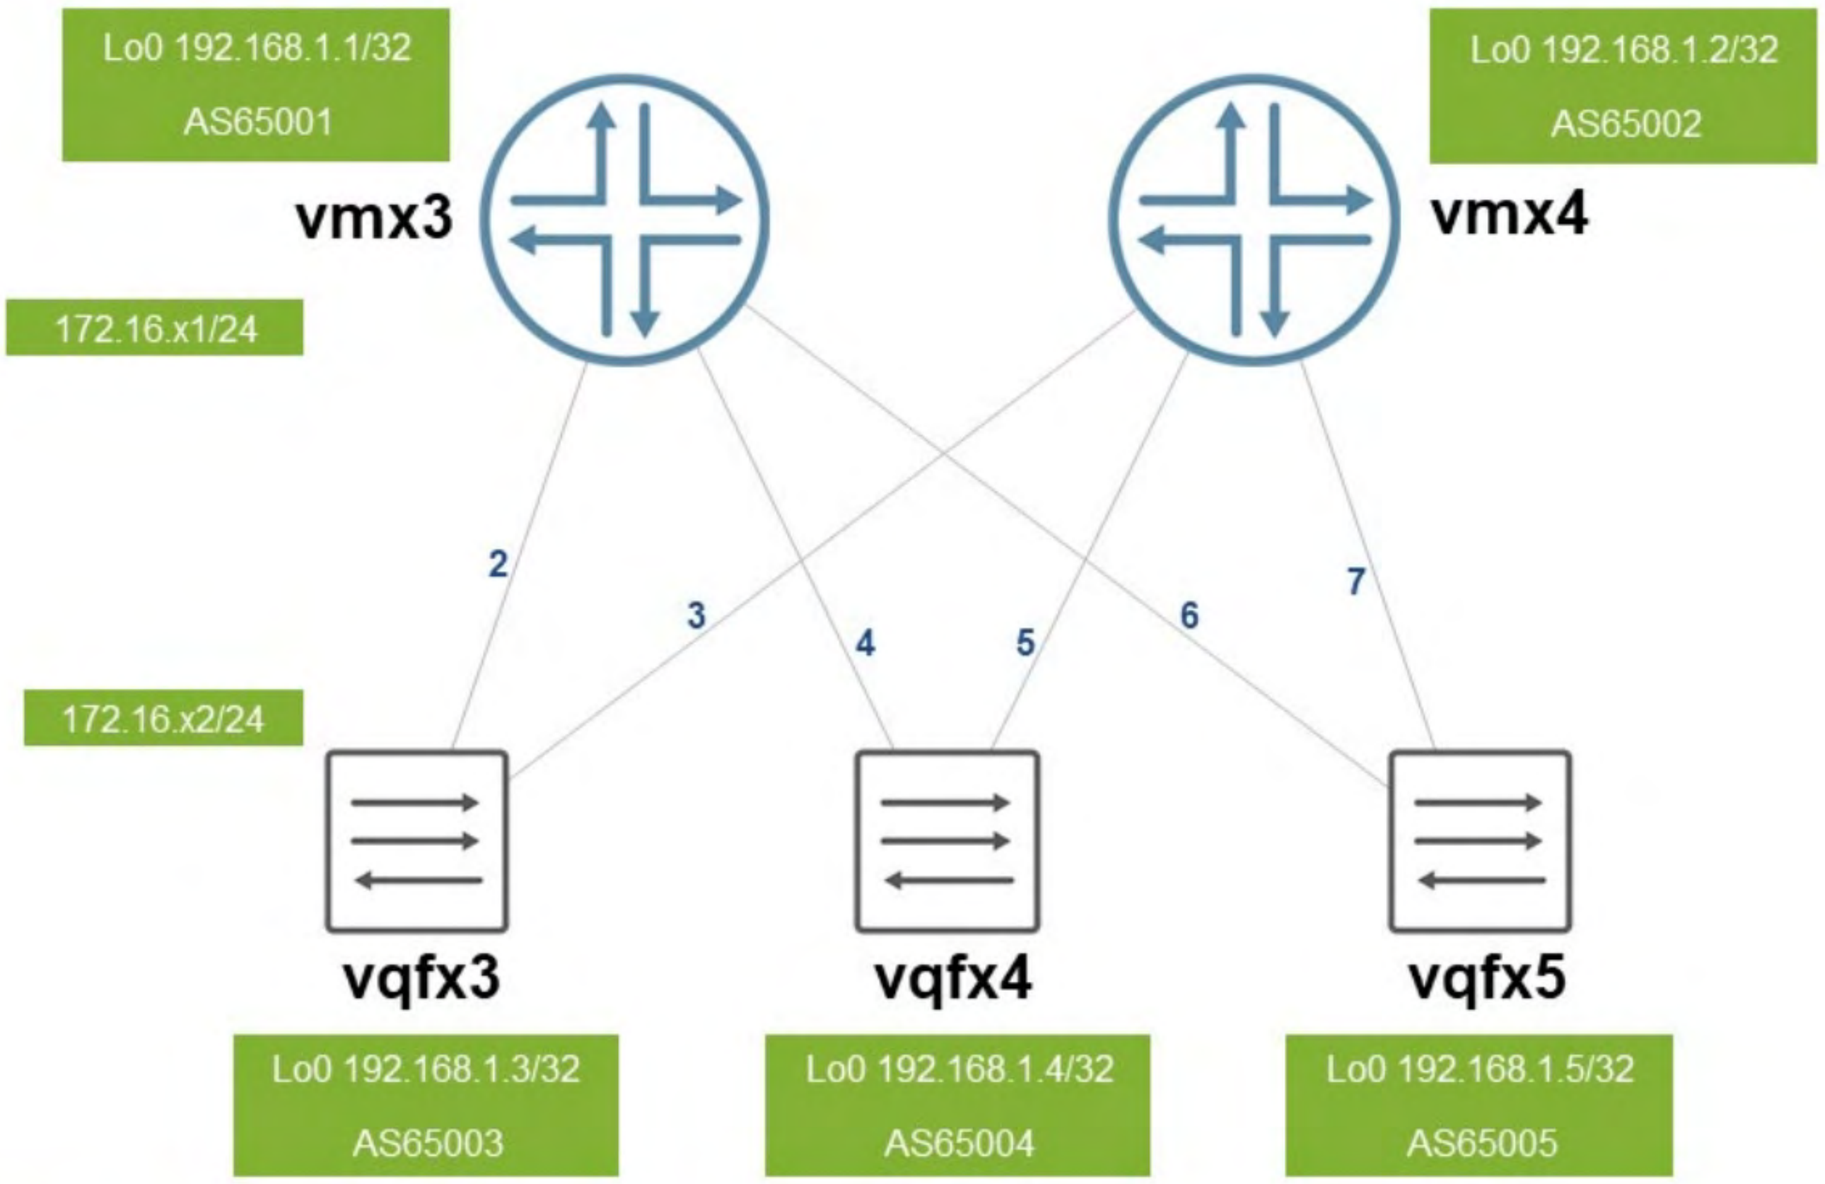

Figure 2 - Logical topology

Note: the starting topology is configured with IP-addressing only.

Task 3.1: Configure the EBGP underlay

- Configure the underlay network according to the information provided in the logical topology.
- Ensure IP-connectivity between all of the loopback IP addresses.
- Make sure all BGP sessions are 'Established' and that sessions state changes are logged to syslog.

Solution

vMX3:

We configure ‘log-updown’ directly under BGP, this way it effects all configured BGP sessions:

```
set protocols bgp log-updown
```

An export policy is applied at the group level:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay export export-loopback
```

Next is the configuration of the different BGP neighbors:

```
set protocols bgp group underlay neighbor 172.16.2.2 peer-as 65003
set protocols bgp group underlay neighbor 172.16.4.2 peer-as 65004
set protocols bgp group underlay neighbor 172.16.6.2 peer-as 65005
```

Then we configure the export policy. When we reference the lo0.0 interface in the export policy, we can re-use it on every device:

```
set policy-options policy-statement export-loopback term lo0 from interface lo0.0
set policy-options policy-statement export-loopback term lo0 then accept
```

The configuration of the autonomous-system:

```
set routing-options autonomous-system 65001
```

vMX4:

```
set protocols bgp log-updown
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay neighbor 172.16.3.2 peer-as 65003
set protocols bgp group underlay neighbor 172.16.5.2 peer-as 65004
set protocols bgp group underlay neighbor 172.16.7.2 peer-as 65005

set policy-options policy-statement export-loopback term lo0 from interface lo0.0
set policy-options policy-statement export-loopback term lo0 then accept

set routing-options autonomous-system 65002
```

vQFX3:

```
set protocols bgp log-updown
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay neighbor 172.16.2.1 peer-as 65001
set protocols bgp group underlay neighbor 172.16.3.1 peer-as 65002

set policy-options policy-statement export-loopback term lo0 from interface lo0.0
set policy-options policy-statement export-loopback term lo0 then accept

set routing-options autonomous-system 65003
```

vQFX4:

```
set protocols bgp log-updown
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay neighbor 172.16.4.1 peer-as 65001
set protocols bgp group underlay neighbor 172.16.5.1 peer-as 65002
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options policy-statement export-loopback term lo0 from interface lo0.0
set policy-options policy-statement export-loopback term lo0 then accept

set routing-options autonomous-system 65004
```

vQFX5:

```
set protocols bgp log-updown
set protocols bgp group underlay export export-loopback
set protocols bgp group underlay neighbor 172.16.7.1 peer-as 65002
set protocols bgp group underlay neighbor 172.16.6.1 peer-as 65001
set policy-options policy-statement export-loopback term lo0 from interface lo0.0
set policy-options policy-statement export-loopback term lo0 then accept

set routing-options autonomous-system 65005
```

Verification

In this assignment, we ready the underlay to support a controllerless overlay. The main objective is to get all of the BGP sessions established and to make sure that there is IP-connectivity between the loopback IPs.

After configuring all of the BGP sessions, we verify whether or not the sessions reached the state 'Established'. Since it can take some time for BGP to properly setup the session, we issue the following command: **'show bgp summary | refresh 2'**. We then wait to see if and when the neighbor relationships form.

```
jncie@vMX3> show bgp summary | refresh 2

---(refreshed at 2017-03-24 10:03:55 UTC)---
Groups: 1 Peers: 3 Down peers: 2
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         4          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2           65003          6          3       0       0           1 4/4/4/0
0/0/0/0
172.16.4.2           65004          0          0       0       0           5 Active
172.16.6.2           65005          0          0       0       0           5 Active
---(refreshed at 2017-03-24 10:03:57 UTC)---
Groups: 1 Peers: 3 Down peers: 2
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         4          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.2.2              65003          6          3       0       0           3 4/4/4/0
0/0/0/0
172.16.4.2              65004          0          0       0       0           7 Active
172.16.6.2              65005          0          0       0       0           7 Active
---(refreshed at 2017-03-24 10:03:59 UTC)---
Groups: 1 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       8          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2              65003          6          4       0       0           5 3/4/4/0
0/0/0/0
172.16.4.2              65004          6          8       0       0           1 1/4/4/0
0/0/0/0
172.16.6.2              65005          0          0       0       0           9 Active
---(refreshed at 2017-03-24 10:04:01 UTC)---
Groups: 1 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       8          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2              65003          6          4       0       0           7 3/4/4/0
0/0/0/0
172.16.4.2              65004          6          8       0       0           3 1/4/4/0
0/0/0/0
172.16.6.2              65005          0          0       0       0          11 Active
---(refreshed at 2017-03-24 10:04:03 UTC)---
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      12          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2              65003          6          5       0       0           9 2/4/4/0
0/0/0/0
172.16.4.2              65004          6          9       0       0           5 1/4/4/0
0/0/0/0
172.16.6.2              65005          8          8       0       0           1 1/4/4/0
0/0/0/0
---(*more 100%)---[abort]

jncie@vMX3>
```

The last column displays the state of the BGP session or, when the state is Established, the amount of Active/Received/Accepted and Damped routes. Finally, when all the BGP sessions are established, the output should look something like this:

```
jncie@vMX3> show bgp summary
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Groups: 1 Peers: 3 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        10          4          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.2.2            65003          8          6       0       0          41 2/3/3/0
0/0/0/0
172.16.4.2            65004          8          7       0       0          33 1/3/3/0
0/0/0/0
172.16.6.2            65005          7          8       0       0          21 1/4/4/0
0/0/0/0
```

To see what routes we have learned, we issue the following command:

```
jncie@vMX3> show route protocol bgp

inet.0: 17 destinations, 23 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.2/32     *[BGP/170] 00:01:08, localpref 100
                      AS path: 65003 65002 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:01:00, localpref 100
                      AS path: 65004 65002 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
                    [BGP/170] 00:00:53, localpref 100
                      AS path: 65005 65002 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
192.168.1.3/32     *[BGP/170] 00:01:12, localpref 100
                      AS path: 65003 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:00:53, localpref 100
                      AS path: 65005 65002 65003 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
192.168.1.4/32     *[BGP/170] 00:01:04, localpref 100
                      AS path: 65004 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
                    [BGP/170] 00:00:53, localpref 100
                      AS path: 65005 65002 65004 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
192.168.1.5/32     *[BGP/170] 00:00:53, localpref 100
                      AS path: 65005 I, validation-state: unverified
                    > to 172.16.6.2 via xe-0/0/4.0
                    [BGP/170] 00:00:57, localpref 100
                      AS path: 65003 65002 65005 I, validation-state: unverified
                    > to 172.16.2.2 via xe-0/0/2.0
                    [BGP/170] 00:00:57, localpref 100
                      AS path: 65004 65002 65005 I, validation-state: unverified
                    > to 172.16.4.2 via xe-0/0/3.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To verify IP-connectivity, we issue a ping to all of the loopback IPs configured on the different nodes. We source this ping from the loopback IP:

```
jncie@vMX3> ping source 192.168.1.1 192.168.1.2 rapid

PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 8.273/8.935/11.238/1.154 ms
jncie@vMX3> ping source 192.168.1.1 192.168.1.3 rapid
PING 192.168.1.3 (192.168.1.3): 56 data bytes
!!!!!
--- 192.168.1.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.144/8.241/9.812/1.089 ms

jncie@vMX3> ping source 192.168.1.1 192.168.1.4 rapid
PING 192.168.1.4 (192.168.1.4): 56 data bytes
!!!!!
--- 192.168.1.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.614/9.390/10.957/1.206 ms

jncie@vMX3> ping source 192.168.1.1 192.168.1.5 rapid
PING 192.168.1.5 (192.168.1.5): 56 data bytes
!!!!!
--- 192.168.1.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 68.604/116.097/191.370/51.883 ms
```

We can issue the same commands on the other vMX router in order to verify all of the BGP sessions and routing tables.

Task 3.2: EBGP Underlay Optimization

- Make sure that the devices load-balance traffic.
- Authenticate BGP sessions with the MD5-hashed password ‘jncie’.
- Make sure that the loss of a BGP neighbor is detected by BFD in exactly 10 seconds. You are not allowed to change the BGP hold-time.

Solution

On every device:

The configuration of the BGP authentication key:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay authentication-key jncie
```

Activating BFD for all BGP sessions that are part of the underlay group configuration stanza:

```
set protocols bgp group underlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group underlay bfd-liveness-detection multiplier 5
```

By default, BGP will only use only a single best path. The following configuration will have BGP consider multiple-paths from multiple AS:

```
set protocols bgp group underlay multipath multiple-as
```

Apart from having BGP consider multiple paths, these paths will need to be programmed in the forwarding table as well. To enable a Juniper device to install multiple paths into the forwarding table, a policy needs to be applied under the [routing-options] stanza :

```
set routing-options forwarding-table export lbpp
```

The policy needs to be created as well. To this end, we configure the following:

```
set policy-options policy-statement lbpp term 1 then load-balance per-packet
```

Note that even though ‘per-packet’ is configured, Juniper devices will load-balance per-flow.

Verification

BFD, authentication and load-balancing are options that we can verify by issuing the following command:

```
jncie@vQFX3> show bgp neighbor 172.16.2.1
Peer: 172.16.2.1+179 AS 65001  Local: 172.16.2.2+50041 AS 65003
  Group: underlay              Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ export-loopback ]
  Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
  Options: <MultipathAs BfdEnabled>
  Authentication key is configured
  Holdtime: 90 Preference: 170
  Number of flaps: 19
  Last flap event: BfdDown
  Peer ID: 192.168.1.1     Local ID: 192.168.1.3       Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
BFD: enabled, up
Local Interface: xe-0/0/2.0
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65001)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              3
  Received prefixes:            3
  Accepted prefixes:            3
  Suppressed due to damping:    0
  Advertised prefixes:          2
Last traffic (seconds): Received 4    Sent 26   Checked 77
Input messages:  Total 24     Updates 6       Refreshes 0     Octets 601
Output messages: Total 22     Updates 2       Refreshes 0     Octets 547
Output Queue[0]: 0            (inet.0, inet-unicast)
```

When we limit the device output, the output becomes easier to scan and we can speed up the verification process while we check all of the BGP sessions in the network:

```
jncie@vQFX3> show bgp neighbor | match "65003|option"

Peer: 172.16.2.1+179 AS 65001  Local: 172.16.2.2+50041 AS 65003
  Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
  Options: <MultipathAs BfdEnabled>
Peer: 172.16.3.1+179 AS 65002  Local: 172.16.3.2+58934 AS 65003
  Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
  Options: <MultipathAs BfdEnabled>
```

To verify and examine the BFD sessions, we issue the following command:

```
jncie@vQFX3> show bfd session extensive
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.2.1               Up        xe-0/0/2.0     10.000    2.000        5
 Client BGP, TX interval 2.000, RX interval 2.000
 Session up time 00:01:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Session type: Single hop BFD
Min async interval 2.000, min slow interval 2.000
Adaptive async TX interval 2.000, RX interval 2.000
Local min TX interval 2.000, minimum RX interval 2.000, multiplier 5
Remote min TX interval 2.000, min RX interval 2.000, multiplier 5
Local discriminator 66, remote discriminator 250
Echo mode disabled/inactive
Remote is control-plane independent
 Session ID: 0x0
```

We can see that the values we previously configured have led to a detect time of 10 seconds. Note that we have, for stability reasons, chosen a very loose BFD time. On a physical QFX, BFD can be used to detect failures a lot faster.

On to the verification of the load-balancing part. Before starting this assignment, we had not yet activated load-balancing . If we were to have examined the route towards 192.168.1.4 on vQFX3, we would have seen the following:

```
jncie@vQFX3> show route 192.168.1.4

inet.0: 13 destinations, 15 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.4/32     *[BGP/170] 00:09:36, localpref 100
                      AS path: 65002 65004 I, validation-state: unverified
                    > to 172.16.3.1 via xe-0/0/4.0
                    [BGP/170] 00:06:28, localpref 100
                      AS path: 65001 65004 I, validation-state: unverified
                    > to 172.16.2.1 via xe-0/0/2.0
```

The output shows that there are two available paths in the routing-table, but only one of them is active. As result of this, the forwarding table would have been programmed with a single unicast next-hop:

```
{master:0}
jncie@vQFX3> show route forwarding-table destination 192.168.1.4

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.4/32     user     0 172.16.3.1         ucst     1709     6 xe-0/0/4.0
```

After activating load-balancing on vQFX3 we can see the following:

```
jncie@vQFX3> show route 192.168.1.4

inet.0: 13 destinations, 15 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
192.168.1.4/32     *[BGP/170] 00:00:09, localpref 100, from 172.16.2.1
                      AS path: 65001 65004 I, validation-state: unverified
                      to 172.16.2.1 via xe-0/0/2.0
                    > to 172.16.3.1 via xe-0/0/4.0
                    [BGP/170] 00:00:05, localpref 100
                      AS path: 65002 65004 I, validation-state: unverified
                    > to 172.16.3.1 via xe-0/0/4.0
```

Due to the configuration of the 'multipath multiple-as', the routing table now displays two next-hops. We also configured, and applied, a load-balancing policy to the routing-table. Due to this policy that is applied to the forwarding table, we can now see the following entry:

```
jncie@vQFX3> show route forwarding-table destination 192.168.1.4

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.4/32     user     0                    ulst   131070     3
                              172.16.2.1         ucst     1708     6 xe-0/0/2.0
                              172.16.3.1         ucst     1709     5 xe-0/0/4.0
```

Instead of a single next-hop, we see a 'ulst'. This stands for unicast list of next-hops. The vQFX will be able to use both links when traffic needs to be sent towards the 192.168.1.4 destination.

Leaf switches will load-share traffic that is destined towards the loopback IP address of the other leaf switches. Spine nodes will load-share traffic that is destined to the loopback IP of the other spine node.

Task 3.3: IBGP Overlay Configuration

- Configure a full-mesh of IBGP sessions sourced from the local loopback interface.
- The IBGP sessions should:
    - use AS65500
    - be authenticated using the MD5-hashed password 'jncie'
    - detect the loss of a BGP peer in 6 seconds
    - only support the exchange of EVPN routes
    - be able to load-balance traffic

Solution

In this task, we configure the signaling part of the overlay network. A similar configuration is applied to all the devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We'll first discuss the configuration in general. After this, we'll provide the configuration that is required on every single device.

The BGP sessions should use AS 65500. We can configure this on a group level, like this:

```
set protocols bgp group overlay local-as 65500
```

Since all the BGP peers will use the same AS number, we can define the group type as internal:

```
set protocols bgp group overlay type internal
```

To source the IBGP session from a specific IP address:

```
set protocols bgp group overlay local-address 192.168.1.x
```

A full-mesh of IBGP peers need to be configured, so we will need to create 4 IBGP neighbors on every device:

```
set protocols bgp group overlay neighbor 192.168.1.x
```

To make sure the BGP sessions uses the correct address family:

```
set protocols bgp group overlay family evpn signaling
```

To authenticate the BGP sessions configured in the overlay group:

```
set protocols bgp group overlay authentication-key jncie
```

To ensure the detection of a lost neighbor within 6 seconds, we configure BFD with the following parameters:

```
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
```

Multipathing can be configured without the 'multiple-as' options. This is because all IBGP session will use the same AS. We can also skip the forwarding-table policy because that was already configured in a previous task:

```
set protocols bgp group overlay multipath
```

The complete configuration that is required for this task is as the following:

vMX3:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.4
set protocols bgp group overlay neighbor 192.168.1.5
set protocols bgp group overlay neighbor 192.168.1.2
```

vMX4:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.4
set protocols bgp group overlay neighbor 192.168.1.5
set protocols bgp group overlay neighbor 192.168.1.1
```

vQFX3:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.3
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.2
set protocols bgp group overlay neighbor 192.168.1.4
set protocols bgp group overlay neighbor 192.168.1.5
```

vQFX4:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.2
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.5
```

vQFX5:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.5
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 3
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.2
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.4
```

Verification

To check the characteristics of the overlay group, we issue the following command:

```
jncie@vQFX3> show bgp group overlay

Group Type: Internal  AS: 65500         Local AS: 65500
 Name: overlay      Index: 1          Flags: <Export Eval>
 Options: <LocalAS>
 Holdtime: 0 Local AS: 65500 Local System AS: 65003
 Total peers: 4    Established: 4
 192.168.1.1+179
 192.168.1.2+62121
 192.168.1.4+179
 192.168.1.5+179
 bgp.evpn.0: 0/0/0/0
```

The BGP sessions that will be used for the overlay network are using AS65500. To examine additional information about the BGP neighbor, we can check the following output:

```
jncie@vQFX3> show bgp neighbor 192.168.1.4
Peer: 192.168.1.4+179 AS 65500 Local: 192.168.1.3+56862 AS 65500
 Group: overlay        Routing-Instance: master
 Forwarding routing-instance: master
 Type: Internal  State: Established  Flags: <Sync>
 Last State: OpenConfirm  Last Event: RecvKeepAlive
 Last Error: None
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>
 Authentication key is configured
 Address families configured: evpn
 Local Address: 192.168.1.3 Holdtime: 90 Preference: 170 Local AS: 65500 Local System AS:
65003
 Number of flaps: 0
 Peer ID: 192.168.1.4   Local ID: 192.168.1.3    Active Holdtime: 90
 Keepalive Interval: 30     Group index: 1  Peer index: 2
 BFD: enabled, down
 NLRI for restart configured on peer: evpn
 NLRI advertised by peer: evpn
 NLRI for this session: evpn
 Peer supports Refresh capability (2)
 Stale routes from peer are kept for: 300
 Peer does not support Restarter functionality
 Restart flag received from the peer: Notification
 NLRI that restart is negotiated for: evpn
 NLRI of received end-of-rib markers: evpn
 NLRI of all end-of-rib markers sent: evpn
 Peer does not support LLGR Restarter functionality
 Peer supports 4 byte AS extension (peer-as 65500)
 Peer does not support Addpath
 Table bgp.evpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:       0
  Received prefixes:      0
  Accepted prefixes:      0
  Suppressed due to damping:  0
 Last traffic (seconds): Received 14  Sent 14  Checked 63
 Input messages: Total 7   Updates 1    Refreshes 0   Octets 144
 Output messages: Total 8   Updates 0    Refreshes 0   Octets 226
```

We highlighted the options and the NLRI that is currently active for this session. We can narrow down the information that was displayed to present the same information in a more concise overview, like this:

```
jncie@vQFX3> show bgp neighbor 192.168.1.4 | match option
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>

jncie@vQFX3> show bgp neighbor 192.168.1.4 | match NLRI
 NLRI for restart configured on peer: evpn
 NLRI advertised by peer: evpn
 NLRI for this session: evpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

After combining match statements together and using some regular expressions, we can further reduce the information displayed and examine every BGP session on all of the nodes in a matter of seconds. We simply issue **show bgp neighbor | match "^Peer|option|NLRI.*this"** on every device:

```
jncie@vMX3> show bgp neighbor | match "^Peer|option|NLRI.*this"
Peer: 172.16.2.2+179 AS 65003 Local: 172.16.2.1+51713 AS 65001
 Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
 Options: <MultipathAs BfdEnabled>
 NLRI for this session: inet-unicast
Peer: 172.16.4.2+52335 AS 65004 Local: 172.16.4.1+179 AS 65001
 Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
 Options: <MultipathAs BfdEnabled>
 NLRI for this session: inet-unicast
Peer: 172.16.6.2+60598 AS 65005 Local: 172.16.6.1+179 AS 65001
 Options: <Preference AuthKey LogUpDown PeerAS Multipath Refresh>
 Options: <MultipathAs BfdEnabled>
 NLRI for this session: inet-unicast
Peer: 192.168.1.2+50947 AS 65500 Local: 192.168.1.1+179 AS 65500
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>
 NLRI for this session: evpn
Peer: 192.168.1.3+64946 AS 65500 Local: 192.168.1.1+179 AS 65500
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>
 NLRI for this session: evpn
Peer: 192.168.1.4+61612 AS 65500 Local: 192.168.1.1+179 AS 65500
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>
 NLRI for this session: evpn
Peer: 192.168.1.5+56293 AS 65500 Local: 192.168.1.1+179 AS 65500
 Options: <Preference LocalAddress AuthKey LogUpDown AddressFamily Multipath LocalAS Rib-group
Refresh>
 Options: <BfdEnabled>
 NLRI for this session: evpn
```

Even though we may have activated a new address family, we will not see any EVPN routes just yet:

```
jncie@vMX3> show route protocol bgp

inet.0: 17 destinations, 19 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.2/32     *[BGP/170] 00:04:47, localpref 100, from 172.16.2.2
                    AS path: 65003 65002 I, validation-state: unverified
                    to 172.16.2.2 via xe-0/0/2.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
           to 172.16.4.2 via xe-0/0/3.0
          > to 172.16.6.2 via xe-0/0/4.0
          [BGP/170] 00:04:47, localpref 100
           AS path: 65004 65002 I, validation-state: unverified
          > to 172.16.4.2 via xe-0/0/3.0
          [BGP/170] 00:04:47, localpref 100
           AS path: 65005 65002 I, validation-state: unverified
          > to 172.16.6.2 via xe-0/0/4.0
192.168.1.3/32   *[BGP/170] 00:04:47, localpref 100
           AS path: 65003 I, validation-state: unverified
          > to 172.16.2.2 via xe-0/0/2.0
192.168.1.4/32   *[BGP/170] 00:04:47, localpref 100
           AS path: 65004 I, validation-state: unverified
          > to 172.16.4.2 via xe-0/0/3.0
192.168.1.5/32   *[BGP/170] 00:04:47, localpref 100
           AS path: 65005 I, validation-state: unverified
          > to 172.16.6.2 via xe-0/0/4.0
```

EVPN routes will only start to appear after we configure VXLAN segments.

Task 3.4: Configure QFX Switch-Level EVPN

- Use ingress-replication and VXLAN encapsulation
- Use Lo0 as the VTEP source interface
- Use route-target 1:1 for the fabric

Solution

Configuring QFX switch-level EVPN means we are finalizing our controllerless overlay leaf switches and we are preparing them to support VXLAN segments.

We first define the switch-options. Here we instruct the switch to use the lo0 interface as the source for the VTEP. The IP connectivity between the different loopback interfaces was provided by the underlay. Therefore, the switches should be able to tunnel packets to each other.

Mandatory under the [**switch-options**] stanza is the configuration of a route-distinguisher. This route-distinguisher (RD), as we will see later on, is a value that is prepended to each EVPN route. For the sake of this task, the value of the route-distinguisher is not that relevant. We only need to ensure it is unique. Therefore, we configure a type 1 RD where we use the loopback IP and an extra arbitrary value of 1.

The last configuration tasks under the [**switch-options**] require us to configure the vrf-import and vrf-target. Through the vrf-import policy, which we will name ‘fabric-import’, we decide what EVPN routes we will accept from our BGP peers. For now, we only import the community that we configure as the vrf-target. Because we will configure ‘vrf-target

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

target:1:1', all the advertised EVPN routes will have the extended community route-target of 'target:1:1' attached to it.

We conclude the task configuring VXLAN encapsulation and ingress-replication under the [**protocols evpn**] stanza on every switch.

vQFX3:

```
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.3:1
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1

set policy-options policy-statement fabric-import term fabric from protocol bgp
set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1

set routing-options router-id 192.168.1.3

set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
```

vQFX4:

```
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.4:1
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1

set policy-options policy-statement fabric-import term fabric from protocol bgp
set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1

set routing-options router-id 192.168.1.4

set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX5:

```
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.5:1
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1

set policy-options policy-statement fabric-import term fabric from protocol bgp
set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1

set routing-options router-id 192.168.1.5

set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
```

Verification

To verify what source-address is being used for the VTEP interface:

```
jncie@vQFX3> show ethernet-switching vxlan-tunnel-end-point source

Logical System Name       Id  SVTEP-IP         IFL   L3-Idx
<default>                 0   192.168.1.3      lo0.0    0
```

For a more complete overview of the QFX switch-level EVPN configuration, we issue the following command:

```
jncie@vQFX3> show evpn instance extensive
Instance: __default_evpn__
  Route Distinguisher: 192.168.1.3:0
  Number of bridge domains: 0
  Number of neighbors: 0

  Instance: default-switch
    Route Distinguisher: 192.168.1.3:1
    Encapsulation type: VXLAN
    MAC database status                         Local  Remote
      MAC advertisements:                           0       0
      MAC+IP advertisements:                        0       0
      Default gateway MAC advertisements:           0       0
    Number of local interfaces: 0 (0 up)
    Number of IRB interfaces: 0 (0 up)
    Number of bridge domains: 0
    Number of neighbors: 0
    Number of ethernet segments: 0
    Router-ID: 192.168.1.3
    Source VTEP interface IP: 192.168.1.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We can see the route-distinguisher, the configured encapsulation type and the router-ID.

Still, even after enabling QFX-switch level EVPN, we will not see any EVPN routes:

```
jncie@vQFX3> show route protocol bgp

inet.0: 13 destinations, 15 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[BGP/170] 00:10:20, localpref 100
                      AS path: 65001 I, validation-state: unverified
                    > to 172.16.2.1 via xe-0/0/2.0
192.168.1.2/32     *[BGP/170] 00:13:57, localpref 100
                      AS path: 65002 I, validation-state: unverified
                    > to 172.16.3.1 via xe-0/0/4.0
192.168.1.4/32     *[BGP/170] 00:13:57, localpref 100, from 172.16.3.1
                      AS path: 65002 65004 I, validation-state: unverified
                    > to 172.16.2.1 via xe-0/0/2.0
                      to 172.16.3.1 via xe-0/0/4.0
                    [BGP/170] 00:01:21, localpref 100
                      AS path: 65001 65004 I, validation-state: unverified
                    > to 172.16.2.1 via xe-0/0/2.0
192.168.1.5/32     *[BGP/170] 00:10:20, localpref 100, from 172.16.2.1
                      AS path: 65001 65005 I, validation-state: unverified
                      to 172.16.2.1 via xe-0/0/2.0
                    > to 172.16.3.1 via xe-0/0/4.0
                    [BGP/170] 00:03:27, localpref 100
                      AS path: 65002 65005 I, validation-state: unverified
                    > to 172.16.3.1 via xe-0/0/4.0

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

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

After configuring a single switch, we can see the instantiation of the internal VTEP interface:

```
jncie@vQFX3> show interfaces vtep

Physical interface: vtep   , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 506
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed:
Unlimited
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Logical interface vtep.32768 (Index 564) (SNMP ifIndex 529)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.1.3, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

jncie@vQFX3>
```

This interface is used for internal purposes. When we configure the other vQFX switches with VXLAN segments, we can witness the instantiation of an additional VTEP interface per discovered endpoint. These interfaces, as we will see later on, will be used to tunnel data as well as learn MAC addresses.

Task 3.5: VXLAN Segment Configuration

- Configure a VXLAN segment for S1, S2 and S3 on vQFX3 and vQFX4.
- Traffic to and from the servers should be tagged with VLAN 100.
- Use route-target 100:100 for the VNI.

Solution

Under the switch-options, we configured 'vrf-target target:1:1'. If we were to simply configure a VXLAN segment and add the VNI to the list, the routes for the VXLAN segment would be advertised with 'target:1:1'. Since we are instructed to use 'target:100:100', we need to specify this route-target under the [**protocols evpn vni-options**] configuration stanza. We also need to add a term to our import policy, to make sure that routes with target:100:100 will be accepted when they are advertised by the other leaf switches.

After taking care of the [**protocols evpn**] stanza and after adding a term to the import policy, we are left with two more things to do: creating a VXLAN segment and assigning VXLAN/VLAN membership to the appropriate interfaces.

Note that we were instructed to create VLAN 100 and use VNI 100. It is important to understand that these two values are completely unrelated. For instance, we could also choose to use VNI 4123 for VLAN 100. The VNI is the VXLAN Network Identifier (or VXLAN segment ID). In the data plane, the VNI is a field in the VXLAN header. In the control plane, the VNI is sent along in the EVPN route to indicate to other routers that a certain MAC routes belongs to a certain VXLAN segment. The VLAN number and the VNI number do not need to match.

vQFX3:

```
set protocols evpn vni-options vni 100 vrf-target export target:100:100
set protocols evpn extended-vni-list 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept
set policy-options community vni-100 members target:100:100

set vlans vxlan-100 vlan-id 100 vxlan vni 100

set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 100
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode trunk vlan members 100
```

vQFX4:

```
set protocols evpn vni-options vni 100 vrf-target export target:100:100
set protocols evpn extended-vni-list 100

set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept
set policy-options community vni-100 members target:100:100

set vlans vxlan-100 vlan-id 100 vxlan vni 100

set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 100
```

Verification

After the configuration of the VXLAN segments, the vQFX switches will discover each other's VTEP interface:

```
jncie@vQFX3> show ethernet-switching vxlan-tunnel-end-point remote summary

Logical System Name       Id  SVTEP-IP         IFL   L3-Idx
<default>                 0   192.168.1.3      lo0.0    0
 RVTEP-IP         IFL-Idx   NH-Id
 192.168.1.4      553       1723
```

Before we verify the configuration on the vQFX switches, we first move to the vMX6 router to generate some traffic:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.2 rapid

PING 10.200.100.2 (10.200.100.2): 56 data bytes
!!!!!
--- 10.200.100.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 19.968/24.565/31.485/3.843 ms

jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.3 rapid
PING 10.200.100.3 (10.200.100.3): 56 data bytes
!!!!!
--- 10.200.100.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
round-trip min/avg/max/stddev = 36.804/50.792/66.485/11.599 ms
```

We can exchange traffic inside the VLAN that is configured locally on vQFX3 as well as within the VXLAN that is configured on both vQFX3 and vQFX4. Before moving over to the vQFX switches, we check the address information on the vMX6:

```
jncie@vMX6> show interfaces xe-0/0/2 | match hard

  Current address: 00:0c:29:8c:00:a8, Hardware address: 00:0c:29:8c:00:a8

jncie@vMX6> show arp vpn DC1-S1-VL100

MAC Address       Address         Name                Interface           Flags
00:0c:29:8c:00:b2 10.200.100.2    10.200.100.2        xe-0/0/2.100        none
00:0c:29:8c:00:bc 10.200.100.3    10.200.100.3        xe-0/0/2.100        none
Total entries: 2
```

We can see the MAC that was used to source the PING and we can see the MAC addresses used by the addresses that we send the PING to.

We now move to vQFX3 and start with the verification of the newly created VLAN/VXLAN:

```
jncie@vQFX3> show vlans 100 extensive

Routing instance: default-switch
  VLAN Name: vxlan-100                      State: Active
Tag: 100
Internal index: 3, Generation Index: 3, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : Yes
Interfaces:
    vtep.32769*,tagged,trunk
    xe-0/0/1.0*,tagged,trunk
    xe-0/0/3.0*,tagged,trunk
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 3
```

We can see that the VLAN is listed as active on three interfaces. It is active on the VTEP interface and on the interfaces xe-0/0/1 and xe-0/0/3. On those last two interfaces, we statically defined the membership.

Since we already generated several packets from the vMX6 router, we should see several entries active in the Ethernet switching table of the device:

```
jncie@vQFX3> show ethernet-switching table vlan-id 100

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vxlan-100           00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vxlan-100           00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vxlan-100           00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
```

We see that the top two entries were learned locally via interfaces xe-0/0/1 and xe-0/0/3. The third entry is learned via the VTEP interface. Listed in the last column is the 'Active source'. This is the IP address of the device that originally learned the MAC address in hardware before advertising it to the rest of the network. In this case, vQFX4 learned the MAC locally and updated vQFX3 using EVPN.

Moving over to vQFX4 will allow us to verify this and check to see if vQFX4 is learning MAC addresses from vQFX3 as well:

```
jncie@vQFX4> show ethernet-switching table vlan-id 100

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vxlan-100           00:0c:29:8c:00:a8   D        vtep.32769               192.168.1.3
   vxlan-100           00:0c:29:8c:00:b2   D        vtep.32769               192.168.1.3
   vxlan-100           00:0c:29:8c:00:bc   D        xe-0/0/1.0
```

For verification purposes, we move back to vQFX3. In case we only want to see what MAC address information was learned from remote switches, we can issue the following command:

```
jncie@vQFX3> show ethernet-switching vxlan-tunnel-end-point remote mac-table

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Logical system   : <default>
Routing instance : default-switch
 Bridging domain : vxlan-100+100, VLAN : 100, VNID : 100
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
   00:0c:29:8c:00:bc   D        vtep.32769       192.168.1.4
```

The newly instantiated VTEP interface, vtep.32769, is used for both MAC learning as well as tunneling data to the remote VTEP. The interface also keeps track of packets that have

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

been encapsulated/decapsulated. Let's clear the counters on the VTEP interface and check them again after sending some traffic from vMX6:

Note: You may see traffic traversing the tunnel, but the volume should be minimal any you should be able to validate that traffic counters were cleared, and you can see your traffic being sent through the tunnel.

```
jncie@vQFX3> clear interfaces statistics all

jncie@vQFX3> show interfaces vtep

Physical interface: vtep    , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 506
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed:
Unlimited
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 564) (SNMP ifIndex 529)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.1.3, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32769 (Index 553) (SNMP ifIndex 531)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.4, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode
```

The counters have been set to 0. Now we generate some traffic from vMX6. To make sure traffic passes through the VTEP interface, we send traffic from server 1 to server 3 inside VLAN 100:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.3 rapid count 10

PING 10.200.100.3 (10.200.100.3): 56 data bytes
!!!!!!!!!!
--- 10.200.100.3 ping statistics ---
10 packets transmitted, 10 packets received, 0% packet loss
round-trip min/avg/max/stddev = 28.325/34.313/42.419/5.017 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6>
```

We have been able to send and receive 10 packets.

```
jncie@vQFX3> show interfaces vtep

Physical interface: vtep    , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 506
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed:
Unlimited
  Device flags    : Present Running
  Link type       : Full-Duplex
  Link flags      : None
  Last flapped    : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 564) (SNMP ifIndex 529)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.1.3, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32769 (Index 553) (SNMP ifIndex 531)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.4, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 10
    Output packets: 10
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode
```

After configuring the QFX switches with a VXLAN segment, the switches will use EVPN signaling to ‘discover’ other switches interested in the same VXLAN segment. We used the command before in a previous task at a time when no VXLAN segments were active. Now, we issue the ‘show evpn instance extensive’ command again and observe the following:

```
jncie@vQFX3> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.3:0
  Number of bridge domains: 0
  Number of neighbors: 0

Instance: default-switch
  Route Distinguisher: 192.168.1.3:1
  Encapsulation type: VXLAN
  MAC database status                 Local  Remote
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  MAC advertisements:                          2        1
  MAC+IP advertisements:                       0        0
  Default gateway MAC advertisements:          0        0
Number of local interfaces: 2 (2 up)
  Interface name  ESI                                Mode             Status
  xe-0/0/1.0      00:00:00:00:00:00:00:00:00:00      single-homed     Up
  xe-0/0/3.0      00:00:00:00:00:00:00:00:00:00      single-homed     Up
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 1
  VLAN  VNI    Intfs / up    IRB intf   Mode             MAC sync  IM route label
  100   100        2   2                Extended         Enabled   100
Number of neighbors: 1
  192.168.1.4
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          0
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
  Number of ethernet segments: 0
  Router-ID: 192.168.1.3
  Source VTEP interface IP: 192.168.1.3
```

The command offers quite some output. It informs us about the characteristics of the switch we are currently logged into and we can see EVPN neighbors that with locally configured VXLAN segments. The output also tells us whether or not the EVPN neighbors have sent any MAC routes.

To examine the address information that is inside the EVPN database, we issue the following command:

```
jncie@vQFX3> show evpn database

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      100  00:0c:29:8c:00:a8  xe-0/0/1.0                     Mar 24 12:22:46
      100  00:0c:29:8c:00:b2  xe-0/0/3.0                     Mar 24 12:22:46
      100  00:0c:29:8c:00:bc  192.168.1.4                    Mar 24 12:23:40
```

Here we see that the addresses learned locally (in hardware) are placed inside the EVPN database. They are eligible for advertisement using BGP. We can see this using the following command:

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.4

default-switch.evpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.3:1::100::00:0c:29:8c:00:a8/304
*                         Self                         100        I
  2:192.168.1.3:1::100::00:0c:29:8c:00:b2/304
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*                              Self                          100        I
  3:192.168.1.3:1::100::192.168.1.3/304
*                              Self                          100        I
```

The two locally learned MAC addresses are advertised to the vQFX4 switch as an EVPN type 2 MAC route. We also see the EVPN type 3 route, the inclusive multicast route. This route is sent to indicate to the other switch (vQFX4) that the local switch (vQFX3) is an endpoint for VNI 100.

If we issue the same command with the keyword 'detail', additional characteristics are revealed:

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.4 detail

default-switch.evpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 2:192.168.1.3:1::100::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:1
     Route Label: 100
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: target:100:100 encapsulation0:0:0:0:vxlan

* 2:192.168.1.3:1::100::00:0c:29:8c:00:b2/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:1
     Route Label: 100
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: target:100:100 encapsulation0:0:0:0:vxlan

* 3:192.168.1.3:1::100::192.168.1.3/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:1
     Route Label: 100
     PMSI: Flags 0x0: Label 0x64: Type INGRESS-REPLICATION 192.168.1.3
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: target:100:100 encapsulation0:0:0:0:vxlan
     PMSI: Flags 0x0: Label 0x64: Type INGRESS-REPLICATION 192.168.1.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The MAC addresses are advertised with the configured route-target. The type 2 route also indicates what encapsulation is used and what the Ethernet Segment Identifier is. In this task, we statically assigned VLAN membership to 'standard' (or non-multihomed) interfaces. Therefore, we can see that the EVPN route is advertised using the default ESI value of '00:00:00:00:00:00:00:00:00:00'.

The type 3 route does not hold any MAC or IP address information. The inclusive multicast route signals to vQFX4 that vQFX3 is configured for VNI100 using target:100:100 and VXLAN encapsulation. Additionally, it indicated that the multicast mode for this VNI is set to ingress-replication. This describes how BUM traffic is handled for this VNI. With the mode set to ingress-replication, the ingress PE (currently, vQFX3 and vQFX4) forward BUM packets to remote PEs that belong to the same VXLAN segment. When EVPN signaling is used, Juniper networks devices only supports ingress replication.

Instead of verifying what routes were send, we can also learn what information vQFX3 has received from vQFX4:

```
jncie@vQFX3> show route receive-protocol bgp 192.168.1.4

inet.0: 13 destinations, 15 routes (13 active, 0 holddown, 0 hidden)

:vxlan.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)

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

bgp.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.4:1::100::00:0c:29:8c:00:bc/304
*                         192.168.1.4                  100        I
  3:192.168.1.4:1::100::192.168.1.4/304
*                         192.168.1.4                  100        I
```

In large scale deployments, the amount of addresses that are signaled can easily run into the tens of thousands. Let's move over to vQFX4 and examine the routing-table using the following command:

```
jncie@vQFX4> show route community-name vni-100

inet.0: 13 destinations, 17 routes (13 active, 0 holddown, 0 hidden)

:vxlan.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)

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

bgp.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:192.168.1.3:1::100::00:0c:29:8c:00:a8/304
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    *[BGP/170] 00:27:18, localpref 100, from 192.168.1.3
                      AS path: I, validation-state: unverified
                      to 172.16.4.1 via xe-0/0/2.0
                    > to 172.16.5.1 via xe-0/0/4.0
2:192.168.1.3:1::100::00:0c:29:8c:00:b2/304
                    *[BGP/170] 00:27:18, localpref 100, from 192.168.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.4.1 via xe-0/0/2.0
                      to 172.16.5.1 via xe-0/0/4.0
3:192.168.1.3:1::100::192.168.1.3/304
                    *[BGP/170] 01:46:55, localpref 100, from 192.168.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.4.1 via xe-0/0/2.0
                      to 172.16.5.1 via xe-0/0/4.0
```

Using the '**`community-name vni-100`**' option, we narrow down the output to routes that belong to VNI 100.

Task 3.6: Active-Active Multihoming

- Configure a VXLAN segment for S4 on vQFX4 and vQFX5.
- Configure the VXLAN segment as a single Ethernet Segment
- Use the VLAN and route-target that was used in the previous assignment.

Solution

To configure a multihomed segment, we need to follow these three steps:

1. Configure the chassis to support aggregated devices
2. Specify the member links that should become a part of the multihomed AE link
3. Define the characteristics of the AE link itself

These steps need to be carried out on all of the devices that are part of the multihomed segment.

Let's examine the AE link configuration that is specific to multihoming. First, we define the ESI. This stands for 'Ethernet Segment Identifier'. The default value is all zeros. We need to define a value that is a network-wide unique value. Furthermore, the value must start with '00'. As we will see later on, this value will be used when advertising EVPN routes.

We also specify a LACP system-ID. The reason for doing this is to make the device that is connected to the two leaf switches believe it is connected to a single device. The LACP system-ID value is send out in LACP packets. Since vMX6 will see LACP packets with the same LACP system ID value, it will regard the AE interface as an interface that is connected to a single device.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

In addition to this, we should not forget to configure VXLAN 100 on vQFX5 as well.

vQFX4:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/3 ether-options 802.3ad ae0

set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 11:11:11:11:11:11
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100
```

vQFX5:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 11:11:11:11:11:11
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set protocols evpn vni-options vni 100 vrf-target export target:100:100
set protocols evpn extended-vni-list 100

set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept
set policy-options community vni-100 members target:100:100

set vlans vxlan-100 vlan-id 100 vxlan vni 100
```

Verification

A lot is going on after completing the configuration tasks. It's best to break things down and verify the different parts here one at a time.

What we will first do is verify the aggregate Ethernet interface on vQFX4 and vQFX5. After this, we will examine how this Ethernet Segment is signaled. We will conclude our verification examining the VXLAN segment on vQFX5.

First, we check LACP by issuing the following command on both switches:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3                  Current   Fast periodic Collecting distributing

jncie@vQFX4> show lacp statistics interfaces

Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/3                 156         149            0            0
```

The 'Actor' refers to the local switch and the 'Partner' refers to the remote switch. We can see that the receive state is 'Current', meaning that we are receiving LACP PDUs. The LACP statistics is an interesting command in case you are having problems with the link. This will show you whether or not LACP packets are being received.

After verifying LACP operations, we check the local interfaces on vQFX4 and on vQFX5:

```
jncie@vQFX4> show interfaces ae0

Physical interface: ae0    , Enabled, Physical link is Up
  Interface index: 660, SNMP ifIndex: 529
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU 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:f2:ef, Hardware address: 02:05:86:72:f2:ef
  Ethernet segment value: 00:11:11:11:11:11:11:11:11:11, Mode: All-active
  Last flapped   : 2017-03-24 12:51:33 UTC (00:03:48 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 1440 bps (0 pps)

  Logical interface ae0.0 (Index 568) (SNMP ifIndex 532)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           221          0         23426          832
        Output:           441          1         51586         1848
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol eth-switch, MTU: 1514
      Flags: Trunk-Mode
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX5> show interfaces ae0

Physical interface: ae0    , Enabled, Physical link is Up
  Interface index: 660, SNMP ifIndex: 526
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU 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:bb:ef, Hardware address: 02:05:86:72:bb:ef
  Ethernet segment value: 00:11:11:11:11:11:11:11:11:11, Mode: All-active
  Last flapped   : 2017-03-24 12:51:30 UTC (00:03:17 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 984 bps (0 pps)

  Logical interface ae0.0 (Index 567) (SNMP ifIndex 530)
    Flags: Up SNMP-Traps 0x24024000 Encapsulation: Ethernet-Bridge
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :            44          0          2942            0
        Output:           199          0         25276         1016
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol eth-switch, MTU: 1514
      Flags: Is-Primary, Trunk-Mode
```

Both vQFX switches have a functioning AE interface that is operating at 10Gbit. The mode is set to ‘All-active’ which is the only supported mode on (v)QFX. Both switches are using the same Ethernet segment value. This is necessary because the vQFX switches need to signal to the rest of the network that these interfaces are one and the same Ethernet Segment. The statically configured LACP id that we configured was to make sure the device on the remote end, vMX6, would see the vQFX4 and vQFX5 as one entity.

Observe the following on vMX6:

```
jncie@vMX6> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 542
  Link-level type: Flexible-Ethernet, MTU: 1522, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:71:74:c0, Hardware address: 00:05:86:71:74:c0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Last flapped   : 2017-03-28 05:15:14 UTC (02:07:16 ago)
  Input rate     : 1984 bps (0 pps)
  Output rate    : 2584 bps (2 pps)
```

Now we check the EVPN instance on both vQFX4 as well as vQFX5:

```
jncie@vQFX4> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.4:0
  Number of bridge domains: 0
  Number of neighbors: 1
    192.168.1.5
      Received routes
        Ethernet Segment:                       1

Instance: default-switch
  Route Distinguisher: 192.168.1.4:1
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         1       3
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                               Mode             Status
    ae0.0           00:11:11:11:11:11:11:11:11:11     all-active       Up
    xe-0/0/1.0      00:00:00:00:00:00:00:00:00:00     single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  VNI    Intfs / up    IRB intf   Mode             MAC sync  IM route label
    100   100        2   2                Extended         Enabled   100
  Number of neighbors: 2
    192.168.1.3
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
    192.168.1.5
      Received routes
        MAC address advertisement:              1
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                2
  Number of ethernet segments: 1
    ESI: 00:11:11:11:11:11:11:11:11:11
      Status: Resolved by IFL ae0.0
      Local interface: ae0.0, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        192.168.1.5      100        0               all-active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Designated forwarder: 192.168.1.4
      Backup forwarder: 192.168.1.5
  Router-ID: 192.168.1.4
  Source VTEP interface IP: 192.168.1.4

jncie@vQFX5> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.5:0
  Number of bridge domains: 0
  Number of neighbors: 1
    192.168.1.4
      Received routes
        Ethernet Segment:                       1

Instance: default-switch
  Route Distinguisher: 192.168.1.5:1
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         1       3
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                               Mode             Status
    ae0.0           00:11:11:11:11:11:11:11:11:11     all-active       Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  VNI    Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100        1   1                Extended        Enabled   100
  Number of neighbors: 2
    192.168.1.3
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
    192.168.1.4
      Received routes
        MAC address advertisement:              1
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                2
  Number of ethernet segments: 1
    ESI: 00:11:11:11:11:11:11:11:11:11
      Status: Resolved by IFL ae0.0
      Local interface: ae0.0, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        192.168.1.4      0          0               all-active
      Designated forwarder: 192.168.1.4
      Backup forwarder: 192.168.1.5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Router-ID: 192.168.1.5
  Source VTEP interface IP: 192.168.1.5
```

Both switches are listing the ESI as active. The way they have 'found' each other is through the use of a community and a (hidden) policy that was automatically created and applied after the configuration of the ESI.

To check this Ethernet Segment (ES) routing information, we issue the following command:

```
jncie@vQFX4> show route advertising-protocol bgp 192.168.1.5 detail

...

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

* 1:192.168.1.4:0::11111111111111111::FFFF:FFFF/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.4:0
     Route Label: 1
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: target:1:1 encapsulation0:0:0:0:vxlan esi-label:all-active (label 0)

* 4:192.168.1.4:0::11111111111111111:192.168.1.4/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.4:0
     Route Label: 1
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: encapsulation0:0:0:0:vxlan es-import-target:11-11-11-11-11-11
```

The type 4 route is used to indicate to the other switches that an Ethernet Segment is active on this switch. At the same time, the switch will also apply a 'hidden' policy that will import routes with the community of the locally configured Ethernet Segment. We cannot examine this policy in the configuration of the switch.

Let's check what policies exist on the device:

```
jncie@vQFX4> show policy ?

Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-default-bd-100-internal__
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
__vrf-export-default-switch-internal__
__vrf-import-__default_evpn__-internal__
__vrf-import-default-switch-internal__
conditions              Show conditions used by policy
damping                 Show state of route flap damping
export-loopback
fabric-import
lbpp
|                       Pipe through a command
```

We issue the following command to see what the '__vrf-import-__default_evpn__-internal__' policy does:

```
jncie@vQFX4> show policy __vrf-import-__default_evpn__-internal__

Policy __vrf-import-__default_evpn__-internal__:
    Term unnamed:
        from community __vrf-community-__default_evpn__-import-internal__ [es-import-
target:11-11-11-11-11-11 ]
        then accept
    Term unnamed:
        then reject
```

The policy imports routes that have an ES-import route target attached to it. This policy will appear on devices that have a multihomed segment configured. The value of the community that will be imported corresponds to the configured ESI. Since both vQFX4 and vQFX5 were configured with the same Ethernet segment, both switches advertise and import a route that carries the same ES-import route target.

Both vQFX4 and 5 advertise the ES route to all other switches. This means they also advertise the route to vQFX3:

```
jncie@vQFX4> show route advertising-protocol bgp 192.168.1.3

default-switch.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.4:1::1111111111111111::0/304
*                         Self                         100        I
  2:192.168.1.4:1::100::00:0c:29:8c:00:bc/304
*                         Self                         100        I
  3:192.168.1.4:1::100::192.168.1.4/304
*                         Self                         100        I

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.4:0::1111111111111111::FFFF:FFFF/304
*                         Self                         100        I
  4:192.168.1.4:0::1111111111111111:192.168.1.4/304
*                         Self                         100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

However, because vQFX3 does not have the ESI configured on any of the interfaces, that switch will not import (or install) the route.

Now we move to vMX6 to verify connectivity:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.4 rapid

PING 10.200.100.4 (10.200.100.4): 56 data bytes
!!!!!
--- 10.200.100.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 34.273/40.358/56.222/8.156 ms
```

```
jncie@vMX6> show arp vpn DC1-S1-VL100 | match 4

00:05:86:71:7b:c0 10.200.100.4    10.200.100.4            xe-0/0/2.100            none
```

After we have generated some traffic, we should see that the vQFX4 and vQFX5 switches have learned this MAC address:

```
jncie@vQFX4> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vxlan-100           00:05:86:71:7b:c0   DR       ae0.0
   vxlan-100           00:0c:29:8c:00:a8   D        vtep.32769               192.168.1.3
   vxlan-100           00:0c:29:8c:00:b2   D        vtep.32769               192.168.1.3
   vxlan-100           00:0c:29:8c:00:bc   D        xe-0/0/1.0
```

```
jncie@vQFX5> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vxlan-100           00:05:86:71:7b:c0   DL       ae0.0
   vxlan-100           00:0c:29:8c:00:a8   D        vtep.32770               192.168.1.3
   vxlan-100           00:0c:29:8c:00:b2   D        vtep.32770               192.168.1.3
   vxlan-100           00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
```

Notice vQFX5 learns the MAC address locally and vQFX4 has learned the route from a remote source. What happened is that vQFX5 learned the MAC after receiving traffic.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX5 then updated vQFX4 through the use of a type 2 MAC route, keep in mind your environment may be opposite, which means the commands will need to be verified on the other device:

```
jncie@vQFX5> show route advertising-protocol bgp 192.168.1.4 detail

...

* 2:192.168.1.5:1::100::00:05:86:71:7b:c0/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.5:1
     Route Label: 100
     ESI: 00:11:11:11:11:11:11:11:11:11
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65500] I
     Communities: target:100:100 encapsulation0:0:0:0:vxlan
...
```

This MAC route contains the usual target and encapsulation. The thing that separates this MAC route from other routes is the non-default ESI. This MAC route is advertised with the ESI that the AE interface is configured with. Because of this additional information, vQFX4 knows that it should update its MAC table and associate the MAC address with that particular ESI.

Notice that vQFX3 has updated its switching table with this MAC address as well:

```
jncie@vQFX3> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vxlan-100           00:05:86:71:7b:c0   DR       esi.1727 00:11:11:11:11:11:11:11:11:11
   vxlan-100           00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vxlan-100           00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vxlan-100           00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
```

Notice that the vQFX3 associated this MAC address with an ESI interface. The vQFX3 does not import the type 4 routes that the vQFX4 and vQFX5 send out. It does however receive, and import, type 1 routes that indicates both vQFX4 as well as vQFX5 are a member of this Ethernet Segment:

```
jncie@vQFX3> show route receive-protocol bgp 192.168.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
inet.0: 13 destinations, 17 routes (13 active, 0 holddown, 0 hidden)

:vxlan.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

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

bgp.evpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.4:0::1111111111111111::FFFF:FFFF/304
*                         192.168.1.4                  100        I
  1:192.168.1.4:1::1111111111111111::0/304
*                         192.168.1.4                  100        I
  2:192.168.1.4:1::100::00:0c:29:8c:00:bc/304
*                         192.168.1.4                  100        I
  3:192.168.1.4:1::100::192.168.1.4/304
*                         192.168.1.4                  100        I

default-switch.evpn.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.4:0::1111111111111111::FFFF:FFFF/304
*                         192.168.1.4                  100        I
  1:192.168.1.4:1::1111111111111111::0/304
*                         192.168.1.4                  100        I
  2:192.168.1.4:1::100::00:0c:29:8c:00:bc/304
*                         192.168.1.4                  100        I
  3:192.168.1.4:1::100::192.168.1.4/304
*                         192.168.1.4                  100        I


jncie@vQFX3> show route receive-protocol bgp 192.168.1.5

inet.0: 13 destinations, 17 routes (13 active, 0 holddown, 0 hidden)

:vxlan.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

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

bgp.evpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.5:0::1111111111111111::FFFF:FFFF/304
*                         192.168.1.5                  100        I
  1:192.168.1.5:1::1111111111111111::0/304
*                         192.168.1.5                  100        I
  2:192.168.1.5:1::100::00:05:86:71:74:c0/304
*                         192.168.1.5                  100        I
  3:192.168.1.5:1::100::192.168.1.5/304
*                         192.168.1.5                  100        I

default-switch.evpn.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.5:0::1111111111111111::FFFF:FFFF/304
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*                          192.168.1.5                    100        I
  1:192.168.1.5:1::11111111111111111::0/304
*                          192.168.1.5                    100        I
  2:192.168.1.5:1::100::00:05:86:71:74:c0/304
*                          192.168.1.5                    100        I
  3:192.168.1.5:1::100::192.168.1.5/304
*                          192.168.1.5                    100        I
```

Task 3.7: Configure vMX Virtual-Switch-Level EVPN

- Configure a VLAN-aware EVPN service on both vMX3 and vMX4.
- Use ingress-replication and VXLAN encapsulation.
- Use Lo0 as the VTEP source interface.
- Use route-target 1:1 for the fabric.
- Import the VNI that is active on the leaf switches.

Solution

The configuration and capabilities between MX and QFX are very different. For instance, a (v)MX offers support for routing-instances of the type ‘virtual-switch’ and a QFX does not. When we need to configure a VLAN-aware EVPN service, we need to configure the EVPN inside a virtual-switch instance.

The router-id and the route-distinguisher-id can be configured outside of the routing-instance. The vMX will use the route-distinguisher-id when building BGP route advertisements for the ‘fabric’ instance. By configuring the route-distinguisher-id under the routing-options, the vMX router will generate a unique route-distinguisher per configured routing-instance. It would have also been possible to configure the route-distinguisher inside the routing-instance. This approach would require us to configure one for every individual routing-instance.

Inside the routing-instance, we can find configuration statements that are similar to the QFX configuration. There is one thing that stands out, and that is the ‘bridge-domain’ configuration. This is just semantics. A VLAN on MX is created by configuring a bridge domain.

vMX3:

```
set routing-options router-id 192.168.1.1
set routing-options route-distinguisher-id 192.168.1.1

set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-target target:1:1

set routing-instances fabric protocols evpn encapsulation vxlan
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances fabric protocols evpn extended-vni-list 100
set routing-instances fabric protocols evpn vni-options vni 100 vrf-target target:100:100
set routing-instances fabric protocols evpn multicast-mode ingress-replication

set routing-instances fabric bridge-domains bd_100 vlan-id 100
set routing-instances fabric bridge-domains bd_100 vxlan vni 100
```

vMX4:

```
set routing-options router-id 192.168.1.2
set routing-options route-distinguisher-id 192.168.1.2

set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-target target:1:1

set routing-instances fabric protocols evpn encapsulation vxlan
set routing-instances fabric protocols evpn extended-vni-list 100
set routing-instances fabric protocols evpn vni-options vni 100 vrf-target target:100:100
set routing-instances fabric protocols evpn multicast-mode ingress-replication

set routing-instances fabric bridge-domains bd_100 vlan-id 100
set routing-instances fabric bridge-domains bd_100 vxlan vni 100
```

Verification

Verifying the local and remote VTEPs:

```
jncie@vMX3> show l2-learning vxlan-tunnel-end-point source

Logical System Name       Id  SVTEP-IP         IFL   L3-Idx
<default>                 0   192.168.1.1      lo0.0    0
    L2-RTT                    Bridge Domain                VNID     MC-Group-IP
    fabric                    bd_100+100                   100      0.0.0.0

jncie@vMX3> show l2-learning vxlan-tunnel-end-point remote

Logical System Name       Id  SVTEP-IP         IFL   L3-Idx
<default>                 0   192.168.1.1      lo0.0    0
 RVTEP-IP         IFL-Idx   NH-Id
 192.168.1.3      341       606
    VNID          MC-Group-IP
    100           0.0.0.0
 RVTEP-IP         IFL-Idx   NH-Id
 192.168.1.4      342       607
    VNID          MC-Group-IP
    100           0.0.0.0
 RVTEP-IP         IFL-Idx   NH-Id
 192.168.1.5      343       609
    VNID          MC-Group-IP
    100           0.0.0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

On both vMX routers, we can see that the vQFX are discovered as neighbors for the newly configured EVPN instance:

```
jncie@vMX3> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 192.168.1.1:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           0       4
    MAC+IP advertisements:                        0       0
    Default gateway MAC advertisements:           0       0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 1
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  100   100            0    0               Extended        Enabled
Number of neighbors: 3
  192.168.1.3
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
  192.168.1.4
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
  192.168.1.5
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.5      100        0               all-active
      192.168.1.4      0          0               all-active
Router-ID: 192.168.1.1
Source VTEP interface IP: 192.168.1.1
```

Notice that the vMX supports routing-instances of the type ‘virtual-switch’. In this case, the instance called ‘fabric’ is such an instance. Due to the lack of virtual-switch support on vQFX, you will only see 1 instance on the vQFX called ‘default-switch’.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Since the vMX routers are not actively advertising anything (yet), they do not list each other as neighbors. Same goes for the vQFX switches, since the vMX routers haven't actively advertised anything, the vQFX switches only list other vQFX switches as neighbors for this EPVN.

On the vMX, we can also see that several VTEP interfaces are instantiated (one for each neighbor):

```
jncie@vMX3> show interfaces vtep

Physical interface: vtep, Enabled, Physical link is Up
  Interface index: 135, SNMP ifIndex: 510
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed:
Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 340) (SNMP ifIndex 521)
    Flags: Up SNMP-Traps 0x4000 Encapsulation: ENET2
    Ethernet segment value: 00:00:00:00:00:00:00:00:00:00, Mode: single-homed, Multi-homed
status: Forwarding
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.1.1, L2 Routing Instance:
fabric, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32769 (Index 341) (SNMP ifIndex 522)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.3, L2 Routing Instance:
fabric, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0
    Protocol bridge, MTU: Unlimited
      Flags: Is-Primary, Trunk-Mode

  Logical interface vtep.32770 (Index 342) (SNMP ifIndex 523)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.4, L2 Routing Instance:
fabric, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0
    Protocol bridge, MTU: Unlimited
      Flags: Trunk-Mode

  Logical interface vtep.32771 (Index 343) (SNMP ifIndex 524)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.5, L2 Routing Instance:
fabric, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0
    Protocol bridge, MTU: Unlimited
      Flags: Trunk-Mode
```

And the vMX is already learning MAC address information across these interfaces as well:

```
jncie@vMX3> show l2-learning vxlan-tunnel-end-point remote mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Logical system   : <default>
Routing instance : fabric
 Bridging domain : bd_100+100, VLAN : 100, VNID : 100
   MAC                 MAC      Logical          Remote VTEP
   address             flags    interface        IP address
   00:05:86:71:74:c0   DR       esi.626          192.168.1.5
   00:0c:29:8c:00:a8   D        vtep.32769       192.168.1.3
   00:0c:29:8c:00:b2   D        vtep.32769       192.168.1.3
   00:0c:29:8c:00:bc   D        vtep.32770       192.168.1.4
```

The vQFX switches only list each other as neighboring devices and they will not have a VTEP interface for the vMX yet:

```
jncie@vQFX5> show interfaces vtep

Physical interface: vtep   , Enabled, Physical link is Up
  Interface index: 641, SNMP ifIndex: 506
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: Unlimited, Speed:
Unlimited
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 565) (SNMP ifIndex 527)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 192.168.1.5, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32769 (Index 568) (SNMP ifIndex 532)
    Flags: Up SNMP-Traps Encapsulation: ENET2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.4, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 2
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode

  Logical interface vtep.32770 (Index 569) (SNMP ifIndex 533)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.1.3, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 11
    Output packets: 5
    Protocol eth-switch, MTU: Unlimited
      Flags: Trunk-Mode

jncie@vQFX5> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.5:0
  Number of bridge domains: 0
  Number of neighbors: 1
    192.168.1.4
      Received routes
        Ethernet Segment:                    1

Instance: default-switch
  Route Distinguisher: 192.168.1.5:1
  Encapsulation type: VXLAN
  MAC database status                   Local  Remote
    MAC advertisements:                     1       3
    MAC+IP advertisements:                  0       0
    Default gateway MAC advertisements:     0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                            Mode             Status
    ae0.0           00:11:11:11:11:11:11:11:11:11  all-active       Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  VNI    Intfs / up    IRB intf   Mode             MAC sync  IM route label
    100   100        1   1                Extended         Enabled   100
  Number of neighbors: 2
    192.168.1.3
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
    192.168.1.4
      Received routes
        MAC address advertisement:              1
        MAC+IP address advertisement:           0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved by IFL ae0.0
    Local interface: ae0.0, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.4      0          0               all-active
    Designated forwarder: 192.168.1.4
    Backup forwarder: 192.168.1.5
Router-ID: 192.168.1.5
Source VTEP interface IP: 192.168.1.5
```

Also notice that on vMX, there is a difference in the way that import, and export policies are applied. We have created a separate virtual-switch called 'fabric' on both spines. Unlike on the vQFX, we did not configure an import policy. Let's look at what policies are currently active on the vMX:

```
jncie@vMX4> show policy ?

Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-bd-override-100-internal__
  __vrf-export-fabric-internal__
  __vrf-import-autoderive-fabric-internal__
  __vrf-import-fabric-internal__
  conditions           Show conditions used by policy
  damping              Show state of route flap damping
  export-loopback
  fabric               Internal fabric state
  lbpp
  logical-system       Name of logical system, or 'all'
  statistics           Statistics for policy
  |                    Pipe through a command
```

We can see that four policies were automatically created by the vMX routers. Two policies, *__vrf-export-fabric-internal__* and *__vrf-export-fabric-internal__*, are the result of the 'vrf-target' statement. These policies can be examined by issuing the following commands:

```
jncie@vMX4> show policy __vrf-import-fabric-internal__

Policy __vrf-import-fabric-internal__:
    Term unnamed:
        from community __vrf-community-fabric-common-internal__ [target:1:1 ]
        then accept
    Term unnamed:
        then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX4> show policy __vrf-export-fabric-internal__

Policy __vrf-export-fabric-internal__:
    Term unnamed:
        then community + __vrf-community-fabric-common-internal__ [target:1:1 ] accept
```

These policies tell us that the vMX will import and export routes for this routing-instance with the configured target value.

The additional policies that were configured were the result of the configuration of the VXLAN segments under the [routing-instances fabric protocols evpn vni-options] stanza.

One of those two policies is an export policy to export VXLAN segments with the configured target:

```
jncie@vMX4> show policy __vrf-export-bd-override-100-internal__

Policy __vrf-export-bd-override-100-internal__:
    Term unnamed:
        then community + __vrf-community-bd-override-100-export-internal__ [target:100:100 ]
accept
```

The other is an import policy that will import routes:

```
jncie@vMX4> show policy __vrf-import-autoderive-fabric-internal__

Policy __vrf-import-autoderive-fabric-internal__:
    Term unnamed:
        from community __vrf-community-bd-override-100-export-internal__ [target:100:100 ]
        then accept
    Term unnamed:
        from community __vrf-community-fabric-common-internal__ [target:1:1 ]
        then accept
    Term unnamed:
        then reject
```

This import policy is made up out of several terms. The first term imports routes that have the target:100:100 community attached to them (VXLAN 100) , the second term is accepts routes with target:1:1 and the last term rejects all other routes.

Task 3.8: Redundant Layer 3 VXLAN Gateway Configuration

- Configure a layer 3 gateway for the VXLAN segment offered to servers located in VLAN 100 on both vMX routers
- Use IP 10.200.100.252/24 on vMX3 and IP 10.200.100.253/24 on vMX4. Configure an IP gateway on the routers for hosts inside VNI 100. Use IP address 10.200.100.254 for this.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Solution

The spines of this overlay can be used to offer a redundant gateway to devices that are connected to the overlay network. To configure a layer 3 interface for a bridging domain, we configure an IRB interface and reference that interface from the bridge-domain configuration.

The IRB interface is configured with two addresses. One is a 'regular' address and the other address is a virtual-gateway-address. The regular address has to be unique on both nodes. This is the IP that the node itself will listen to. The virtual-gateway-address will be the address that other device will use as their default gateway.

Configuring the virtual-gateway-address will make the routers respond to ARP requests for that gateway address when as well as make the vMX routers advertise a virtual MAC address to the other leaf switches.

vMX3:

```
set routing-instances fabric bridge-domains bd_100 routing-interface irb.100

set interfaces irb unit 100 family inet address 10.200.100.252/24 virtual-gateway-address
10.200.100.254
```

vMX4:

```
set routing-instances fabric bridge-domains bd_100 routing-interface irb.100

set interfaces irb unit 100 family inet address 10.200.100.253/24 virtual-gateway-address
10.200.100.254
```

Verification

We configured an IRB interface on both vMX routers and when we examine the EVPN instance, we now see the following:

```
jncie@vMX3> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 192.168.1.1:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           1       7
    MAC+IP advertisements:                        2       2
    Default gateway MAC advertisements:           2       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route
...
```

We check the routing information advertised from vMX to vQFX:

```
jncie@vMX3> show route advertising-protocol bgp 192.168.1.5

fabric.evpn.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.1:8::100::00:00:5e:00:01:01/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.1:8::100::00:05:86:4d:34:f0/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.1:8::100::00:00:5e:00:01:01::10.200.100.254/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.1:8::100::00:05:86:4d:34:f0::10.200.100.252/304 MAC/IP
*                         Self                         100        I
  3:192.168.1.1:8::100::192.168.1.1/304 IM
*                         Self                         100        I

__default_evpn__.evpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.1:0::050000fde90000006400::FFFF:FFFF/304 AD/ESI
*                         Self                         100        I

jncie@vMX4> show route advertising-protocol bgp 192.168.1.5

fabric.evpn.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.2:8::100::00:00:5e:00:01:01/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::100::00:05:86:82:68:f0/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::100::00:00:5e:00:01:01::10.200.100.254/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::100::00:05:86:82:68:f0::10.200.100.253/304 MAC/IP
*                         Self                         100        I
  3:192.168.1.2:8::100::192.168.1.2/304 IM
*                         Self                         100        I

__default_evpn__.evpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  1:192.168.1.2:0::050000fdea0000006400::FFFF:FFFF/304 AD/ESI
*                         Self                         100        I
```

We can see that both vMX routers are advertising their own ‘real’ IP as well as the virtual-gateway address. This virtual-gateway is always bound to the same MAC address, which is 00:00:5e:00:01:01. Additional IRB interfaces configured with a virtual-gateway address will use the exact same address.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Now we check the connectivity from vMX6:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.254 rapid

PING 10.200.100.254 (10.200.100.254): 56 data bytes
.....
--- 10.200.100.254 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

We cannot send traffic to the virtual-gateway address, but this is normal. We can send traffic to the 'real' IP addresses configured on vMX3 and vMX4:

```
jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.252 rapid

PING 10.200.100.252 (10.200.100.252): 56 data bytes
!!!!!
--- 10.200.100.252 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 32.860/35.197/40.930/2.933 ms

jncie@vMX6> ping routing-instance DC1-S1-VL100 10.200.100.253 rapid

PING 10.200.100.253 (10.200.100.253): 56 data bytes
!!!!!
--- 10.200.100.253 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 11.664/23.210/50.805/14.532 ms
```

After this, we can check the ARP entries on the vMX:

```
jncie@vMX3> show arp interface irb.100

MAC Address       Address         Name                      Interface               Flags
00:0c:29:8c:00:a8 10.200.100.1    10.200.100.1              irb.100 [vtep.32769]    none
00:00:5e:00:01:01 10.200.100.254  10.200.100.254            irb.100                 permanent
published gateway

jncie@vMX4> show arp interface irb.100

MAC Address       Address         Name                      Interface               Flags
00:0c:29:8c:00:a8 10.200.100.1    10.200.100.1              irb.100 [vtep.32769]    none
00:00:5e:00:01:01 10.200.100.254  10.200.100.254            irb.100                 permanent
published gateway
```

Task 3.9: Inter-VXLAN Routing

- Configure an additional VXLAN segment for all servers in VLAN 101 using the following parameters:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

  - VLAN 101
  - route-target 101:101
  - Subnet 10.200.101.0/24 (10.200.101.252 on vMX3, 10.200.101.253 on vMX4 and 10.200.100.254 as VIP)

- Add the VXLAN segment on all server-facing interfaces.

Solution

vMX3:

```
set routing-instances fabric protocols evpn extended-vni-list 101
set routing-instances fabric protocols evpn vni-options vni 101 vrf-target target:101:101
set routing-instances fabric bridge-domains bd_101 vlan-id 101
set routing-instances fabric bridge-domains bd_101 vxlan vni 101
set routing-instances fabric bridge-domains bd_101 routing-interface irb.101
set interfaces irb unit 101 family inet address 10.200.101.252/24 virtual-gateway-address
10.200.101.254
```

vMX4:

```
set routing-instances fabric protocols evpn extended-vni-list 101
set routing-instances fabric protocols evpn vni-options vni 101 vrf-target target:101:101
set routing-instances fabric bridge-domains bd_101 vlan-id 101
set routing-instances fabric bridge-domains bd_101 vxlan vni 101
set routing-instances fabric bridge-domains bd_101 routing-interface irb.101
set interfaces irb unit 101 family inet address 10.200.101.253/24 virtual-gateway-address
10.200.101.254
```

vQFX3:

```
set protocols evpn extended-vni-list 101
set protocols evpn vni-options vni 101 vrf-target export target:101:101

set policy-options policy-statement fabric-import term vni-101 from community vni-101
set policy-options policy-statement fabric-import term vni-101 then accept
set policy-options community vni-101 members target:101:101

set vlans vxlan-101 vlan-id 101 vxlan vni 101

set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 101
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode trunk vlan members 101
```

vQFX4:

```
set protocols evpn extended-vni-list 101
set protocols evpn vni-options vni 101 vrf-target export target:101:101

set policy-options policy-statement fabric-import term vni-101 from community vni-101
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options policy-statement fabric-import term vni-101 then accept
set policy-options community vni-101 members target:101:101

set vlans vxlan-101 vlan-id 101 vxlan vni 101

set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 101
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members 101
```

vQFX5:

```
set protocols evpn extended-vni-list 101 vni-options vni 101 vrf-target export target:101:101

set policy-options policy-statement fabric-import term vni-101 from community vni-101
set policy-options policy-statement fabric-import term vni-101 then accept
set policy-options community vni-101 members target:101:101

set vlans vxlan-101 vlan-id 101 vxlan vni 101

set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members 101
```

Verification

From vMX6, first verify the connectivity within the newly configured VXLAN segment:

```
jncie@vMX6> ping routing-instance DC1-S1-VL101 10.200.101.4 rapid

PING 10.200.101.4 (10.200.101.4): 56 data bytes
!!!!!
--- 10.200.101.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 38.486/50.916/69.207/12.735 ms

jncie@vMX6> ping routing-instance DC1-S1-VL101 10.200.101.252 rapid

PING 10.200.101.252 (10.200.101.252): 56 data bytes
!!!!!
--- 10.200.101.252 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 27.883/44.716/65.957/12.839 ms

jncie@vMX6> ping routing-instance DC1-S1-VL101 10.200.101.253 rapid

PING 10.200.101.253 (10.200.101.253): 56 data bytes
!!!!!
--- 10.200.101.253 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.919/13.093/15.764/1.936 ms

jncie@vMX6>
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We then verify the connectivity to the other VXLAN segment by sending traffic to one of the hosts inside the 10.200.100.0/24 subnet:

```
jncie@vMX6> ping routing-instance DC1-S1-VL101 10.200.100.1 rapid

PING 10.200.100.1 (10.200.100.1): 56 data bytes
!!!!!
--- 10.200.100.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 31.908/36.039/41.477/3.180 ms
```

Part 2: IBGP-based Overlay Network

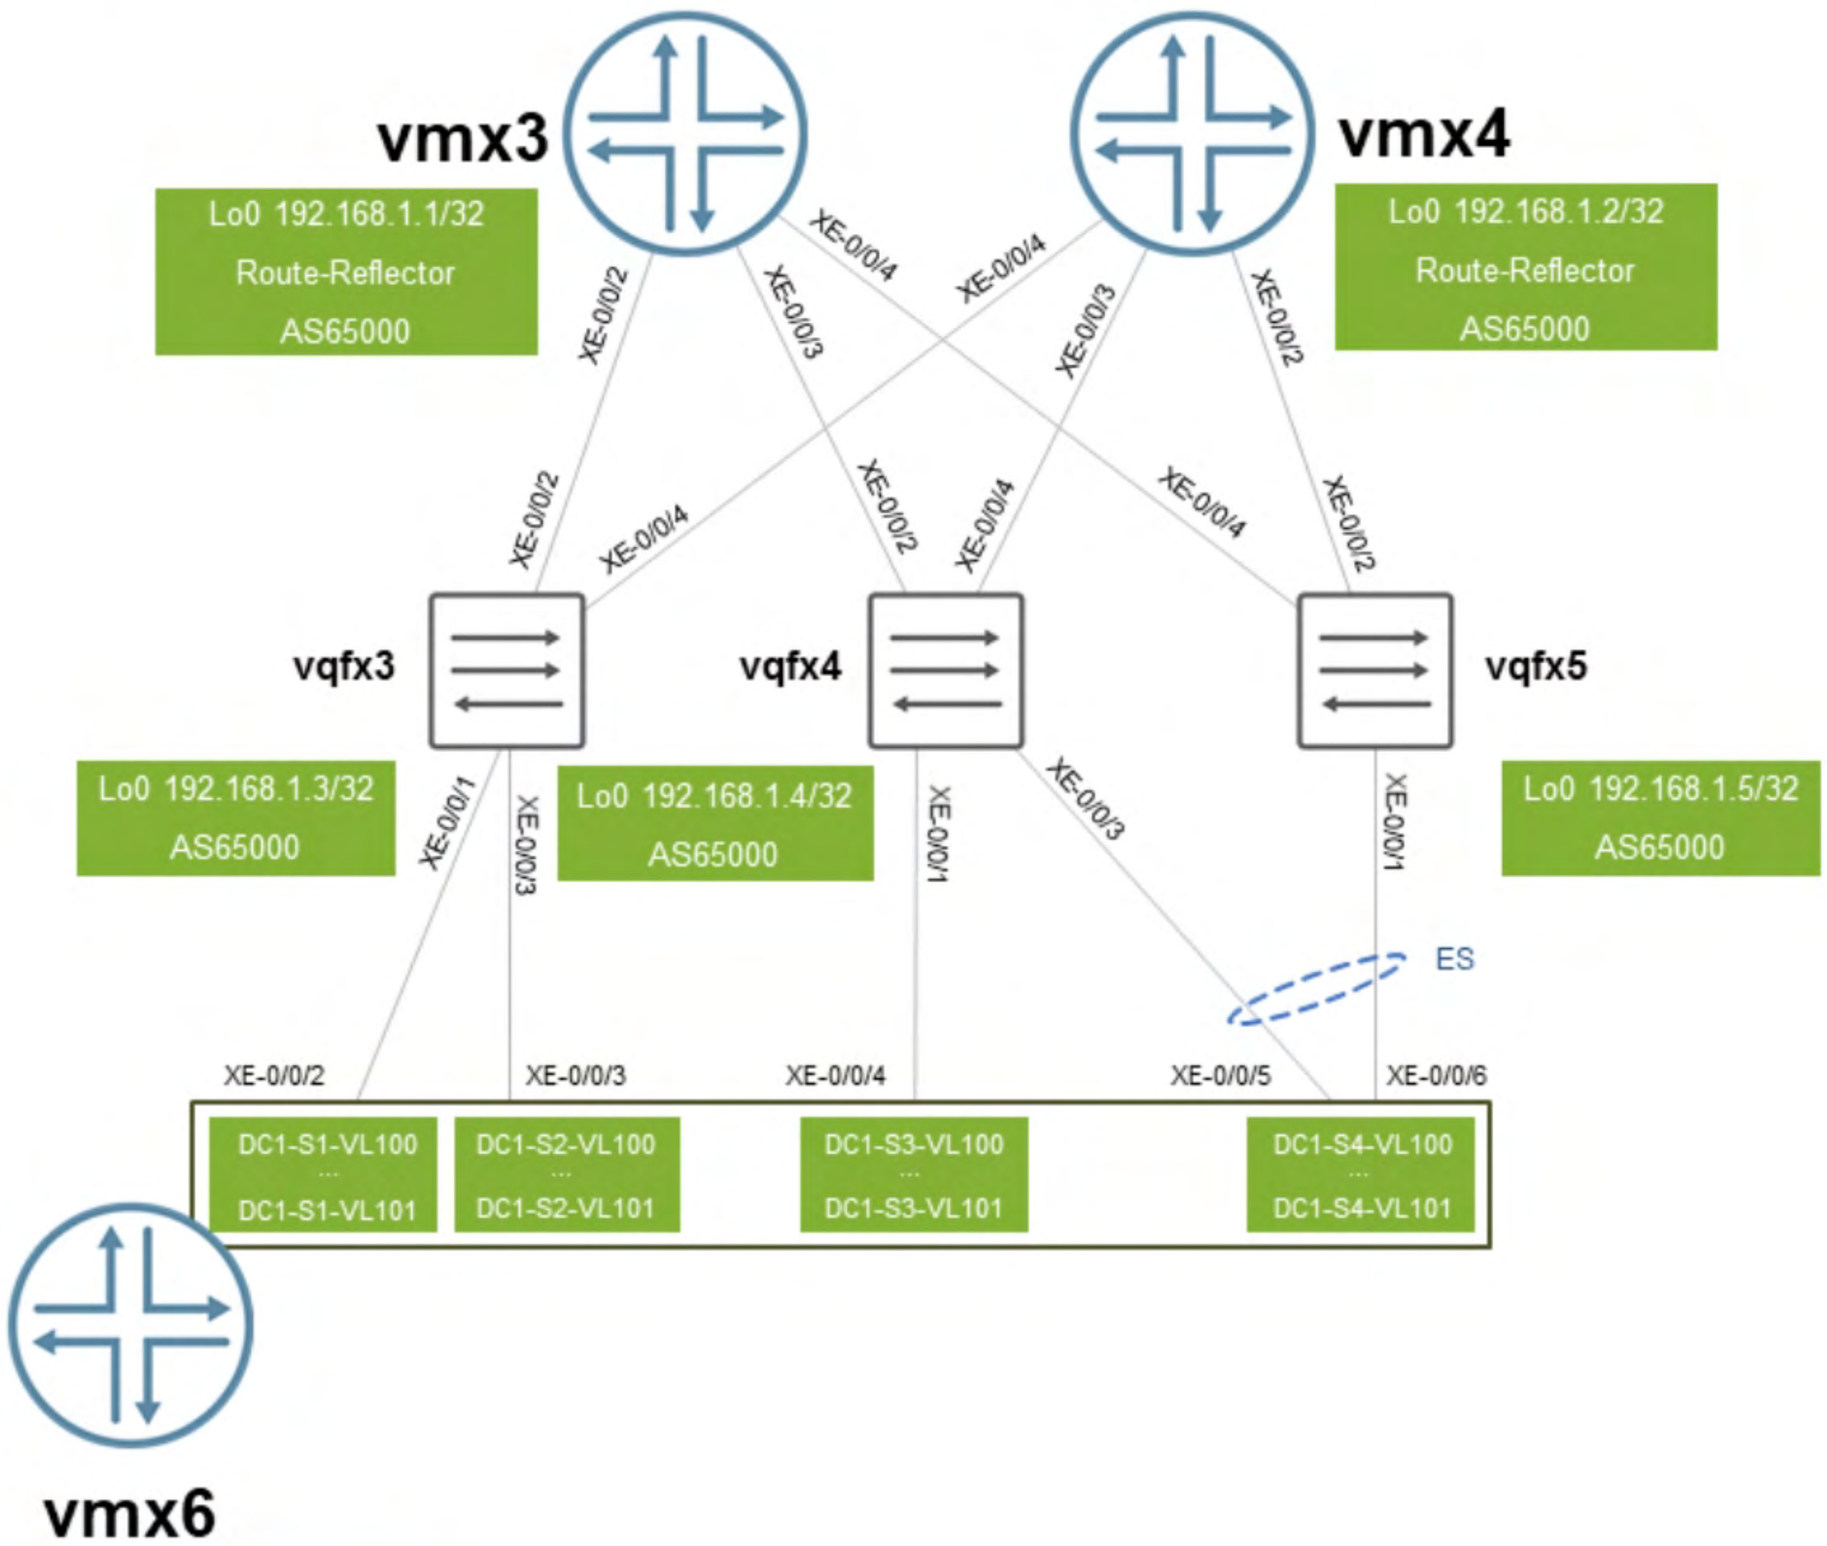

Figure 20 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

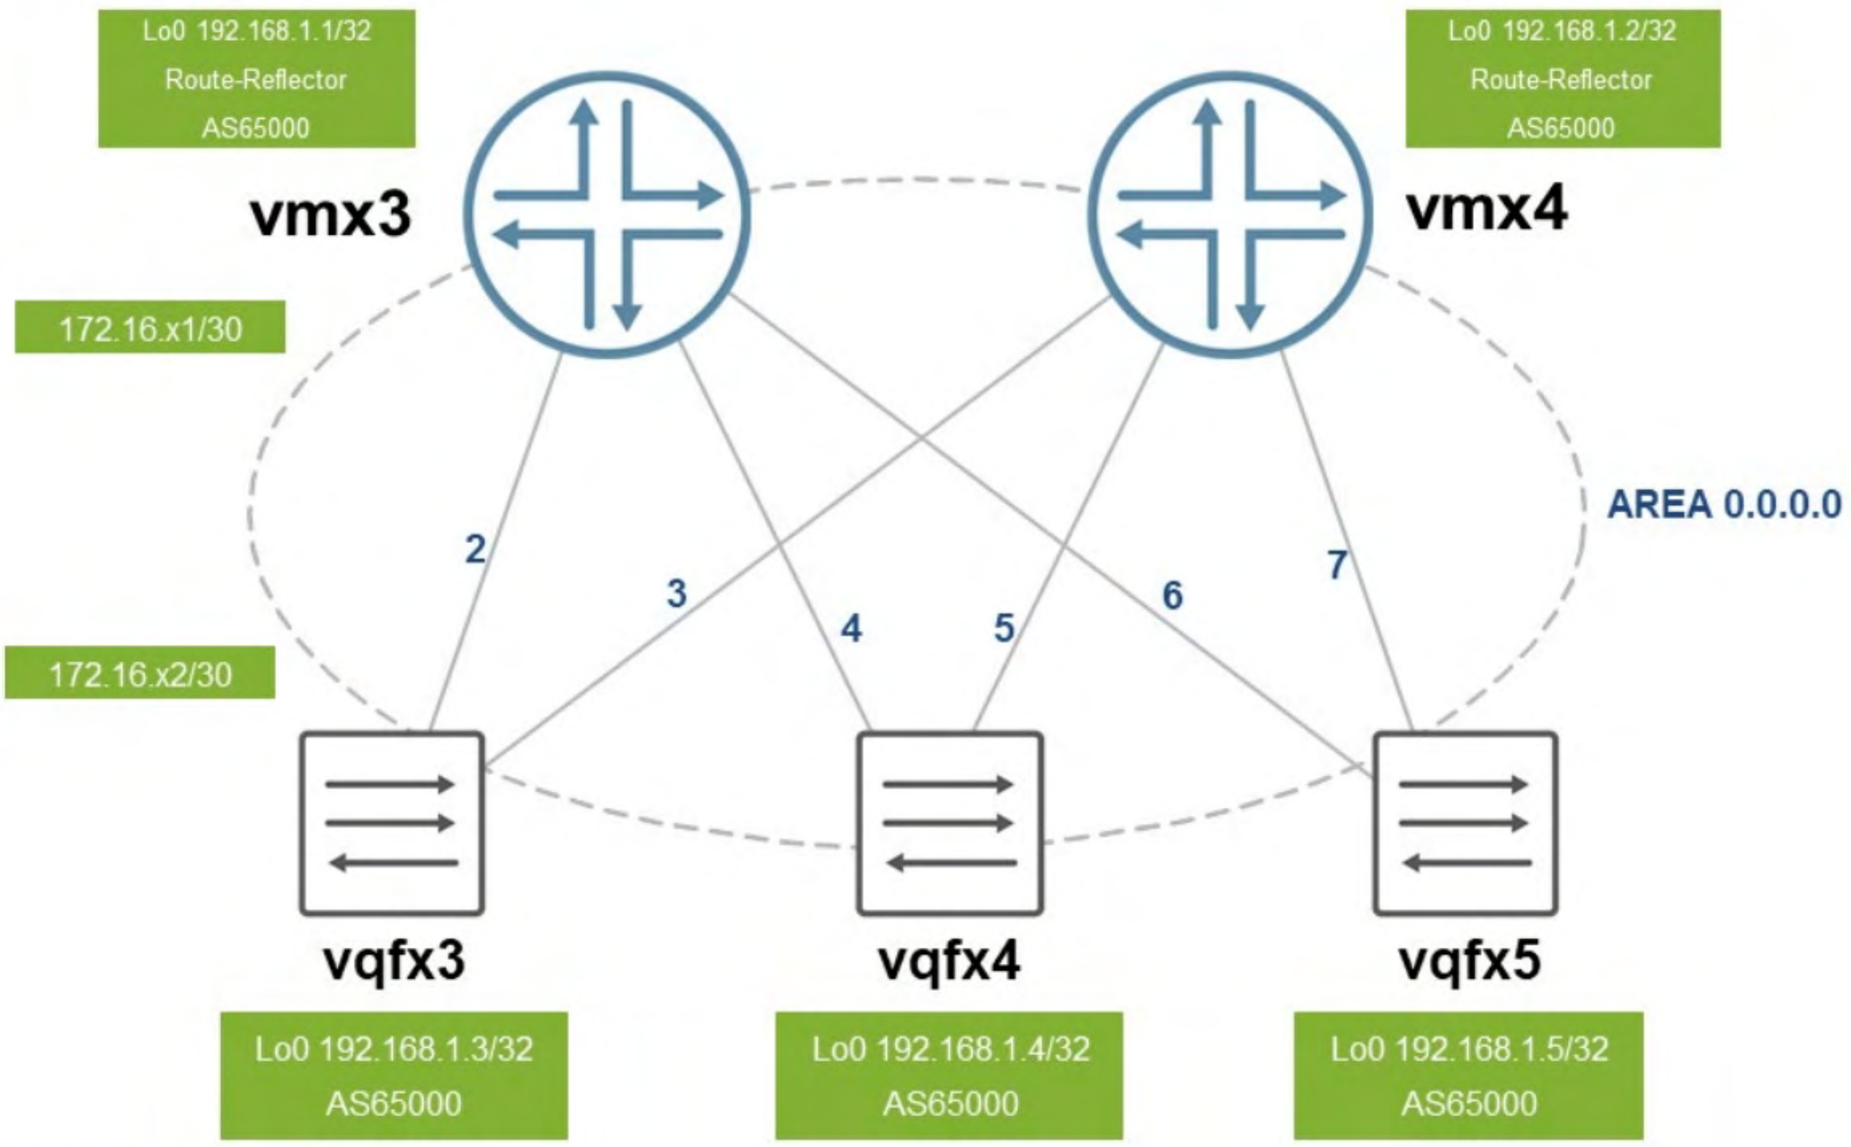

Figure 2 - Logical topology

Note: the starting topology is configured with IP-addressing only.

Task 3.10: IGP Configuration

- Configure a single area OSPF the devices listed in the topology.
- Ensure IP-connectivity between the loopback IP addresses.

Solution

For this task, we activate OSPF on all of the links between the vMX and vQFX. Additionally, we activate OSPF for the loopback interface. By doing so, we inject the loopback IP into the OSPF database. Since we do not aim to form OSPF neighbor relationships across the loopback interface, best-practices are to configure the lo0 interface as a passive interface.

vMX3:

```
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0
```

vQFX3:

```
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0
```

vQFX4:

```
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0
```

vQFX5:

```
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0
```

Verification

All of the leaf switches should form an OSPF neighbor relationship with both of the spine routers. Therefore, the easiest place to check the OSPF neighbor relationships is on vMX3 and vMX4:

```
jncie@vMX3> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
172.16.2.2       xe-0/0/2.0             Full      192.168.1.3      128    36
172.16.4.2       xe-0/0/3.0             Full      192.168.1.4      128    31
172.16.6.2       xe-0/0/4.0             Full      192.168.1.5      128    31

jncie@vMX4> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
172.16.7.2       xe-0/0/2.0             Full      192.168.1.5      128    30
172.16.5.2       xe-0/0/3.0             Full      192.168.1.4      128    37
172.16.3.2       xe-0/0/4.0             Full      192.168.1.3      128    39
```

On every device, we included the lo0 interface in the OSPF stanza configuration. To verify if every device is advertising its configured loopback IP, we can issue the following command:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show route protocol ospf terse

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

A V Destination         P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.3.0/30       O  10          2             >172.16.2.2
* ? 172.16.5.0/30       O  10          2             >172.16.4.2
* ? 172.16.7.0/30       O  10          2             >172.16.6.2
* ? 192.168.1.2/32      O  10          2              172.16.2.2
                                                     172.16.4.2
                                                    >172.16.6.2
* ? 192.168.1.3/32      O  10          1             >172.16.2.2
* ? 192.168.1.4/32      O  10          1             >172.16.4.2
* ? 192.168.1.5/32      O  10          1             >172.16.6.2
* ? 224.0.0.5/32        O  10          1              MultiRecv
```

We see that the loopback IP addresses configured on the other device have appeared in the routing table. To see if vMX3 is also advertising its own loopback, we can check the routing table on another device:

```
jncie@vMX4> show route protocol ospf terse

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

A V Destination         P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.2.0/30       O  10          2             >172.16.3.2
* ? 172.16.4.0/30       O  10          2             >172.16.5.2
* ? 172.16.6.0/30       O  10          2             >172.16.7.2
* ? 192.168.1.1/32      O  10          2              172.16.7.2
                                                    >172.16.5.2
                                                     172.16.3.2
* ? 192.168.1.3/32      O  10          1             >172.16.3.2
* ? 192.168.1.4/32      O  10          1             >172.16.5.2
* ? 192.168.1.5/32      O  10          1             >172.16.7.2
* ? 224.0.0.5/32        O  10          1              MultiRecv
```

In addition to this, we can also choose to issue the following command:

```
jncie@vMX3> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.1.1      192.168.1.1      0x80000007   223  0x22 0xa389  72
  bits 0x0, link count 4
  id 192.168.1.1, data 255.255.255.255, Type Stub (3)

…
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

This way, we could have verified everything from the vMX3 router itself.

Task 3.11: IGP Optimization

- Make sure traffic will be load-balanced across equal cost paths.
- OSPF adjacency loss should be detected within 10 seconds.
- Ensure that no type 2 LSAs are found within Area 0.0.0.0.
- Authenticate OSPF using MD5 authentication. Use the password 'jncie'.

Solution

No OSPF specific configuration is required for load-balancing. To prepare a Juniper device for load-balancing, we need to configure a policy that includes the 'load-balance per-packet' statement. This policy needs to be applied to the forwarding table.

To make sure we detect an adjacency loss within 10 seconds, we configure BFD with a minimum interval of 2 seconds. As we will see later on, the multiplier of 4 will leave us with an 8 second detection time.

When OSPF is activated on a link, the device will generate a type 2 LSA for that segment. To ensure that type 2 LSAs are not generated for these links, we specify them as point-to-point links. This also speeds up the OSPF neighbor adjacency formation since the routers will not have to elect a DR and a BDR router.

Finally, we enable authentication for OSPF. When using MD5 authentication, we need to make sure that the key-id and the password match on both OSPF routers.

On every vMX:

```
set routing-options forwarding-table export lbpp

set policy-options policy-statement lbpp term lbpp then load-balance per-packet

set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 bfd-liveness-detection minimum-interval
2000
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 authentication md5 1 key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/3.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0 bfd-liveness-detection minimum-interval
2000
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0 authentication md5 1 key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 bfd-liveness-detection minimum-interval
2000
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 authentication md5 1 key jncie
```

On every vQFX:

```
set routing-options forwarding-table export lbpp

set policy-options policy-statement lbpp term lbpp then load-balance per-packet

set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 bfd-liveness-detection minimum-interval
2000
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0 authentication md5 1 key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 bfd-liveness-detection minimum-interval
2000
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface xe-0/0/4.0 authentication md5 1 key jncie
```

Verification

First, we verify whether or not every OSPF neighbor relationship recovered after making the configuration changes:

```
jncie@vMX3> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
172.16.2.2       xe-0/0/2.0             Full      192.168.1.3      128    33
172.16.4.2       xe-0/0/3.0             Full      192.168.1.4      128    33
172.16.6.2       xe-0/0/4.0             Full      192.168.1.5      128    35
```

We do not want any type 2 LSAs in the OSPF database. To this end, we configured the OSPF links as point-to-point links.

Before configuring ‘interface-type p2p’, the OSPF database included a Network LSA for all of the links between the leaf and spine devices:

```
jncie@vMX3> show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.1.1      192.168.1.1      0x80000007   423  0x22 0xa389  72
Router   192.168.1.2      192.168.1.2      0x80000007   420  0x22 0xaf74  72
Router   192.168.1.3      192.168.1.3      0x80000005   417  0x22 0x9b26  60
Router   192.168.1.4      192.168.1.4      0x80000004   426  0x22 0x8037  60
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Router    192.168.1.5      192.168.1.5      0x80000004   434  0x22 0x4b61  60
Network   172.16.2.2       192.168.1.3      0x80000002   422  0x22 0xb381  32
Network   172.16.3.2       192.168.1.3      0x80000001   419  0x22 0xb87b  32
Network   172.16.4.2       192.168.1.4      0x80000001   436  0x22 0xa38e  32
Network   172.16.5.2       192.168.1.4      0x80000001   427  0x22 0xa689  32
Network   172.16.6.2       192.168.1.5      0x80000001   444  0x22 0x919c  32
Network   172.16.7.2       192.168.1.5      0x80000001   436  0x22 0x9497  32
```

After applying the configuration changes, we can see that these LSAs are no longer in de OSPF database:

```
jncie@vMX3> show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.1.1      192.168.1.1      0x80000020    17  0x22 0x7136 108
Router   192.168.1.2      192.168.1.2      0x8000000c    10  0x22 0xf0c1 108
Router   192.168.1.3      192.168.1.3      0x80000020     9  0x22 0x7b1   84
Router   192.168.1.4      192.168.1.4      0x80000009    20  0x22 0x784c  84
Router   192.168.1.5      192.168.1.5      0x8000000a    18  0x22 0x7147  84
```

Instead of checking the OSPF database, we can also verify the same thing by issuing the following command:

```
jncie@vMX3> show ospf interface xe-0/0/2.0 detail

Interface          State   Area            DR ID           BDR ID          Nbrs
xe-0/0/2.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 172.16.2.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

The output tells us that the link is of the type P2P (point-to-point). There is no DR and BDR and because of this, their value is listed as ‘0.0.0.0’. Notice also that this output includes the authentication type configured for this link.

The verification of BFD can be done as follows:

```
jncie@vMX3> show bfd session

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.2.2               Up        xe-0/0/2.0     8.000     2.000        4
172.16.4.2               Up        xe-0/0/3.0     8.000     2.000        4
172.16.6.2               Up        xe-0/0/4.0     8.000     2.000        4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX4> show bfd session

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.3.2               Up        xe-0/0/4.0     8.000     2.000        4
172.16.5.2               Up        xe-0/0/3.0     8.000     2.000        4
172.16.7.2               Up        xe-0/0/2.0     8.000     2.000        4
```

The verification of the load-balancing policy can be done by examining the forwarding table. We need to make sure we check a destination IP address that can be load-balanced. So, from any of the leaf switches, we can check the loopback IP of another switch:

```
jncie@vQFX3> show route forwarding-table destination 192.168.1.4

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.4/32     user     0                    ulst   131070     3
                              172.16.2.1         ucst     1708     6 xe-0/0/2.0
                              172.16.3.1         ucst     1709     6 xe-0/0/4.0

jncie@vQFX3> show route forwarding-table destination 192.168.1.5

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.1.5/32     user     0                    ulst   131070     3
                              172.16.2.1         ucst     1708     6 xe-0/0/2.0
                              172.16.3.1         ucst     1709     6 xe-0/0/4.0
```

To quickly check load balancing, consider the following:

```
jncie@vQFX4> show route forwarding-table | match ulst
192.168.1.3/32     user     0                    ulst   131070     3
192.168.1.5/32     user     0                    ulst   131070     3

jncie@vMX3> show route forwarding-table | match ulst

192.168.1.2/32     user     0                    ulst  1048574     2

jncie@vMX4> show route forwarding-table | match ulst

192.168.1.1/32     user     0                    ulst  1048575     2
```

Notice how the vMX routers will only be able to load-balance towards the loopback IP address of the other vMX. The spines can only use multiple links when they forward traffic to each other. When a spine forwards traffic to a leaf, it will use only 1 link.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.12: IBGP Configuration

- Configure IBGP for AS65000 based on the information provided in the logical topology.
- Ensure that:
    - vMX3 and vMX4 function as route-reflectors
    - the loss of a BGP neighbor is detected within 10 seconds
    - BGP sessions only support the exchange of EVPN routes
    - BGP sessions are authenticated using the MD5-hased password 'jncie'

Solution

Here, we configure the signaling for the overlay network. Normally, a full mesh of IBGP sessions is required. The reason why this is required is because of the way BGP handles loop detection. Whenever a route is advertised across an EBGP session, the BGP AS path is updated. The local AS is prepended as the route is advertised to the BGP neighbor. When a route is advertised across an IBGP session, the AS path is not updated. Since this makes it impossible for BGP to detect a routing loop, IBGP learned routes are never advertised across IBGP sessions.

To lift this restriction, we configure vMX3 and vMX4 as Route-Reflectors. RRs are allowed to advertise, or reflect, IBGP learned routes across IBGP sessions. RRs introduce a new BGP attribute, the CLUSTER_LIST. This is a sequence of CLUSTER_ID values that represent the reflection path that the route has passed. In other words, instead of detecting a loop by examining the AS path, the RR can detect a loop by examining the cluster list. On Juniper devices, you can configure a BGP speaker to be a RR by simply configuring the device with a 'cluster' value.

Since this task did not mention anything specifically, you can choose to use different or similar cluster values.

Since vMX3 and vMX4 are the RRs, they pear with every other device (including each other). The vQFX devices only set up a BGP session with each of the RRs.

Apart from the RR configuration, the rest of the overlay configuration is similar to the one discussed previously in part 1 of this chapter.

vMX3:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay cluster 192.168.1.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 4
set protocols bgp group overlay neighbor 192.168.1.2
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.4
set protocols bgp group overlay neighbor 192.168.1.5
```

vMX4:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay cluster 192.168.1.2
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 4
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.3
set protocols bgp group overlay neighbor 192.168.1.4
set protocols bgp group overlay neighbor 192.168.1.5
```

vQFX3:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.3
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 4
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.2
```

vQFX4:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 4
set protocols bgp group overlay neighbor 192.168.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay neighbor 192.168.1.2
```

vQFX5:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.5
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay authentication-key jncie
set protocols bgp group overlay bfd-liveness-detection minimum-interval 2000
set protocols bgp group overlay bfd-liveness-detection multiplier 4
set protocols bgp group overlay neighbor 192.168.1.1
set protocols bgp group overlay neighbor 192.168.1.2
```

Verification

Verifying the overlay is similar to what was discussed earlier during chapter 3 task 1. We issue the same command to verify that all of the BGP sessions are established:

```
jncie@vMX3> show bgp summary

Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.evpn.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.1.2           65000         12         11       0       0        4:21 Establ
  bgp.evpn.0: 0/0/0/0
192.168.1.3           65000          7          5       0       0        1:48 Establ
  bgp.evpn.0: 0/0/0/0
192.168.1.4           65000          7          5       0       0        1:40 Establ
  bgp.evpn.0: 0/0/0/0
192.168.1.5           65000          7          4       0       0        1:33 Establ
  bgp.evpn.0: 0/0/0/0
```

To check the options and address-family active for the BGP sessions, we issue the following command:

```
jncie@vQFX3> show bgp neighbor 192.168.1.1 | match "options| session"

  Options: <Preference LocalAddress AuthKey AddressFamily Rib-group Refresh>
  Options: <BfdEnabled>
  NLRI for this session: evpn
```

BGP is active only for EVPN. Since we have no VXLAN or EVPN related configuration applied to the devices, we currently have not active BGP routes:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show route protocol bgp

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

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

Task 3.13: Configure QFX Switch-Level EVPN

- Use ingress-replication, VXLAN encapsulation and Lo0 as the VTEP source interface.
- Configure the QFX to use route-target 1:1 for the overlay.
- VXLAN segments should carry a target community that differs from the fabric community.
- Prepare the overlay to allow the import of VXLAN segments that do not have explicitly defined target communities.

Solution

The QFX switch-level EVPN configuration was discussed before. The only new concept that is introduced here is the automatic assignment of target communities to VXLAN segments. This is possible when all routers are configured with the same autonomous system value. To complete this task, we configure the [**switch-options vrf-target**] hierarchy to import a target that is associated with the overlay network itself. In addition to this, we configure the 'auto' statement.

vQFX3:

```
set routing-options router-id 192.168.1.3

set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list all

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.3:65000
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1
```

vQFX4:

```
set routing-options router-id 192.168.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list all

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.4:65000
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1
```

vQFX5:

```
set routing-options router-id 192.168.1.5

set protocols evpn encapsulation vxlan
set protocols evpn multicast-mode ingress-replication
set protocols evpn extended-vni-list all

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 192.168.1.5:65000
set switch-options vrf-import fabric-import
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set policy-options policy-statement fabric-import term fabric from community fabric
set policy-options policy-statement fabric-import term fabric then accept

set policy-options community fabric members target:1:1
```

Verification

We start out verifying the QFX switch-level EVPN parameters by issuing the following command:

```
jncie@vQFX3> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.3:0
  Number of bridge domains: 0
  Number of neighbors: 0

Instance: default-switch
  Route Distinguisher: 192.168.1.3:65000
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  MAC advertisements:                          0        0
  MAC+IP advertisements:                       0        0
  Default gateway MAC advertisements:          0        0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 0
Number of neighbors: 0
Number of ethernet segments: 0
Router-ID: 192.168.1.3
Source VTEP interface IP: 192.168.1.3
```

Through the configuration of 'vrf-target auto', we enabled automatic route-targets. Since automatically generated route-targets are based on the autonomous system value, every device that participates in the overlay will need to have the same autonomous system configured.

Even though no VXLAN segments have been defined yet, we can already see that a policy to import automatically generated targets has been created by the system. In addition to this, we can see that the policy we have defined is also referenced inside that automatically generated policy:

```
jncie@vQFX3> show policy ?

Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-autoderive-default-switch-internal__
  __vrf-export-default-switch-internal__
  __vrf-import-autoderive-default-switch-internal__
  __vrf-import-default-switch-internal__
  conditions           Show conditions used by policy
  damping              Show state of route flap damping
  fabric-import
  lbpp
  |                    Pipe through a command
{master:0}
```

The following policy was configured:

```
jncie@vQFX3> show policy fabric-import

Policy fabric-import:
    Term fabric:
        from community fabric [target:1:1 ]
        then accept
```

Using the following command, we can see that the 'fabric-import' policy is referenced by the automatically generated (and applied) policy in the first term:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show policy __vrf-import-autoderive-default-switch-internal__

Policy __vrf-import-autoderive-default-switch-internal__:
    Term unnamed:
        from policy fabric-import
        then accept
    Term unnamed:
        then reject
```

The final term, or term 2, is a 'then reject'. After completing subsequent tasks, we will see that for every VXLAN segment, another term is added before this reject term.

Task 3.14: VXLAN Segment Configuration and Active-Active Multihoming

- Enable layer 2 communication for devices in VLAN 100 using VNI 100.
- Do not explicitly configure the VXLAN segments with a target community.
- Enable active-active multihoming on vQFX4 and vQFX5. Assume that the connected device does not support LACP.

Solution

In the previous task, though not strictly necessary, we configured the [protocols evpn extended-vni-list] with the 'all' statement. This way, when we define a new VXLAN, we only need to configure a VLAN with the 'vxlan' statement. If we would not have used the 'all' statement, we would have to add a VXLAN number every time we configured an additional VXLAN segment.

vQFX3:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 100
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode trunk vlan members 100
set vlans vlan-100 vlan-id 100 vxlan vni 100
```

vQFX4:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk vlan members 100

set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/3 ether-options 802.3ad ae0

set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp system-id 00:11:11:11:11:11
set interfaces ae0 aggregated-ether-options lacp force-up
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set vlans vlan-100 vlan-id 100 vxlan vni 100
```

vQFX5:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp system-id 00:11:11:11:11:11
set interfaces ae0 aggregated-ether-options lacp force-up
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set vlans vlan-100 vlan-id 100 vxlan vni 100
```

Verification

We start out verifying the multihomed segment:

```
jncie@vQFX4> show lacp statistics interfaces

Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/3                   0           0            0            0

{master:0}
jncie@vQFX4> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3 FUP    Actor   No   Yes   Yes  Yes  Yes   Yes     Fast   Passive
      xe-0/0/3 FUP  Partner   No   Yes    No   No  Yes   Yes     Fast   Passive
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3 FUP              Current     No periodic Collecting distributing
```

No LACP packets are exchanged and the interface state is Fore UP. This knob can be useful when you need to connect device that do not support the LACP protocol.

We continue our verification by checking the EVPN instance on the vQFX:

```
jncie@vQFX4> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.4:0
  Number of bridge domains: 0
  Number of neighbors: 1
    192.168.1.5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Received routes
        Ethernet Segment:                       1

Instance: default-switch
  Route Distinguisher: 192.168.1.4:65000
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         1       3
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                               Mode             Status
    ae0.0           00:11:11:11:11:11:11:11:11:11     all-active       Up
    xe-0/0/1.0      00:00:00:00:00:00:00:00:00:00     single-homed     Up
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 1
  VLAN  VNI    Intfs / up    IRB intf   Mode             MAC sync  IM route label
  100   100         2   2                Extended         Enabled   100
Number of neighbors: 2
  192.168.1.3
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
  192.168.1.5
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved by IFL ae0.0
    Local interface: ae0.0, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.5      100        0               all-active
    Designated forwarder: 192.168.1.4
    Backup forwarder: 192.168.1.5
  Router-ID: 192.168.1.4
  Source VTEP interface IP: 192.168.1.4
```

We can see that both vQFX3 as well as vQFX5 are seen as neighbors for the EVPN. The printout tells us that vQFX4 has a local interface actively participating in the multihomed segment. vQFX4 has also learned about vQFX5's involvement in this ES.

To check the automatic assignment of route-targets, we can check the hidden policies:

```
jncie@vQFX3> show policy ?
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-autoderive-default-switch-internal__
  __vrf-export-default-switch-internal__
  __vrf-import-autoderive-default-switch-internal__
  __vrf-import-default-switch-internal__
  conditions           Show conditions used by policy
  damping              Show state of route flap damping
  fabric-import
  lbpp
  |                    Pipe through a command
```

```
jncie@vQFX3> show policy __vrf-export-autoderive-default-switch-internal__

    Term unnamed:
        then community + __vrf-community-default-switch-100-internal__ [target:65000:268435556
] accept
```

We can see that a new export policy was updated. The switch will export routes with target:65000:268435556. The import policy was also updated:

```
jncie@vQFX3> show policy __vrf-import-autoderive-default-switch-internal__

Policy __vrf-import-autoderive-default-switch-internal__:
    Term 100:
        from community __vrf-community-default-switch-100-internal__ [target:65000:268435556 ]
        then accept
    Term unnamed:
        from policy fabric-import
        then accept
    Term unnamed:
        then reject
```

The newly added term inside the import policy will ensure that that the switch imports the VXLAN routing information that use the same target community.

When we look at the advertised route using the 'detail' keyword, we can see what target community is attached to the route:

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.1 detail

default-switch.evpn.0: 11 destinations, 19 routes (11 active, 0 holddown, 0 hidden)
* 2:192.168.1.3:65000::100::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 100
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan

* 2:192.168.1.3:65000::100::00:0c:29:8c:00:b2/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 100
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan

* 3:192.168.1.3:65000::100::192.168.1.3/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 100
     PMSI: Flags 0x0: Label 0x64: Type INGRESS-REPLICATION 192.168.1.3
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan
     PMSI: Flags 0x0: Label 0x64: Type INGRESS-REPLICATION 192.168.1.3
```

Before verifying the switching table and EVPN database, we first generate some traffic:

```
jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.100.2

PING 10.200.100.2 (10.200.100.2): 56 data bytes
!!!!!
--- 10.200.100.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 19.198/21.851/28.937/3.588 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.100.3

PING 10.200.100.3 (10.200.100.3): 56 data bytes
!!!!!
--- 10.200.100.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 29.712/31.410/36.364/2.507 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.100.4

PING 10.200.100.4 (10.200.100.4): 56 data bytes
!!!!!
--- 10.200.100.4 ping statistics ---
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 28.816/38.519/56.649/11.028 ms
```

Communication within the VXLAN segment is possible. Now we have some entries in the switching and EVPN table to examine:

```
jncie@vQFX3> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-100             00:05:86:71:7b:c0   DR       esi.1728
00:11:11:11:11:11:11:11:11:11
   vlan-100             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-100             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-100             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4

jncie@vQFX4> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-100             00:05:86:71:7b:c0   DR       ae0.0
   vlan-100             00:0c:29:8c:00:a8   D        vtep.32769               192.168.1.3
   vlan-100             00:0c:29:8c:00:b2   D        vtep.32769               192.168.1.3
   vlan-100             00:0c:29:8c:00:bc   D        xe-0/0/1.0

jncie@vQFX5> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-100             00:05:86:71:7b:c0   DL       ae0.0
   vlan-100             00:0c:29:8c:00:a8   D        vtep.32770               192.168.1.3
   vlan-100             00:0c:29:8c:00:b2   D        vtep.32770               192.168.1.3
   vlan-100             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4

jncie@vQFX3> show evpn database
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Instance: default-switch
VLAN  VNI  MAC address        Active source                    Timestamp        IP address
      100   00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11   Mar 25 08:29:15
      100   00:0c:29:8c:00:a8  xe-0/0/1.0                       Mar 25 08:26:25
      100   00:0c:29:8c:00:b2  xe-0/0/3.0                       Mar 25 08:26:18
      100   00:0c:29:8c:00:bc  192.168.1.4                      Mar 25 08:26:53

jncie@vQFX4> show evpn database

Instance: default-switch
VLAN  VNI  MAC address        Active source                    Timestamp        IP address
      100   00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11   Mar 25 08:29:15
      100   00:0c:29:8c:00:a8  192.168.1.3                      Mar 25 08:26:39
      100   00:0c:29:8c:00:b2  192.168.1.3                      Mar 25 08:26:39
      100   00:0c:29:8c:00:bc  xe-0/0/1.0                       Mar 25 08:26:52

jncie@vQFX5> show evpn database

Instance: default-switch
VLAN  VNI  MAC address        Active source                    Timestamp        IP address
      100   00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11   Mar 25 08:29:14
      100   00:0c:29:8c:00:a8  192.168.1.3                      Mar 25 08:27:06
      100   00:0c:29:8c:00:b2  192.168.1.3                      Mar 25 08:27:06
      100   00:0c:29:8c:00:bc  192.168.1.4                      Mar 25 08:27:06
```

Task 3.15: VXLAN Segment Configuration for VNI 101

- Enable layer 2 communication for devices located in VLAN 101.

Solution

Through the use of the automatically generated target community and the '**protocols evpn extended-vni-list all**' statement, adding new VXLANs to the controllerless overlay becomes a menial task. The only required configuration is the creation of the VLAN/VXLAN and assigning it to an interface.

vQFX3:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 101
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members 101

set vlans vlan-101 vlan-id 101 vxlan vni 101
```

vQFX4:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 101
set interfaces ae0 unit 0 family ethernet-switching vlan members 101
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set vlans vlan-101 vlan-id 101 vxlan vni 101
```

vQFX5:

```
set interfaces ae0 unit 0 family ethernet-switching vlan members 101

set vlans vlan-101 vlan-id 101 vxlan vni 101
```

Verification

Verifying connectivity within the VXLAN segment:

```
jncie@vMX6> ping rapid routing-instance DC1-S1-VL101 10.200.101.2

PING 10.200.101.2 (10.200.101.2): 56 data bytes
!!!!!
--- 10.200.101.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 24.727/27.013/31.790/2.563 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL101 10.200.101.3

PING 10.200.101.3 (10.200.101.3): 56 data bytes
!!!!!
--- 10.200.101.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 38.698/51.163/90.383/19.714 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL101 10.200.101.4

PING 10.200.101.4 (10.200.101.4): 56 data bytes
!!!!!
--- 10.200.101.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 34.280/75.531/233.139/78.819 ms
```

Verifying the switching table:

```
jncie@vQFX3> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 8 entries, 8 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-100             00:05:86:71:7b:c0   DR       esi.1728
00:11:11:11:11:11:11:11:11:11
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
   vlan-100            00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-100            00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-100            00:0c:29:8c:00:bc   D        vtep.32769          192.168.1.4
   vlan-101            00:05:86:71:7b:c0   DR       esi.1728
00:11:11:11:11:11:11:11:11:11
   vlan-101            00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-101            00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-101            00:0c:29:8c:00:bc   D        vtep.32769          192.168.1.4

jncie@vQFX4> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
```

```
Ethernet switching table : 8 entries, 8 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                  Active
   name                address             flags    interface                source
   vlan-100            00:05:86:71:7b:c0   DR       ae0.0
   vlan-100            00:0c:29:8c:00:a8   D        vtep.32769          192.168.1.3
   vlan-100            00:0c:29:8c:00:b2   D        vtep.32769          192.168.1.3
   vlan-100            00:0c:29:8c:00:bc   D        xe-0/0/1.0
   vlan-101            00:05:86:71:7b:c0   DL       ae0.0
   vlan-101            00:0c:29:8c:00:a8   D        vtep.32769          192.168.1.3
   vlan-101            00:0c:29:8c:00:b2   D        vtep.32769          192.168.1.3
   vlan-101            00:0c:29:8c:00:bc   D        xe-0/0/1.0
```

When we look at the export and import policy, we can see that an additional term has been added for the newly created VXLAN segment:

```
jncie@vQFX3> show policy __vrf-import-autoderive-default-switch-internal__

Policy __vrf-import-autoderive-default-switch-internal__:
    Term 100:
        from community __vrf-community-default-switch-100-internal__ [target:65000:268435556 ]
        then accept
    Term 101:
        from community __vrf-community-default-switch-101-internal__ [target:65000:268435557 ]
        then accept
    Term unnamed:
        from policy fabric-import
        then accept
    Term unnamed:
        then reject

jncie@vQFX3> show policy __vrf-export-autoderive-default-switch-internal__

    Term unnamed:
        then community + __vrf-community-default-switch-100-internal__ [target:65000:268435556
] accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Term unnamed:
        then community + __vrf-community-default-switch-101-internal__ [target:65000:268435557
] accept
```

We can see that the auto-target is also applied to the exported routes:

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.1 detail

default-switch.evpn.0: 18 destinations, 30 routes (18 active, 0 holddown, 0 hidden)
* 2:192.168.1.3:65000::100::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 100
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan

...

* 2:192.168.1.3:65000::101::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 101
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435557 encapsulation0:0:0:0:vxlan
```

Task 3.16: Additional VXLAN Segment Configuration

- Enable layer 2 communication for devices located in the following VLANs:

VXLAN / VNI	Route Target
102	
103	
104	
105	target:65000:105
106	target:65000:106
107	target:65000:107
108	target:65000:108
109	target:65000:109
110	target:65000:110

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Only configure route-targets for VXLAN segments when they are mentioned.

Solution

Using automatically assigned target communities can ease both configuration as well as automation. In certain scenarios it may still be desirable to override the automatically generated target communities. In very large deployments, you may want to limit the routing-table and forwarding table entries of some of devices that participate in the overlay network. In other cases, you may want to attribute different qualities to the virtual-switch instance on a vMX and only import the segments that should enjoy those different qualities.

You can override the automatically assigned route-target by explicitly configuring one for the VXLAN segment.

On every vQFX:

First, we configure the missing VLANs and VXLANs:

```
set vlans vlan-102 vlan-id 102 vxlan vni 102
set vlans vlan-103 vlan-id 103 vxlan vni 103
set vlans vlan-104 vlan-id 104 vxlan vni 104
set vlans vlan-105 vlan-id 105 vxlan vni 105
set vlans vlan-106 vlan-id 106 vxlan vni 106
set vlans vlan-107 vlan-id 107 vxlan vni 107
set vlans vlan-108 vlan-id 108 vxlan vni 108
set vlans vlan-109 vlan-id 109 vxlan vni 109
set vlans vlan-110 vlan-id 110 vxlan vni 110
```

Then we configure the target-community for the VXLAN segments under the [**protocols evpn**] configuration stanza:

```
set protocols evpn vni-options vni 105 vrf-target export target:65000:105
set protocols evpn vni-options vni 106 vrf-target export target:65000:106
set protocols evpn vni-options vni 107 vrf-target export target:65000:107
set protocols evpn vni-options vni 108 vrf-target export target:65000:108
set protocols evpn vni-options vni 109 vrf-target export target:65000:109
set protocols evpn vni-options vni 110 vrf-target export target:65000:110

set protocols evpn extended-vni-list 105
set protocols evpn extended-vni-list 106
set protocols evpn extended-vni-list 107
set protocols evpn extended-vni-list 108
set protocols evpn extended-vni-list 109
set protocols evpn extended-vni-list 110
```

The import policy that is active on the vQFX devices only import the target for the overlay (target:1:1) and the automatically assigned target-communities. We can add a term for

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

every VXLAN to the already existing policy ('fabric-import'), or we can create an additional policy.

In this example, we will create, and apply an import policy that specifically handles VXLAN segments:

```
set switch-options vrf-import vxlan-import

set policy-options policy-statement vxlan-import term target-105 from community target-105
set policy-options policy-statement vxlan-import term target-105 then accept
set policy-options policy-statement vxlan-import term target-106 from community target-106
set policy-options policy-statement vxlan-import term target-106 then accept
set policy-options policy-statement vxlan-import term target-107 from community target-107
set policy-options policy-statement vxlan-import term target-107 then accept
set policy-options policy-statement vxlan-import term target-108 from community target-108
set policy-options policy-statement vxlan-import term target-108 then accept
set policy-options policy-statement vxlan-import term target-109 from community target-109
set policy-options policy-statement vxlan-import term target-109 then accept
set policy-options policy-statement vxlan-import term target-110 from community target-110
set policy-options policy-statement vxlan-import term target-110 then accept

set policy-options community target-105 members target:65000:105
set policy-options community target-106 members target:65000:106
set policy-options community target-107 members target:65000:107
set policy-options community target-108 members target:65000:108
set policy-options community target-109 members target:65000:109
set policy-options community target-110 members target:65000:110
```

Lastly, we assign the new VLANs to all of the trunk interfaces:

vQFX3:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 102-110
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members 102-110
```

vQFX4:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 102-110
set interfaces ae0 unit 0 family ethernet-switching vlan members 102-110
```

vQFX5:

```
set interfaces ae0 unit 0 family ethernet-switching vlan members 102-110
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

First, we check the policies on the switch. We can see that several additional policies have been created:

```
jncie@vQFX3> show policy ?

Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-autoderive-default-switch-internal__
  __vrf-export-default-bd-105-internal__
  __vrf-export-default-bd-106-internal__
  __vrf-export-default-bd-107-internal__
  __vrf-export-default-bd-108-internal__
  __vrf-export-default-bd-109-internal__
  __vrf-export-default-bd-110-internal__
  __vrf-export-default-switch-internal__
  __vrf-import-autoderive-default-switch-internal__
  __vrf-import-default-switch-internal__
  conditions           Show conditions used by policy
  damping              Show state of route flap damping
  fabric-import
  lbpp
  vxlan-import
  |                    Pipe through a command
```

We start by checking the import policy on the vQFX:

```
jncie@vQFX3> show policy __vrf-import-autoderive-default-switch-internal__

Policy __vrf-import-autoderive-default-switch-internal__:
    Term 100:
        from community __vrf-community-default-switch-100-internal__ [target:65000:268435556 ]
        then accept
    Term 101:
        from community __vrf-community-default-switch-101-internal__ [target:65000:268435557 ]
        then accept
    Term 102:
        from community __vrf-community-default-switch-102-internal__ [target:65000:268435558 ]
        then accept
    Term 103:
        from community __vrf-community-default-switch-103-internal__ [target:65000:268435559 ]
        then accept
    Term 104:
        from community __vrf-community-default-switch-104-internal__ [target:65000:268435560 ]
        then accept
    Term unnamed:
        from policy [ fabric-import vxlan-import ]
        then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Term unnamed:
        then reject
```

We see that the import policy now has additional terms for the newly created VXLAN segments. In addition to this, the 'unnamed' term was updated. We can see that the newly created 'vxlan-import' policy was added to this term.

After this, we check to see how the export policy was updated:

```
jncie@vQFX3> show policy __vrf-export-autoderive-default-switch-internal__

    Term unnamed:
        then community + __vrf-community-default-switch-100-internal__ [target:65000:268435556
] accept
    Term unnamed:
        then community + __vrf-community-default-switch-101-internal__ [target:65000:268435557
] accept
    Term unnamed:
        then community + __vrf-community-default-switch-102-internal__ [target:65000:268435558
] accept
    Term unnamed:
        then community + __vrf-community-default-switch-103-internal__ [target:65000:268435559
] accept
    Term unnamed:
        then community + __vrf-community-default-switch-104-internal__ [target:65000:268435560
] accept
```

As expected, the policies were updated in the same way.

When we examined the policies on the switch, we also noticed a new set of policies. These were called '__vrf-export-default-bd-105-internal__' up until '__vrf-export-default-bd-110-internal__'. These policies are the result of the configuration of VNI 105 to 110:

```
jncie@vQFX3> show policy __vrf-export-default-bd-105-internal__

Policy __vrf-export-default-bd-105-internal__:
    Term unnamed:
        then community + __vrf-community-default-bd-105-export-internal__ [target:65000:105 ]
accept
```

To verify that we are actually attaching these target communities, we can check the advertised routes on one of the vQFX switches:

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.1 detail

default-switch.evpn.0: 81 destinations, 129 routes (81 active, 0 holddown, 0 hidden)
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 2:192.168.1.3:65000::103::00:0c:29:8c:00:b2/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 103
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435559 encapsulation0:0:0:0:vxlan

* 2:192.168.1.3:65000::104::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 104
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:268435560 encapsulation0:0:0:0:vxlan

...

* 2:192.168.1.3:65000::105::00:0c:29:8c:00:b2/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 105
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:105 encapsulation0:0:0:0:vxlan

...

* 3:192.168.1.3:65000::109::192.168.1.3/304 (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.3:65000
     Route Label: 109
     PMSI: Flags 0x0: Label 0x6d: Type INGRESS-REPLICATION 192.168.1.3
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:109 encapsulation0:0:0:0:vxlan
     PMSI: Flags 0x0: Label 0x6d: Type INGRESS-REPLICATION 192.168.1.3
...
```

Consider the use of "| match" to present the same information in a more concise way:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show route advertising-protocol bgp 192.168.1.1 detail | match
"65000::|Communities"

* 2:192.168.1.3:65000::100::00:0c:29:8c:00:a8/304 (1 entry, 1 announced)
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan
* 2:192.168.1.3:65000::100::00:0c:29:8c:00:b2/304 (1 entry, 1 announced)
     Communities: target:65000:268435556 encapsulation0:0:0:0:vxlan
...
     Communities: target:65000:108 encapsulation0:0:0:0:vxlan
* 3:192.168.1.3:65000::109::192.168.1.3/304 (1 entry, 1 announced)
     Communities: target:65000:109 encapsulation0:0:0:0:vxlan
* 3:192.168.1.3:65000::110::192.168.1.3/304 (1 entry, 1 announced)
     Communities: target:65000:110 encapsulation0:0:0:0:vxlan
```

The rest of the verification that can be done in this task is similar to what was discussed before.

We can issue several pings on vMX6 to see if we can pass traffic through the VXLAN segment:

```
jncie@vMX6> ping rapid routing-instance DC1-S1-VL102 10.200.102.4

PING 10.200.102.4 (10.200.102.4): 56 data bytes
!!!!!
--- 10.200.102.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 34.398/55.777/131.103/37.724 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL106 10.200.106.3

PING 10.200.106.3 (10.200.106.3): 56 data bytes
!!!!!
--- 10.200.106.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 32.975/58.918/127.279/34.528 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL110 10.200.110.2

PING 10.200.110.2 (10.200.110.2): 56 data bytes
!!!!!
--- 10.200.110.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 11.947/25.194/54.135/15.031 ms
```

We could consider issuing ‘show evpn instance extensive’ and the ‘show ethernet-switching table’ commands:

```
jncie@vQFX3> show evpn instance extensive

Instance: __default_evpn__
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Route Distinguisher: 192.168.1.3:0
  Number of bridge domains: 0
  Number of neighbors: 0

Instance: default-switch
  Route Distinguisher: 192.168.1.3:65000
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                         22      23
    MAC+IP advertisements:                       0       0
    Default gateway MAC advertisements:          0       0
  Number of local interfaces: 2 (2 up)
    Interface name  ESI                            Mode             Status
    xe-0/0/1.0      00:00:00:00:00:00:00:00:00:00  single-homed     Up
    xe-0/0/3.0      00:00:00:00:00:00:00:00:00:00  single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 11
    VLAN  VNI    Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100        2   2                Extended        Enabled   100
    101   101        2   2                Extended        Enabled   101
    102   102        2   2                Extended        Enabled   102
    103   103        2   2                Extended        Enabled   103
    104   104        2   2                Extended        Enabled   104
    105   105        2   2                Extended        Enabled   105
    106   106        2   2                Extended        Enabled   106
    107   107        2   2                Extended        Enabled   107
    108   108        2   2                Extended        Enabled   108
    109   109        2   2                Extended        Enabled   109
    110   110        2   2                Extended        Enabled   110
  Number of neighbors: 2
    192.168.1.4
      Received routes
        MAC address advertisement:             22
        MAC+IP address advertisement:           0
        Inclusive multicast:                   11
        Ethernet auto-discovery:                2
    192.168.1.5
      Received routes
        MAC address advertisement:              1
        MAC+IP address advertisement:           0
        Inclusive multicast:                   11
        Ethernet auto-discovery:                2
  Number of ethernet segments: 1
    ESI: 00:11:11:11:11:11:11:11:11:11
      Status: Resolved
      Number of remote PEs connected: 2
        Remote PE        MAC label  Aliasing label  Mode
        192.168.1.5      108        0               all-active
        192.168.1.4      102        0               all-active
  Router-ID: 192.168.1.3
  Source VTEP interface IP: 192.168.1.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 44 entries, 44 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-100             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-100             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-100             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-100             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-101             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-101             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-101             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-101             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-102             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-102             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-102             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-102             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-103             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-103             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-103             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-103             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-104             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-104             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-104             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-104             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-105             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-105             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-105             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-105             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-106             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-106             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-106             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-106             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
   vlan-107             00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-107             00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-107             00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-107             00:0c:29:8c:00:bc   D        vtep.32769               192.168.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
   vlan-108            00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-108            00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-108            00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-108            00:0c:29:8c:00:bc   D        vtep.32769          192.168.1.4
   vlan-109            00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-109            00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-109            00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-109            00:0c:29:8c:00:bc   D        vtep.32769          192.168.1.4
   vlan-110            00:05:86:71:74:c0   DR       esi.1724
00:11:11:11:11:11:11:11:11:11
   vlan-110            00:0c:29:8c:00:a8   D        xe-0/0/1.0
   vlan-110            00:0c:29:8c:00:b2   D        xe-0/0/3.0
   vlan-110            00:0c:29:8c:00:bc   D        vtep.32769          192.168.1.4
```

Task 3.17: MX Virtual-Switch with Selective VXAN Import

- Import the VXLAN segments 100, 101, 102, 103 and 104 into a routing-instance called 'fabric'.
- Make sure that no other address information is placed into the newly created instance.

Solution

We need to import the VXLAN segments that the leaf switches have advertised using the automatically generated target community. To this end, we configure a virtual-switch with a configuration that is similar to the one we applied on the QFX.

vMX3:

```
set routing-options route-distinguisher-id 192.168.1.1
set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-target target:1:1
set routing-instances fabric vrf-target auto
set routing-instances fabric protocols evpn encapsulation vxlan
set routing-instances fabric protocols evpn extended-vni-list all
set routing-instances fabric protocols evpn multicast-mode ingress-replication
set routing-instances fabric bridge-domains bd_100 vlan-id 100 vxlan vni 100
set routing-instances fabric bridge-domains bd_101 vlan-id 101 vxlan vni 101
set routing-instances fabric bridge-domains bd_102 vlan-id 102 vxlan vni 102
set routing-instances fabric bridge-domains bd_103 vlan-id 103 vxlan vni 103
set routing-instances fabric bridge-domains bd_104 vlan-id 104 vxlan vni 104
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set routing-options route-distinguisher-id 192.168.1.2

set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-target target:1:1
set routing-instances fabric vrf-target auto
set routing-instances fabric protocols evpn encapsulation vxlan
set routing-instances fabric protocols evpn extended-vni-list all
set routing-instances fabric protocols evpn multicast-mode ingress-replication
set routing-instances fabric bridge-domains bd_100 vlan-id 100 vxlan vni 100
set routing-instances fabric bridge-domains bd_101 vlan-id 101 vxlan vni 101
set routing-instances fabric bridge-domains bd_102 vlan-id 102 vxlan vni 102
set routing-instances fabric bridge-domains bd_103 vlan-id 103 vxlan vni 103
set routing-instances fabric bridge-domains bd_104 vlan-id 104 vxlan vni 104
```

Verification

First, we check the EVPN extensive information on both vMX routers:

```
jncie@vMX3> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 192.168.1.1:8
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         0      20
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 5
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0                Extended        Enabled
    101   101            0    0                Extended        Enabled
    102   102            0    0                Extended        Enabled
    103   103            0    0                Extended        Enabled
    104   104            0    0                Extended        Enabled
  Number of neighbors: 3
    192.168.1.3
      Received routes
        MAC address advertisement:              10
        MAC+IP address advertisement:            0
        Inclusive multicast:                     5
        Ethernet auto-discovery:                 0
    192.168.1.4
      Received routes
        MAC address advertisement:               9
        MAC+IP address advertisement:            0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
  192.168.1.5
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.5      100        0               all-active
      192.168.1.4      103        0               all-active
  Router-ID: 192.168.1.1
  Source VTEP interface IP: 192.168.1.1

jncie@vMX4> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 192.168.1.2:8
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         0      20
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 5
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0                Extended        Enabled
    101   101            0    0                Extended        Enabled
    102   102            0    0                Extended        Enabled
    103   103            0    0                Extended        Enabled
    104   104            0    0                Extended        Enabled
Number of neighbors: 3
  192.168.1.3
    Received routes
      MAC address advertisement:             10
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                0
  192.168.1.4
    Received routes
      MAC address advertisement:              9
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
  192.168.1.5
    Received routes
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.5      100        0               all-active
      192.168.1.4      103        0               all-active
Router-ID: 192.168.1.2
Source VTEP interface IP: 192.168.1.2
```

Both routers have discovered the vQFX switches and both vMX are receiving different types of routes from those leaf switches.

We then check the import and export policies that are active on the vMX:

```
jncie@vMX3> show policy ?

Possible completions:
  <[Enter]>            Execute this command
  <policy>             Name of policy
  __vrf-export-fabric-internal__
  __vrf-import-autoderive-fabric-internal__
  __vrf-import-fabric-internal__
  conditions           Show conditions used by policy
  damping              Show state of route flap damping
  fabric               Internal fabric state
  lbpp
  logical-system       Name of logical system, or 'all'
  statistics           Statistics for policy
  |                    Pipe through a command
```

Same as on vQFX, the vMX generated and applied several policies. We can check the content of the policy by issuing the following command:

```
jncie@vMX3> show policy __vrf-import-autoderive-fabric-internal__

Policy __vrf-import-autoderive-fabric-internal__:
    Term 100:
        from community __vrf-community-fabric-100-internal__ [target:65000:268435556 ]
        then accept
    Term 101:
        from community __vrf-community-fabric-101-internal__ [target:65000:268435557 ]
        then accept
    Term 102:
        from community __vrf-community-fabric-102-internal__ [target:65000:268435558 ]
        then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Term 103:
    from community __vrf-community-fabric-103-internal__ [target:65000:268435559 ]
    then accept
Term 104:
    from community __vrf-community-fabric-104-internal__ [target:65000:268435560 ]
    then accept
Term unnamed:
    from community __vrf-community-fabric-common-internal__ [target:1:1 ]
    then accept
Term unnamed:
    then reject
```

We see that the import policy imports VNIs 100-104. In addition to this, target:1:1 (called 'vrf-community-fabric-common-internal') is also imported. Notice the last term, called 'unnamed'. This term rejects all other addresses from being imported into the routing-table.

Check the address information that was placed inside the routing-instance:

```
jncie@vMX3> show route table fabric.evpn.0

fabric.evpn.0: 39 destinations, 78 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:192.168.1.4:0::11111111111111111::FFFF:FFFF/304 AD/ESI
                   *[BGP/170] 00:09:32, localpref 100, from 192.168.1.4
                      AS path: I, validation-state: unverified
                     > to 172.16.4.2 via xe-0/0/3.0
                     [BGP/170] 00:09:32, localpref 100, from 192.168.1.2
                      AS path: I, validation-state: unverified
                     > to 172.16.4.2 via xe-0/0/3.0
1:192.168.1.4:65000::11111111111111111::0/304 AD/EVI
                   *[BGP/170] 00:09:32, localpref 100, from 192.168.1.4
                      AS path: I, validation-state: unverified
                     > to 172.16.4.2 via xe-0/0/3.0
                     [BGP/170] 00:09:32, localpref 100, from 192.168.1.2
                      AS path: I, validation-state: unverified
                     > to 172.16.4.2 via xe-0/0/3.0
…
3:192.168.1.5:65000::104::192.168.1.5/304 IM
                   *[BGP/170] 00:09:32, localpref 100, from 192.168.1.5
                      AS path: I, validation-state: unverified
                     > to 172.16.6.2 via xe-0/0/4.0
                     [BGP/170] 00:09:32, localpref 100, from 192.168.1.2
                      AS path: I, validation-state: unverified
                     > to 172.16.6.2 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Check the EVPN instance on the vMX:

```
jncie@vMX3> show evpn database instance fabric

Instance: fabric
VLAN  DomainId  MAC address        Active source                       Timestamp        IP address
     100        00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11       Mar 26 11:43:43
     100        00:0c:29:8c:00:a8  192.168.1.3                         Mar 26 11:43:43
     100        00:0c:29:8c:00:b2  192.168.1.3                         Mar 26 11:43:43
     100        00:0c:29:8c:00:bc  192.168.1.4                         Mar 26 11:43:43
     101        00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11       Mar 26 11:43:43
     101        00:0c:29:8c:00:a8  192.168.1.3                         Mar 26 11:43:43
     101        00:0c:29:8c:00:b2  192.168.1.3                         Mar 26 11:43:43
     101        00:0c:29:8c:00:bc  192.168.1.4                         Mar 26 11:43:43
     102        00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11       Mar 26 11:43:43
     102        00:0c:29:8c:00:a8  192.168.1.3                         Mar 26 11:43:43
     102        00:0c:29:8c:00:b2  192.168.1.3                         Mar 26 11:43:43
     102        00:0c:29:8c:00:bc  192.168.1.4                         Mar 26 11:43:43
     103        00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11       Mar 26 11:43:43
     103        00:0c:29:8c:00:a8  192.168.1.3                         Mar 26 11:43:43
     103        00:0c:29:8c:00:b2  192.168.1.3                         Mar 26 11:43:43
     103        00:0c:29:8c:00:bc  192.168.1.4                         Mar 26 11:43:43
     104        00:05:86:71:7b:c0  00:11:11:11:11:11:11:11:11:11       Mar 26 11:43:43
     104        00:0c:29:8c:00:a8  192.168.1.3                         Mar 26 11:43:43
     104        00:0c:29:8c:00:b2  192.168.1.3                         Mar 26 11:43:43
     104        00:0c:29:8c:00:bc  192.168.1.4                         Mar 26 11:43:43
```

Check the bridging table on the vMX:

```
jncie@vMX3> show bridge mac-table

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : fabric
 Bridging domain : bd_100, VLAN : 100
   MAC                 MAC      Logical                 Active
   address             flags    interface               source
   00:05:86:71:7b:c0   DR       esi.614                 00:11:11:11:11:11:11:11:11:11
   00:0c:29:8c:00:a8   D        vtep.32769              192.168.1.3
   00:0c:29:8c:00:b2   D        vtep.32769              192.168.1.3
   00:0c:29:8c:00:bc   D        vtep.32770              192.168.1.4

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

...

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Routing instance : fabric
 Bridging domain : bd_104, VLAN : 104
   MAC                 MAC      Logical                 Active
   address             flags    interface               source
   00:05:86:71:7b:c0   DR       esi.614                 00:11:11:11:11:11:11:11:11:11
   00:0c:29:8c:00:a8   D        vtep.32769              192.168.1.3
   00:0c:29:8c:00:b2   D        vtep.32769              192.168.1.3
   00:0c:29:8c:00:bc   D        vtep.32770              192.168.1.4
```

To check what MAC addresses were learned from remote devices, we can also issue the following command:

```
jncie@vMX4> show l2-learning vxlan-tunnel-end-point remote mac-table

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Logical system   : <default>
Routing instance : fabric
 Bridging domain : bd_100+100, VLAN : 100, VNID : 100
   MAC                 MAC      Logical           Remote VTEP
   address             flags    interface         IP address
   00:05:86:71:7b:c0   DR       esi.614           192.168.1.5
   00:0c:29:8c:00:a8   D        vtep.32769        192.168.1.3
   00:0c:29:8c:00:b2   D        vtep.32769        192.168.1.3
   00:0c:29:8c:00:bc   D        vtep.32770        192.168.1.4

...

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

 Bridging domain : bd_104+104, VLAN : 104, VNID : 104
   MAC                 MAC      Logical           Remote VTEP
   address             flags    interface         IP address
   00:05:86:71:7b:c0   DR       esi.614           192.168.1.4
   00:0c:29:8c:00:a8   D        vtep.32769        192.168.1.3
   00:0c:29:8c:00:b2   D        vtep.32769        192.168.1.3
   00:0c:29:8c:00:bc   D        vtep.32770        192.168.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.18: Layer 3 VXLAN Gateway Configuration

- Configure the following layer 3 gateways on vMX3 and vMX4:

VXLAN / VNI	Gateway
100	10.200.100.254/24
101	10.200.101.254/24
102	10.200.102.254/24
103	10.200.103.254/24
104	10.200.104.254/24

- Configure vMX3 and vMX4 with the IP addresses that precede the virtual-gateway-address.

Solution

vMX3:

```
set interfaces irb unit 100 family inet address 10.200.100.252/24 virtual-gateway-address
10.200.100.254
set interfaces irb unit 101 family inet address 10.200.101.252/24 virtual-gateway-address
10.200.101.254
set interfaces irb unit 102 family inet address 10.200.102.252/24 virtual-gateway-address
10.200.102.254
set interfaces irb unit 103 family inet address 10.200.103.252/24 virtual-gateway-address
10.200.103.254
set interfaces irb unit 104 family inet address 10.200.104.252/24 virtual-gateway-address
10.200.104.254

set routing-instances fabric bridge-domains bd_100 routing-interface irb.100
set routing-instances fabric bridge-domains bd_101 routing-interface irb.101
set routing-instances fabric bridge-domains bd_102 routing-interface irb.102
set routing-instances fabric bridge-domains bd_103 routing-interface irb.103
set routing-instances fabric bridge-domains bd_104 routing-interface irb.104
```

vMX4:

```
set interfaces irb unit 100 family inet address 10.200.100.253/24 virtual-gateway-address
10.200.100.254
set interfaces irb unit 101 family inet address 10.200.101.253/24 virtual-gateway-address
10.200.101.254
set interfaces irb unit 102 family inet address 10.200.102.253/24 virtual-gateway-address
10.200.102.254
set interfaces irb unit 103 family inet address 10.200.103.253/24 virtual-gateway-address
10.200.103.254
set interfaces irb unit 104 family inet address 10.200.104.253/24 virtual-gateway-address
10.200.104.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances fabric bridge-domains bd_100 routing-interface irb.100
set routing-instances fabric bridge-domains bd_101 routing-interface irb.101
set routing-instances fabric bridge-domains bd_102 routing-interface irb.102
set routing-instances fabric bridge-domains bd_103 routing-interface irb.103
set routing-instances fabric bridge-domains bd_104 routing-interface irb.104
```

Verification

Check IP-connectivity from vMX6 by sending traffic from one subnet to several other subnets:

```
jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.200.101.4

PING 10.200.101.4 (10.200.101.4): 56 data bytes
!!!!!
--- 10.200.101.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 38.228/46.329/63.100/8.853 m

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.200.102.3

PING 10.200.102.3 (10.200.102.3): 56 data bytes
!!!!!
--- 10.200.102.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 40.135/58.838/79.732/13.058 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.200.103.2

PING 10.200.103.2 (10.200.103.2): 56 data bytes
!!!!!
--- 10.200.103.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 34.788/40.408/45.675/4.249 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL100 10.200.104.1

PING 10.200.104.1 (10.200.104.1): 56 data bytes
!!!!!
--- 10.200.104.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 31.704/38.597/48.408/5.880 ms
```

Check IP connectivity from vMX3:

```
jncie@vMX3> show arp interface irb.100

MAC Address       Address         Name                    Interface               Flags
00:0c:29:8c:00:b2 10.200.100.2    10.200.100.2            irb.100 [vtep.32770]    none
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
00:00:5e:00:01:01 10.200.100.254  10.200.100.254          irb.100                 permanent
published gateway
Total entries: 2

jncie@vMX4> show arp interface irb.100

MAC Address       Address         Name                    Interface               Flags
00:0c:29:8c:00:a8 10.200.100.1    10.200.100.1            irb.100 [vtep.32769]    none
00:0c:29:8c:00:b2 10.200.100.2    10.200.100.2            irb.100 [vtep.32769]    none
00:0c:29:8c:00:bc 10.200.100.3    10.200.100.3            irb.100 [vtep.32770]    none
00:05:86:71:74:c0 10.200.100.4    10.200.100.4            irb.100 [.local..8]     none
00:00:5e:00:01:01 10.200.100.254  10.200.100.254          irb.100                 permanent
published gateway
```

Check the EVPN instance:

```
jncie@vMX3> show evpn instance extensive

Instance: __default_evpn__
  Route Distinguisher: 192.168.1.1:0
  Number of bridge domains: 0
  Number of neighbors: 0

Instance: fabric
  Route Distinguisher: 192.168.1.1:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           5      34
    MAC+IP advertisements:                       10      10
    Default gateway MAC advertisements:          10       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 5 (5 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
    irb.101                101     Up     master
    irb.102                102     Up     master
    irb.103                103     Up     master
    irb.104                104     Up     master
  Number of bridge domains: 5
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0     irb.100    Extended        Enabled   100
    101   101            0    0     irb.101    Extended        Enabled   101
    102   102            0    0     irb.102    Extended        Enabled   102
    103   103            0    0     irb.103    Extended        Enabled   103
    104   104            0    0     irb.104    Extended        Enabled   104
  Number of neighbors: 4
    192.168.1.2
      Received routes
        MAC address advertisement:             10
        MAC+IP address advertisement:          10
        Inclusive multicast:                    5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Ethernet auto-discovery:                5
192.168.1.3
  Received routes
    MAC address advertisement:               10
    MAC+IP address advertisement:             0
    Inclusive multicast:                      5
    Ethernet auto-discovery:                  0
192.168.1.4
  Received routes
    MAC address advertisement:               10
    MAC+IP address advertisement:             0
    Inclusive multicast:                      5
    Ethernet auto-discovery:                  2
  192.168.1.5
    Received routes
      MAC address advertisement:              4
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
Number of ethernet segments: 6
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.4      100        0               all-active
      192.168.1.5      103        0               all-active
  ESI: 05:00:00:fd:e8:00:00:00:64:00
    Status: Resolved by IFL irb.100
    Local interface: irb.100, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.2      100        0               all-active
  ESI: 05:00:00:fd:e8:00:00:00:65:00
    Status: Resolved by IFL irb.101
    Local interface: irb.101, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.2      101        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:66:00
  Status: Resolved by IFL irb.102
  Local interface: irb.102, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.2      102        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:67:00
  Status: Resolved by IFL irb.103
  Local interface: irb.103, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.2      103        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:68:00
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Status: Resolved by IFL irb.104
      Local interface: irb.104, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        192.168.1.2      104        0               all-active
  Router-ID: 192.168.1.1
  Source VTEP interface IP: 192.168.1.1

jncie@vMX4> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 192.168.1.2:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           5      34
    MAC+IP advertisements:                       10      10
    Default gateway MAC advertisements:          10       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 5 (5 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
    irb.101                101     Up     master
    irb.102                102     Up     master
    irb.103                103     Up     master
    irb.104                104     Up     master
  Number of bridge domains: 5
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0     irb.100    Extended        Enabled   100
    101   101            0    0     irb.101    Extended        Enabled   101
    102   102            0    0     irb.102    Extended        Enabled   102
    103   103            0    0     irb.103    Extended        Enabled   103
    104   104            0    0     irb.104    Extended        Enabled   104
  Number of neighbors: 4
    192.168.1.1
      Received routes
        MAC address advertisement:             10
        MAC+IP address advertisement:          10
        Inclusive multicast:                    5
        Ethernet auto-discovery:                5
    192.168.1.3
      Received routes
        MAC address advertisement:             10
        MAC+IP address advertisement:           0
        Inclusive multicast:                    5
        Ethernet auto-discovery:                0
    192.168.1.4
      Received routes
        MAC address advertisement:             10
        MAC+IP address advertisement:           0
        Inclusive multicast:                    5
        Ethernet auto-discovery:                2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  192.168.1.5
    Received routes
      MAC address advertisement:              4
      MAC+IP address advertisement:           0
      Inclusive multicast:                    5
      Ethernet auto-discovery:                2
Number of ethernet segments: 6
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.4      104        0               all-active
      192.168.1.5      103        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:64:00
  Status: Resolved by IFL irb.100
  Local interface: irb.100, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.1      100        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:65:00
  Status: Resolved by IFL irb.101
  Local interface: irb.101, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.1      101        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:66:00
  Status: Resolved by IFL irb.102
  Local interface: irb.102, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.1      102        0               all-active
ESI: 05:00:00:fd:e8:00:00:00:67:00
  Status: Resolved by IFL irb.103
  Local interface: irb.103, Status: Up/Forwarding
  Number of remote PEs connected: 1
    Remote PE        MAC label  Aliasing label  Mode
    192.168.1.1      103        0               all-active
  ESI: 05:00:00:fd:e8:00:00:00:68:00
    Status: Resolved by IFL irb.104
    Local interface: irb.104, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      192.168.1.1      104        0               all-active
Router-ID: 192.168.1.2
Source VTEP interface IP: 192.168.1.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.19: Layer 3 VXLAN Gateway Configuration

- Configure the following layer 3 gateways on vMX3:

VXLAN / VNI	Gateway
105	10.200.105.254
106	10.200.106.254

- Configure the following layer 3 gateways on vMX4:

VXLAN / VNI	Gateway
107	10.200.107.254
108	10.200.108.254

- Only import VXLAN segments into the routing-instance when the vMX offers a gateway for the VXLAN segment.
- Make sure there is IP connectivity between the different subnets

Solution

We import the relevant VXLAN segment into the virtual-switch instance by configuring the VNI inside the routing instance under the [routing-instances fabric protocols evpn vni-options vni xxx vrf-target target:65000:xxx] stanza.

To make sure that there is IP connectivity between the different subnets, we also need to advertise the subnet we configure on the IRB interface into OSPF. We can do this by including the IRB interface in the OSPF configuration. Since we do not want to discover any OSPF neighbor adjacencies in this subnet, we use the keyword 'passive'.

vMX3:

```
set interfaces irb unit 105 family inet address 10.200.105.254/24
set interfaces irb unit 106 family inet address 10.200.106.254/24

set protocols ospf area 0.0.0.0 interface irb.105 passive
set protocols ospf area 0.0.0.0 interface irb.106 passive

set routing-instances fabric protocols evpn vni-options vni 105 vrf-target target:65000:105
set routing-instances fabric protocols evpn vni-options vni 106 vrf-target target:65000:106

set routing-instances fabric bridge-domains bd_105 vlan-id 105
set routing-instances fabric bridge-domains bd_105 routing-interface irb.105
set routing-instances fabric bridge-domains bd_105 vxlan vni 105
set routing-instances fabric bridge-domains bd_106 vlan-id 106
set routing-instances fabric bridge-domains bd_106 routing-interface irb.106
set routing-instances fabric bridge-domains bd_106 vxlan vni 106
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set interfaces irb unit 107 family inet address 10.200.107.254/24
set interfaces irb unit 108 family inet address 10.200.108.254/24

set protocols ospf area 0.0.0.0 interface irb.107 passive
set protocols ospf area 0.0.0.0 interface irb.108 passive

set routing-instances fabric protocols evpn vni-options vni 107 vrf-target target:65000:107
set routing-instances fabric protocols evpn vni-options vni 108 vrf-target target:65000:108

set routing-instances fabric bridge-domains bd_107 vlan-id 107
set routing-instances fabric bridge-domains bd_107 routing-interface irb.107
set routing-instances fabric bridge-domains bd_107 vxlan vni 107
set routing-instances fabric bridge-domains bd_108 vlan-id 108
set routing-instances fabric bridge-domains bd_108 routing-interface irb.108
set routing-instances fabric bridge-domains bd_108 vxlan vni 108
```

Verification

For every configured VXLAN, a separate term was added. Now we check IP connectivity from vMX6:

```
jncie@vMX6> ping rapid routing-instance DC1-S2-VL105 10.200.105.254

PING 10.200.105.254 (10.200.105.254): 56 data bytes
!!!!!
--- 10.200.105.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 17.369/28.920/65.192/18.201 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL106 10.200.106.254

PING 10.200.106.254 (10.200.106.254): 56 data bytes
!!!!!
--- 10.200.106.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 20.794/23.043/29.529/3.270 ms

jncie@vMX6> ping rapid routing-instance DC1-S2-VL107 10.200.107.254

PING 10.200.107.254 (10.200.107.254): 56 data bytes
.!!!!
--- 10.200.107.254 ping statistics ---
5 packets transmitted, 4 packets received, 20% packet loss
round-trip min/avg/max/stddev = 17.563/22.045/25.658/2.895 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping rapid routing-instance DC1-S2-VL108 10.200.108.254

PING 10.200.108.254 (10.200.108.254): 56 data bytes
.!!!!
--- 10.200.108.254 ping statistics ---
5 packets transmitted, 4 packets received, 20% packet loss
round-trip min/avg/max/stddev = 16.976/18.507/21.176/1.685 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.105.1

PING 10.200.105.1 (10.200.105.1): 56 data bytes
!!!!!
--- 10.200.105.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 23.938/28.668/33.126/2.996 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.106.1

PING 10.200.106.1 (10.200.106.1): 56 data bytes
!!!!!
--- 10.200.106.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 38.967/50.764/89.988/19.708 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.107.1

PING 10.200.107.1 (10.200.107.1): 56 data bytes
!!!!!
--- 10.200.107.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 28.311/33.819/46.007/6.506 ms

jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.108.1

PING 10.200.108.1 (10.200.108.1): 56 data bytes
!!!!!
--- 10.200.108.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 39.236/42.463/44.663/2.022 ms
```

We check what address information was placed into the EVPN database:

```
jncie@vMX3> show evpn database | match "10[5-8]"

     105        00:05:86:4d:34:f0  irb.105                                 Mar 29 07:02:57
10.200.105.254
     105        00:05:86:71:74:c0  00:11:11:11:11:11:11:11:11:11  Mar 29 07:06:42
     105        00:0c:29:8c:00:a8  192.168.1.3                             Mar 29 07:03:04
     105        00:0c:29:8c:00:b2  192.168.1.3                             Mar 29 07:03:04
     105        00:0c:29:8c:00:bc  192.168.1.4                             Mar 29 07:06:42
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
     106         00:05:86:4d:34:f0  irb.106                              Mar 29 07:02:57
10.200.106.254
     106         00:05:86:71:74:c0  00:11:11:11:11:11:11:11:11:11        Mar 29 07:06:42
     106         00:0c:29:8c:00:a8  192.168.1.3                          Mar 29 07:03:04
     106         00:0c:29:8c:00:b2  192.168.1.3                          Mar 29 07:03:04
     106         00:0c:29:8c:00:bc  192.168.1.4                          Mar 29 07:06:42

jncie@vMX4> show evpn database | match "10[5-8]"
     107         00:05:86:71:74:c0  00:11:11:11:11:11:11:11:11:11        Mar 29 07:01:15
     107         00:05:86:82:68:f0  irb.107                              Mar 29 06:57:58
10.200.107.254
     107         00:0c:29:8c:00:a8  192.168.1.3                          Mar 29 07:06:38
     107         00:0c:29:8c:00:b2  192.168.1.3                          Mar 29 07:06:38
     107         00:0c:29:8c:00:bc  192.168.1.4                          Mar 29 07:01:15
     108         00:05:86:71:74:c0  00:11:11:11:11:11:11:11:11:11        Mar 29 07:01:23
     108         00:05:86:82:68:f0  irb.108                              Mar 29 06:57:58
10.200.108.254
     108         00:0c:29:8c:00:a8  192.168.1.3                          Mar 29 07:06:38
     108         00:0c:29:8c:00:b2  192.168.1.3                          Mar 29 07:06:38
     108         00:0c:29:8c:00:bc  192.168.1.4                          Mar 29 07:01:15
```

To check and see if the routes were advertised correctly using OSPF:

```
jncie@vMX3> show route protocol ospf

inet.0: 37 destinations, 37 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.107.0/24    *[OSPF/10] 00:00:05, metric 3
                      to 172.16.6.2 via xe-0/0/4.0
                    > to 172.16.2.2 via xe-0/0/2.0
10.200.108.0/24    *[OSPF/10] 00:00:05, metric 3
                      to 172.16.6.2 via xe-0/0/4.0
                    > to 172.16.2.2 via xe-0/0/2.0
...

jncie@vMX4> show route protocol ospf

inet.0: 37 destinations, 37 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.200.105.0/24    *[OSPF/10] 00:00:01, metric 3
                      to 172.16.7.2 via xe-0/0/2.0
                      to 172.16.5.2 via xe-0/0/3.0
                    > to 172.16.3.2 via xe-0/0/4.0
10.200.106.0/24    *[OSPF/10] 00:00:01, metric 3
                      to 172.16.7.2 via xe-0/0/2.0
                    > to 172.16.5.2 via xe-0/0/3.0
                      to 172.16.3.2 via xe-0/0/4.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 3.20: Layer 3 VXLAN Gateway Configuration

- Configure the following anycast gateways on vMX3 and vMX4:

VXLAN / VNI	Gateway
109	10.200.109.254
110	10.200.110.254

Solution

For this task, we need to import the VNIs into the fabric instance and we need to configure an anycast gateway. To configure an anycast gateway, we simply need to configure an IRB interface with an identical IP and MAC address on the two vMX routers.

vMX3:

```
set interfaces irb unit 109 family inet address 10.200.109.254/24
set interfaces irb unit 109 mac aa:aa:aa:aa:a1:09
set interfaces irb unit 110 family inet address 10.200.110.254/24
set interfaces irb unit 110 mac aa:aa:aa:aa:a1:10

set routing-instances fabric protocols evpn vni-options vni 109 vrf-target target:65000:109
set routing-instances fabric protocols evpn vni-options vni 110 vrf-target target:65000:110

set routing-instances fabric bridge-domains bd_109 vlan-id 109
set routing-instances fabric bridge-domains bd_109 routing-interface irb.109
set routing-instances fabric bridge-domains bd_109 vxlan vni 109
set routing-instances fabric bridge-domains bd_110 vlan-id 110
set routing-instances fabric bridge-domains bd_110 routing-interface irb.110
set routing-instances fabric bridge-domains bd_110 vxlan vni 110
```

vMX4:

```
set interfaces irb unit 109 family inet address 10.200.109.254/24
set interfaces irb unit 109 mac aa:aa:aa:aa:a1:09
set interfaces irb unit 110 family inet address 10.200.110.254/24
set interfaces irb unit 110 mac aa:aa:aa:aa:a1:10

set routing-instances fabric protocols evpn vni-options vni 109 vrf-target target:65000:109
set routing-instances fabric protocols evpn vni-options vni 110 vrf-target target:65000:110

set routing-instances fabric bridge-domains bd_109 vlan-id 109
set routing-instances fabric bridge-domains bd_109 routing-interface irb.109
set routing-instances fabric bridge-domains bd_109 vxlan vni 109
set routing-instances fabric bridge-domains bd_110 vlan-id 110
set routing-instances fabric bridge-domains bd_110 routing-interface irb.110
set routing-instances fabric bridge-domains bd_110 vxlan vni 110
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

Check IP connectivity from vMX6:

```
jncie@vMX6> ping rapid routing-instance DC1-S4-VL109 10.200.109.254

PING 10.200.109.254 (10.200.109.254): 56 data bytes
!!!!!
--- 10.200.109.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.817/18.328/22.679/2.634 ms

jncie@vMX6> ping rapid routing-instance DC1-S4-VL110 10.200.110.254

PING 10.200.110.254 (10.200.110.254): 56 data bytes
!!!!!
--- 10.200.110.254 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.147/17.166/18.198/1.168 ms
```

Checking connectivity between the different subnets:

```
jncie@vMX6> ping rapid routing-instance DC1-S1-VL100 10.200.110.3

PING 10.200.110.3 (10.200.110.3): 56 data bytes
!!!!!
--- 10.200.110.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 30.174/32.263/33.892/1.341 ms

jncie@vMX6> show arp vpn DC1-S4-VL109

MAC Address       Address         Name                      Interface               Flags
aa:aa:aa:aa:a1:09 10.200.109.254  10.200.109.254            ae0.109                 none

jncie@vMX6> show arp vpn DC1-S4-VL110
MAC Address       Address         Name                      Interface               Flags
aa:aa:aa:aa:a1:10 10.200.110.254  10.200.110.254            ae0.110                 none
```

Check the switching table on the vQFX:

```
jncie@vQFX5> show ethernet-switching table vlan-id 110

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 5 entries, 5 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                 Active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
name                address             flags    interface              source
vlan-110            00:05:86:71:7b:c0   DR       ae0.0
vlan-110            00:0c:29:8c:00:a8   D        vtep.32770             192.168.1.3
vlan-110            00:0c:29:8c:00:b2   D        vtep.32770             192.168.1.3
vlan-110            00:0c:29:8c:00:bc   D        vtep.32769             192.168.1.4
vlan-110            aa:aa:aa:aa:a1:10   D        vtep.32772             192.168.1.2
```

Check the EVPN database on the vQFX:

```
jncie@vQFX5> show evpn database l2-domain-id 110 extensive

Instance: default-switch

VN Identifier: 110, MAC address: 00:05:86:71:74:c0
  Source: 00:11:11:11:11:11:11:11:11:11, Rank: 1, Status: Active
    Remote origin: 192.168.1.4
    Timestamp: Mar 27 12:06:39 (0x58d9004f)
    State: <Remote-To-Local-Adv-Done>

VN Identifier: 110, MAC address: 00:0c:29:8c:00:a8
  Source: 192.168.1.3, Rank: 1, Status: Active
    Timestamp: Mar 27 12:20:28 (0x58d9038c)
    State: <Remote-To-Local-Adv-Done>

VN Identifier: 110, MAC address: 00:0c:29:8c:00:b2
  Source: 192.168.1.3, Rank: 1, Status: Active
    Timestamp: Mar 27 12:20:28 (0x58d9038c)
    State: <Remote-To-Local-Adv-Done>

VN Identifier: 110, MAC address: 00:0c:29:8c:00:bc
  Source: 192.168.1.4, Rank: 1, Status: Active
    Timestamp: Mar 27 12:00:32 (0x58d8fee0)
    State: <Remote-To-Local-Adv-Done>

VN Identifier: 110, MAC address: aa:aa:aa:aa:a1:10
  Source: 192.168.1.1, Rank: 1, Status: Active
    Timestamp: Mar 27 12:20:32 (0x58d90390)
    State: <Remote-To-Local-Adv-Done>
    IP address: 10.200.110.254
  Source: 192.168.1.2, Rank: 2, Status: Inactive
    Timestamp: Mar 27 12:00:32 (0x58d8fee0)
    State: <>
    IP address: 10.200.110.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The anycast gateway is advertised using the default ESI:

```
jncie@vMX4> show route advertising-protocol bgp 192.168.1.5 evpn-ethernet-tag-id 110 detail
table fabric.evpn.0

...

* 2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP (1 entry, 1 announced)
 BGP group overlay type Internal
     Route Distinguisher: 192.168.1.2:8
     Route Label: 110
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:110 encapsulation:vxlan(0x8)

...
```

Therefore, the rest of the switches will use only one vMX. I you want to be able to load-share traffic towards an anycast gateway, you would need advertise the VXLAN segment with a non-default ESI.

Check routes advertised from vMX3:

```
jncie@vMX3> show route advertising-protocol bgp 192.168.1.4 evpn-ethernet-tag-id 110

bgp.evpn.0: 175 destinations, 258 routes (175 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         192.168.1.2                  100        I
  2:192.168.1.3:65000::110::00:0c:29:8c:00:a8/304 MAC/IP
*                         192.168.1.3                  100        I
  2:192.168.1.3:65000::110::00:0c:29:8c:00:b2/304 MAC/IP
*                         192.168.1.3                  100        I
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         192.168.1.2                  100        I
  3:192.168.1.1:8::110::192.168.1.1/304 IM
*                         Self                         100        I
  3:192.168.1.2:8::110::192.168.1.2/304 IM
*                         192.168.1.2                  100        I
  3:192.168.1.3:65000::110::192.168.1.3/304 IM
*                         192.168.1.3                  100        I
  3:192.168.1.5:65000::110::192.168.1.5/304 IM
*                         192.168.1.5                  100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
fabric.evpn.0: 144 destinations, 211 routes (144 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         Self                         100        I
  3:192.168.1.1:8::110::192.168.1.1/304 IM
*                         Self                         100        I

__default_evpn__.evpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
```

Check routes advertised from vMX4:

```
jncie@vMX4> show route advertising-protocol bgp 192.168.1.4 evpn-ethernet-tag-id 110

bgp.evpn.0: 175 destinations, 258 routes (175 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         192.168.1.1                  100        I
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.3:65000::110::00:0c:29:8c:00:a8/304 MAC/IP
*                         192.168.1.3                  100        I
  2:192.168.1.3:65000::110::00:0c:29:8c:00:b2/304 MAC/IP
*                         192.168.1.3                  100        I
  2:192.168.1.1:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         192.168.1.1                  100        I
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         Self                         100        I
  3:192.168.1.1:8::110::192.168.1.1/304 IM
*                         192.168.1.1                  100        I
  3:192.168.1.2:8::110::192.168.1.2/304 IM
*                         Self                         100        I
  3:192.168.1.3:65000::110::192.168.1.3/304 IM
*                         192.168.1.3                  100        I
  3:192.168.1.5:65000::110::192.168.1.5/304 IM
*                         192.168.1.5                  100        I

fabric.evpn.0: 144 destinations, 211 routes (144 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10/304 MAC/IP
*                         Self                         100        I
  2:192.168.1.2:8::110::aa:aa:aa:aa:a1:10::10.200.110.254/304 MAC/IP
*                         Self                         100        I
  3:192.168.1.2:8::110::192.168.1.2/304 IM
*                         Self                         100        I

__default_evpn__.evpn.0: 116 destinations, 183 routes (116 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 3: IP Fabric and Controllerless Overlay

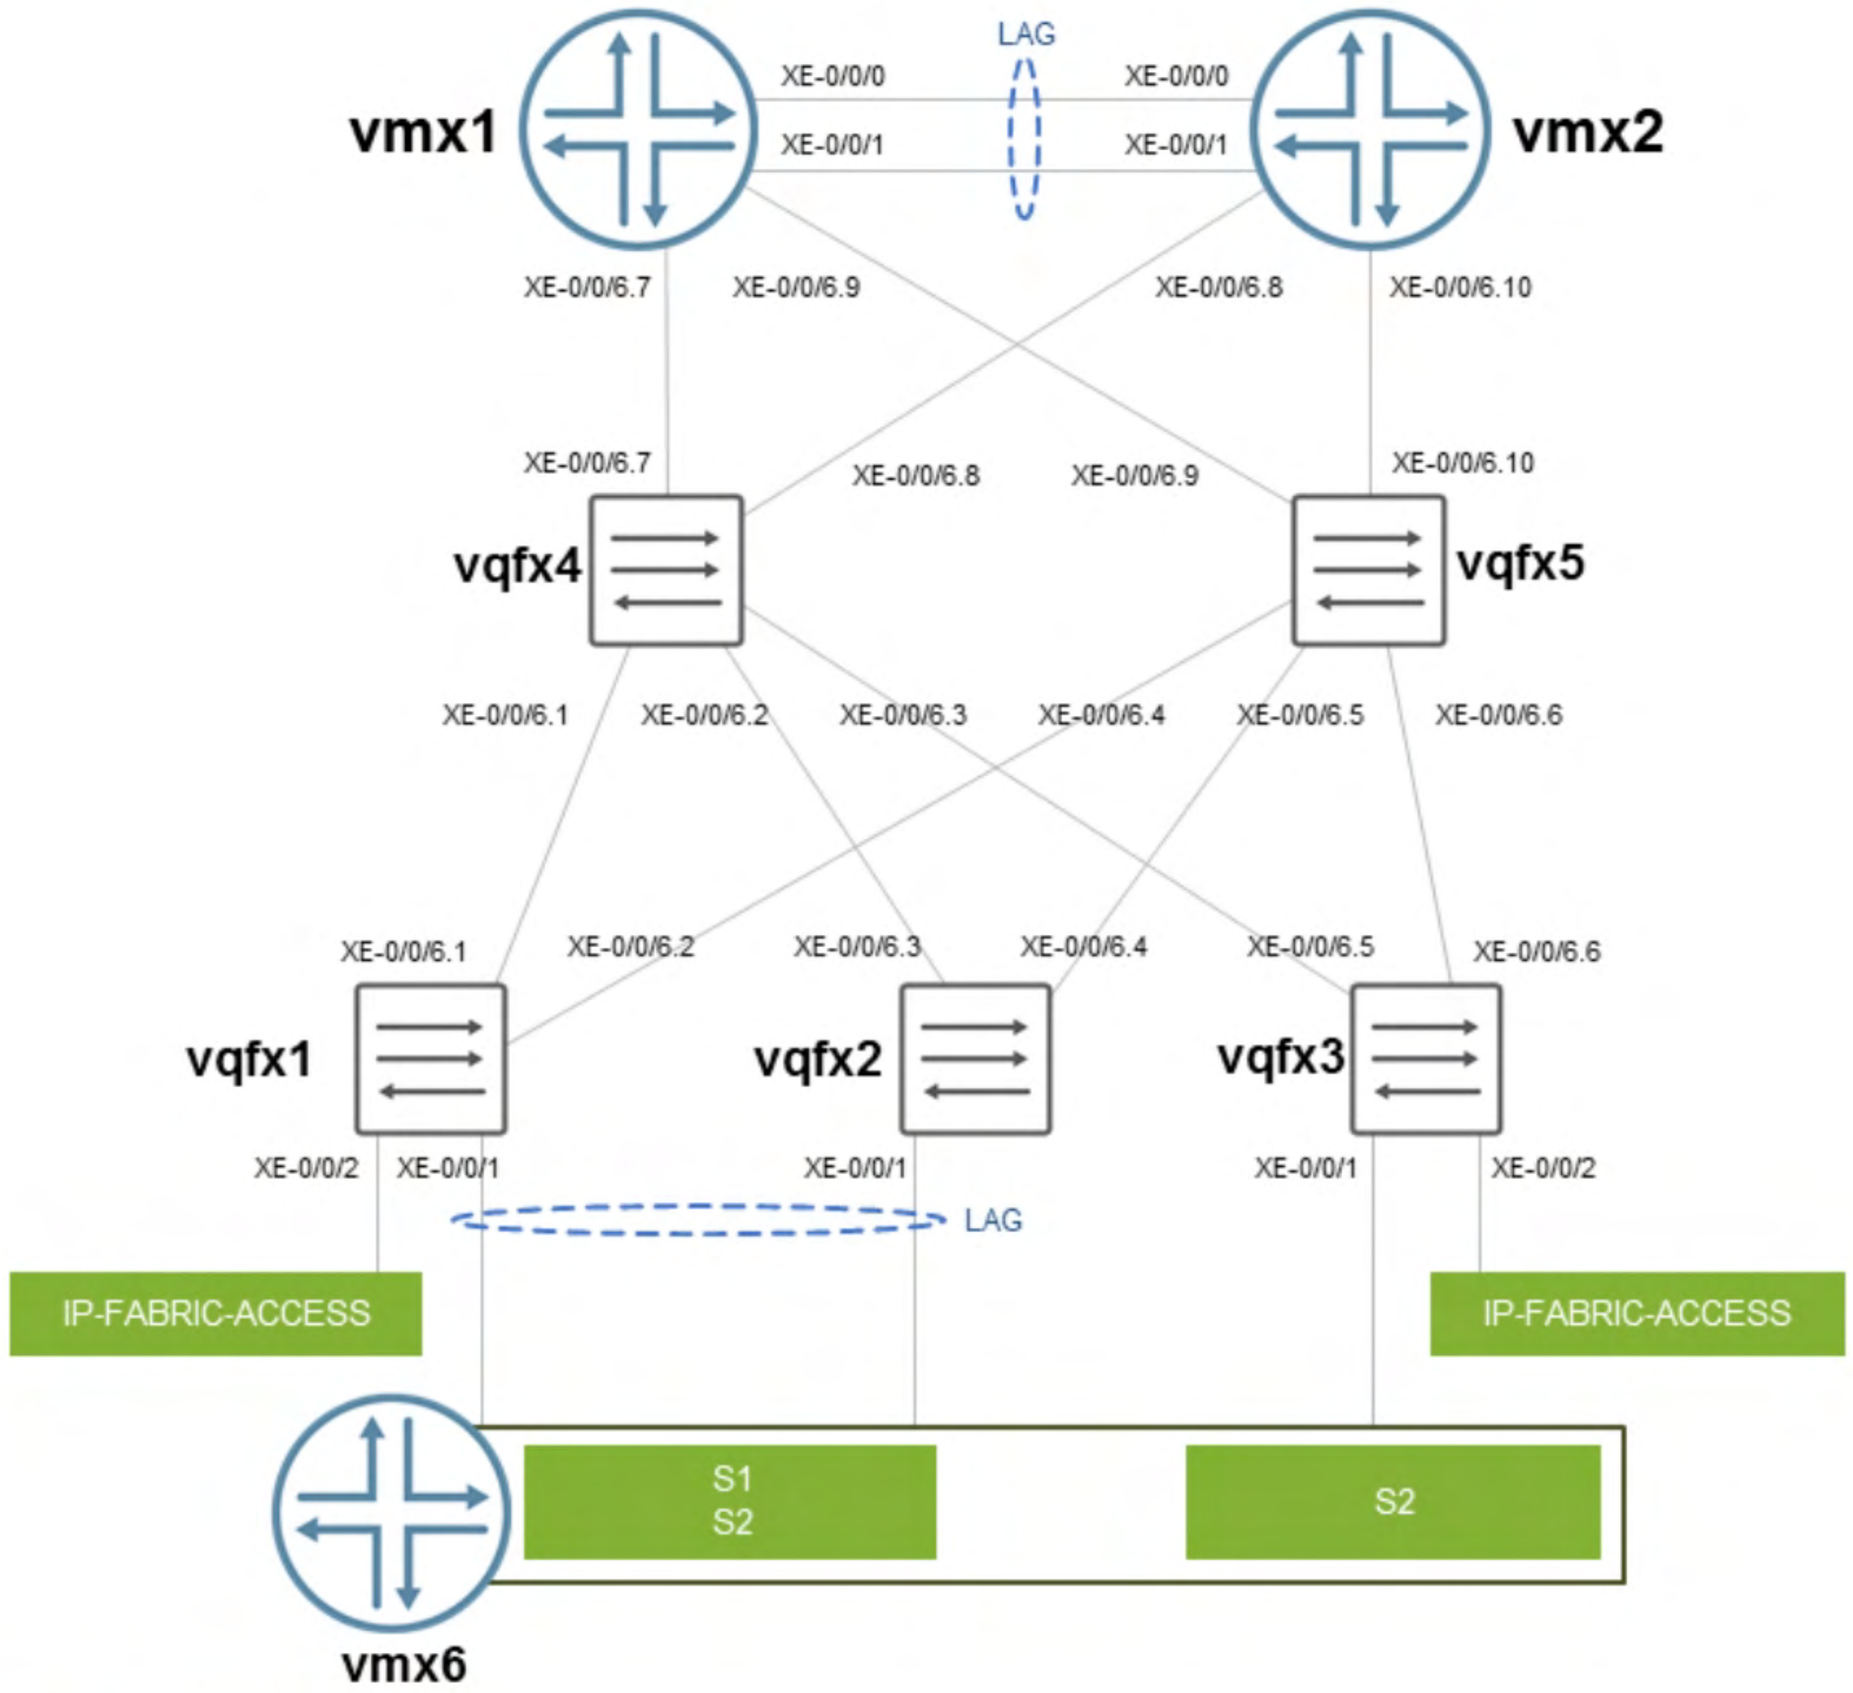

Figure 21 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

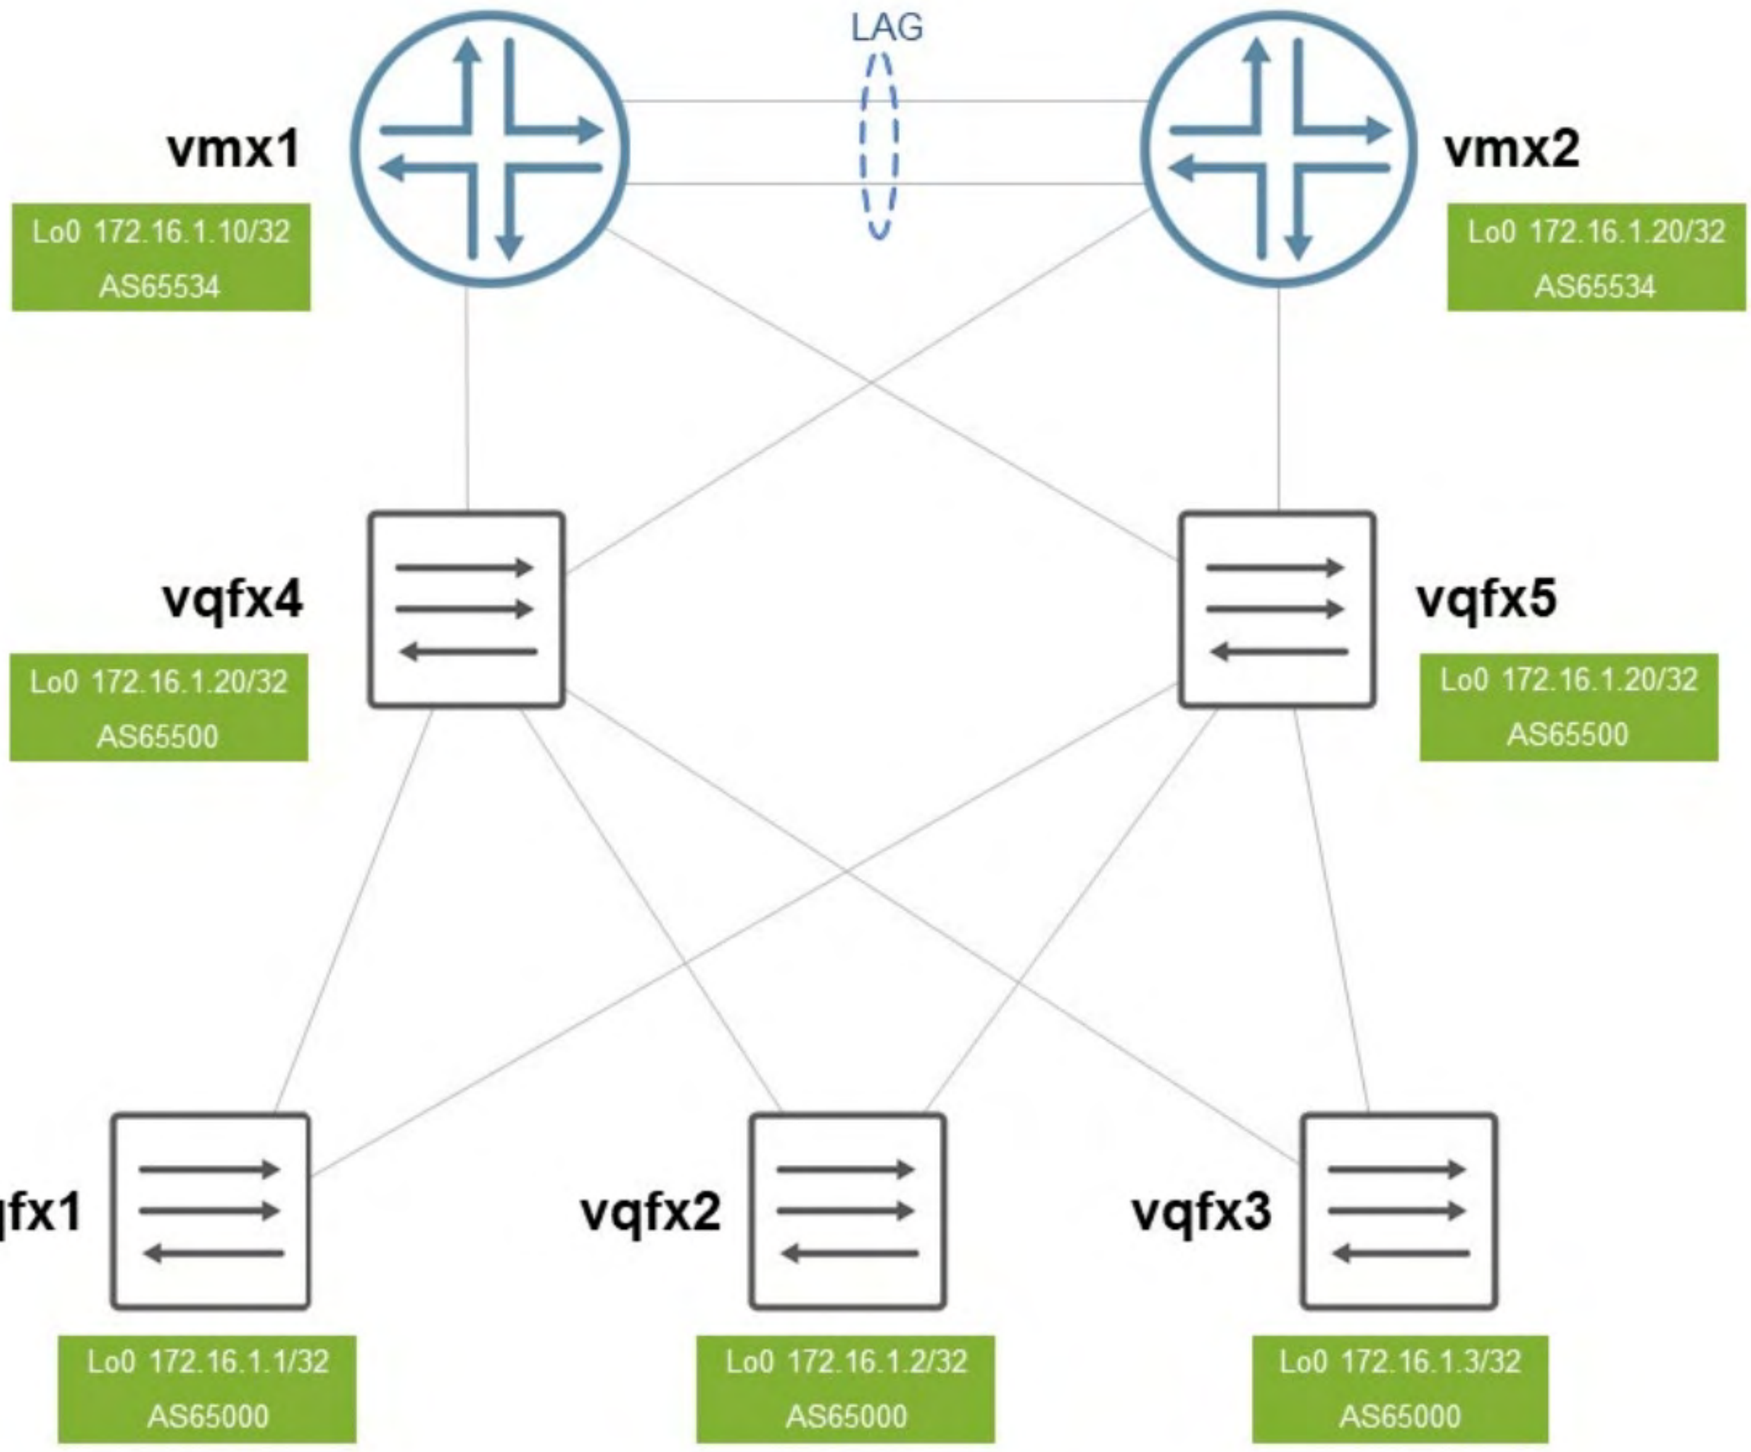

Figure 2 - Logical topology

Note: starting topology is preconfigured with IP-addressing only

Task 3.21: EBGP for the IP-Fabric

- Configure the devices displayed in the topology to form an EBGP-based IP-Fabric.
- Members of the IP-Fabric should advertise their own loopback IP address.
- Configure the same export policy on every device. Do not use the 'interface' statement in any of the match conditions.
- Use the autonomous systems displayed in the logical topology.
- Members of the fabric that have equal-cost paths towards other nodes should utilize all available links.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Solution

vMX1:

```
set routing-options autonomous-system 65534
set routing-options forwarding-table export lbpp

set protocols bgp group core-pod export advertise-lo0
set protocols bgp group core-pod peer-as 65500
set protocols bgp group core-pod multipath
set protocols bgp group core-pod neighbor 192.168.7.1
set protocols bgp group core-pod neighbor 192.168.9.1

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vMX2:

```
set routing-options autonomous-system 65534
set routing-options forwarding-table export lbpp

set protocols bgp group core-pod export advertise-lo0
set protocols bgp group core-pod peer-as 65500
set protocols bgp group core-pod multipath
set protocols bgp group core-pod neighbor 192.168.8.1
set protocols bgp group core-pod neighbor 192.168.10.1

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vQFX4:

```
set routing-options autonomous-system 65500

set protocols bgp group pod-spine-leaf advertise-peer-as
set protocols bgp group pod-spine-leaf export advertise-lo0
set protocols bgp group pod-spine-leaf peer-as 65000
set protocols bgp group pod-spine-leaf neighbor 192.168.1.0
set protocols bgp group pod-spine-leaf neighbor 192.168.3.0
set protocols bgp group pod-spine-leaf neighbor 192.168.5.0

set protocols bgp group pod-spine-core export advertise-lo0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group pod-spine-core peer-as 65534
set protocols bgp group pod-spine-core neighbor 192.168.7.0
set protocols bgp group pod-spine-core neighbor 192.168.8.0

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept
```

vQFX5:

```
set routing-options autonomous-system 65500

set protocols bgp group pod-spine-leaf advertise-peer-as
set protocols bgp group pod-spine-leaf export advertise-lo0
set protocols bgp group pod-spine-leaf peer-as 65000
set protocols bgp group pod-spine-leaf neighbor 192.168.2.0
set protocols bgp group pod-spine-leaf neighbor 192.168.4.0
set protocols bgp group pod-spine-leaf neighbor 192.168.6.0

set protocols bgp group pod-spine-core export advertise-lo0
set protocols bgp group pod-spine-core peer-as 65534
set protocols bgp group pod-spine-core neighbor 192.168.9.0
set protocols bgp group pod-spine-core neighbor 192.168.10.0

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept
```

vQFX1:

```
set routing-options autonomous-system 65000
set routing-options autonomous-system loops 2
set routing-options forwarding-table export lbpp

set protocols bgp group pod-leaf export advertise-lo0
set protocols bgp group pod-leaf peer-as 65500
set protocols bgp group pod-leaf multipath
set protocols bgp group pod-leaf neighbor 192.168.1.1
set protocols bgp group pod-leaf neighbor 192.168.2.1

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX2:

```
set routing-options autonomous-system 65000
set routing-options autonomous-system loops 2
set routing-options forwarding-table export lbpp

set protocols bgp group pod-leaf export advertise-lo0
set protocols bgp group pod-leaf peer-as 65500
set protocols bgp group pod-leaf multipath
set protocols bgp group pod-leaf neighbor 192.168.3.1
set protocols bgp group pod-leaf neighbor 192.168.4.1

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vQFX3:

```
set routing-options autonomous-system 65000
set routing-options autonomous-system loops 2
set routing-options forwarding-table export lbpp

set protocols bgp group pod-leaf export advertise-lo0
set protocols bgp group pod-leaf peer-as 65500
set protocols bgp group pod-leaf multipath
set protocols bgp group pod-leaf neighbor 192.168.5.1
set protocols bgp group pod-leaf neighbor 192.168.6.1

set policy-options policy-statement advertise-lo0 term lo0 from protocol direct
set policy-options policy-statement advertise-lo0 term lo0 from route-filter 172.16.1.0/24
prefix-length-range /32-/32
set policy-options policy-statement advertise-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

Verification

Verifying the BGP neighbor establishment:

```
jncie@vQFX4> show bgp summary

Groups: 2 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
192.168.1.0             65000          8         11          0          0        2:08 1/1/1/0
0/0/0/0
192.168.3.0             65000          6          8          0          0        1:19 1/1/1/0
0/0/0/0
192.168.5.0             65000          6          9          0          0        1:39 1/1/1/0
0/0/0/0
192.168.7.0             65534          8         12          0          0        2:21 1/1/1/0
0/0/0/0
192.168.8.0             65534          6         11          0          0        2:13 1/1/1/0
0/0/0/0

jncie@vQFX5> show bgp summary
Groups: 2 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.2.0             65000          7         11          0          0        2:09 1/1/1/0
0/0/0/0
192.168.4.0             65000          5          8          0          0        1:23 1/1/1/0
0/0/0/0
192.168.6.0             65000          6          9          0          0        1:45 1/1/1/0
0/0/0/0
192.168.9.0             65534          8         12          0          0        2:22 1/1/1/0
0/0/0/0
192.168.10.0            65534          7         10          0          0        2:10 1/1/1/0
0/0/0/0
```

Verifying the IP connectivity between loopback IP addresses:

```
{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.2
PING 172.16.1.2 (172.16.1.2): 56 data bytes
!!!!!
--- 172.16.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 46.453/49.425/55.849/3.628 ms

{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.3

PING 172.16.1.3 (172.16.1.3): 56 data bytes
!!!!!
--- 172.16.1.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 43.306/45.649/52.438/3.412 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.4

PING 172.16.1.4 (172.16.1.4): 56 data bytes
!!!!!
--- 172.16.1.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 31.653/33.397/35.774/1.470 ms

{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.5

PING 172.16.1.5 (172.16.1.5): 56 data bytes
!!!!!
--- 172.16.1.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 30.628/31.978/35.227/1.653 ms

{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.10

PING 172.16.1.10 (172.16.1.10): 56 data bytes
!!!!!
--- 172.16.1.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 28.004/29.230/31.079/1.086 ms

{master:0}
jncie@vQFX1> ping rapid source 172.16.1.1 172.16.1.20

PING 172.16.1.20 (172.16.1.20): 56 data bytes
!!!!!
--- 172.16.1.20 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 20.519/21.715/23.174/1.027 ms
```

Verifying load balancing from the leaf switches to the other leaf switches and from the leaf switches to the vMX routers:

```
jncie@vQFX1> show route forwarding-table destination 172.16.1.2

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.16.1.2/32      user     0                    ulst   131070     5
                              192.168.1.1        ucst     1697     5 xe-0/0/6.1
                              192.168.2.1        ucst     1698     5 xe-0/0/6.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX1> show route forwarding-table destination 172.16.1.10

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.16.1.10/32     user     0                    ulst   131070     5
                              192.168.1.1       ucst     1697     5 xe-0/0/6.1
                              192.168.2.1       ucst     1698     5 xe-0/0/6.2
```

Verifying load balancing from the vMX routers to the leaf switches:

```
jncie@vMX1> show route forwarding-table destination 172.16.1.1
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.16.1.1/32      user     0                    ulst  1048574     4
                              192.168.7.1       ucst      551     5 xe-0/0/6.7
                              192.168.9.1       ucst      550     5 xe-0/0/6.9
```

Task 3.22: Completing the IP-Fabric

- BGP speakers should be able to detect peer loss in 9 seconds without the use of BFD.
- All control plane protocol messages exchanged should use the MD5-hased password ‘jncie’.
- Advertise gateway addresses configured on interface xe-0/0/2 on both vQFX1 as well as vQFX3.
- Whenever an IP address is added to interface xe-0/0/2 on either vQFX1 or vQFX3, it should be advertised without any routing-policy change.

Solution

vMX1:

```
set protocols bgp group core-pod hold-time 9
set protocols bgp group core-pod authentication-key jncie
```

vMX2:

```
set protocols bgp group core-pod hold-time 9
set protocols bgp group core-pod authentication-key jncie
```

vQFX4:

```
set protocols bgp group pod-spine-leaf hold-time 9
set protocols bgp group pod-spine-leaf authentication-key jncie
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group pod-spine-core hold-time 9
set protocols bgp group pod-spine-core authentication-key jncie
```

vQFX5:

```
set protocols bgp group pod-spine-leaf hold-time 9
set protocols bgp group pod-spine-leaf authentication-key jncie

set protocols bgp group pod-spine-core hold-time 9
set protocols bgp group pod-spine-core authentication-key jncie
```

vQFX1:

```
set protocols bgp group pod-leaf hold-time 9
set protocols bgp group pod-leaf authentication-key jncie
set protocols bgp group pod-leaf export advertise-direct

set policy-options policy-statement advertise-direct term xe-0/0/2 from protocol direct
set policy-options policy-statement advertise-direct term xe-0/0/2 from prefix-list fabric-
direct-xe-0/0/2
set policy-options policy-statement advertise-direct term xe-0/0/2 then accept

set policy-options prefix-list fabric-direct-xe-0/0/2 apply-path "interfaces xe-0/0/2 unit <*>
family inet address <*>"
```

vQFX2:

```
set protocols bgp group pod-leaf hold-time 9
set protocols bgp group pod-leaf authentication-key jncie
```

vQFX3:

```
set protocols bgp group pod-leaf hold-time 9
set protocols bgp group pod-leaf authentication-key jncie
set protocols bgp group pod-leaf export advertise-direct

set policy-options prefix-list fabric-direct-xe-0/0/2 apply-path "interfaces xe-0/0/2 unit <*>
family inet address <*>"

set policy-options policy-statement advertise-direct term xe-0/0/2 from protocol direct
set policy-options policy-statement advertise-direct term xe-0/0/2 from prefix-list fabric-
direct-xe-0/0/2
set policy-options policy-statement advertise-direct term xe-0/0/2 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

Verifying the hold time and the use of an MD5-hased password:

```
jncie@vMX1> show bgp neighbor 192.168.7.1

Peer: 192.168.7.1+53353 AS 65500 Local: 192.168.7.0+179 AS 65534
  Group: core-pod              Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ advertise-lo0 ]
  Options: <Preference HoldTime AuthKey PeerAS Multipath Refresh>
  Authentication key is configured
  Holdtime: 9 Preference: 170
  Number of flaps: 1
  Last flap event: RecvNotify
  Error: 'Cease' Sent: 0 Recv: 1
  Peer ID: 172.16.1.4      Local ID: 172.16.1.10       Active Holdtime: 9
  Keepalive Interval: 3          Group index: 0    Peer index: 2
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  Local Interface: xe-0/0/6.7
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65500)
  Peer does not support Addpath
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              16
    Received prefixes:            16
    Accepted prefixes:            16
    Suppressed due to damping:    0
    Advertised prefixes:          1
  Last traffic (seconds): Received 670477 Sent 95 Checked 670477
  Input messages:  Total 43     Updates 7       Refreshes 0     Octets 1083
  Output messages: Total 37     Updates 1       Refreshes 0     Octets 736
  Output Queue[1]: 0            (inet.0, inet-unicast)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verifying the prefixes advertised from vQFX1:

```
jncie@vQFX1> show configuration policy-options prefix-list fabric-direct-xe-0/0/2 | display
inheritance

##
## apply-path was expanded to:
##     192.168.110.0/24;
##     192.168.111.0/24;
##     192.168.112.0/24;
##     192.168.113.0/24;
##     192.168.114.0/24;
##     192.168.115.0/24;
##
apply-path "interfaces xe-0/0/2 unit <*> family inet address <*>";

jncie@vQFX1> show route advertising-protocol bgp 192.168.1.1

inet.0: 33 destinations, 43 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.1/32           Self                                    I
* 192.168.110.0/24        Self                                    I
* 192.168.111.0/24        Self                                    I
* 192.168.112.0/24        Self                                    I
* 192.168.113.0/24        Self                                    I
* 192.168.114.0/24        Self                                    I
* 192.168.115.0/24        Self                                    I
```

Task 3.23: Controllerless Overlay

- Have vQFX1, vQFX2 and vQFX3 form an overlay. Use the IP-Fabric as underlay network.
- The servers S1 and S2 should be able to communicate inside VLAN 200 and VLAN 300. The servers are expecting tagged frames. It should be possible to tunnel VLANs inside VLAN 300.
- Ensure that the connections to S1 and S2 are redundant.
- New VLANs added to the server-facing interfaces should only require configuration in the [edit vlans] and [edit interfaces] stanza.
- Every VNI should have a different RT.

Solution

By default, tunneled frames will have their VLAN stripped. If VLANs are being tunneled inside a VLAN, the VLAN information should also be encapsulated. The configuration knob for this is located inside two different stanzas. We need to configure the VXLAN itself to encapsulate the VLAN:

```
set vlans vlan vxlan encapsulate-inner-vlan
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

In addition to this, we need to configure the QFX to accept VXLAN packets that still contain the inner-VLAN. This is a global setting:

```
set protocols l2-learning decapsulate-accept-inner-vlan
```

These commands have been tested on actual QFX5100 switches. You can ignore notifications about this configuration not working on the vQFX.

vQFX1:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 esi 00:12:12:12:12:12:12:12:12:12
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 12:12:12:12:12:12
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 200
set interfaces ae0 unit 0 family ethernet-switching vlan members 300
set interfaces ae0 unit 0 family ethernet-switching vlan members 400

set routing-options router-id 172.16.1.1

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.16.1.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65250
set protocols bgp group overlay neighbor 172.16.1.2
set protocols bgp group overlay neighbor 172.16.1.3

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list all
set protocols evpn multicast-mode ingress-replication

set protocols l2-learning decapsulate-accept-inner-vlan

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.16.1.1:65000
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set vlans vlan-200 vlan-id 200
set vlans vlan-200 vxlan vni 200
set vlans vlan-300 vlan-id 300
set vlans vlan-300 vxlan vni 300
set vlans vlan-300 vxlan encapsulate-inner-vlan
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX2:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 esi 00:12:12:12:12:12:12:12:12:12
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 12:12:12:12:12:12
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 200
set interfaces ae0 unit 0 family ethernet-switching vlan members 300
set interfaces ae0 unit 0 family ethernet-switching vlan members 400

set routing-options router-id 172.16.1.2

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.16.1.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65250
set protocols bgp group overlay neighbor 172.16.1.3
set protocols bgp group overlay neighbor 172.16.1.1

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list all
set protocols evpn multicast-mode ingress-replication

set protocols l2-learning decapsulate-accept-inner-vlan

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.16.1.2:65000
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set vlans vlan-200 vlan-id 200
set vlans vlan-200 vxlan vni 200
set vlans vlan-300 vlan-id 300
set vlans vlan-300 vxlan vni 300
set vlans vlan-300 vxlan encapsulate-inner-vlan
```

vQFX3:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 200
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 300

set routing-options router-id 172.16.1.3

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.16.1.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65250
set protocols bgp group overlay neighbor 172.16.1.2
set protocols bgp group overlay neighbor 172.16.1.1

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list all
set protocols evpn multicast-mode ingress-replication

set protocols l2-learning decapsulate-accept-inner-vlan

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.16.1.3:65000
set switch-options vrf-target target:1:1
set switch-options vrf-target auto

set vlans vlan-200 vlan-id 200
set vlans vlan-200 vxlan vni 200
set vlans vlan-300 vlan-id 300
set vlans vlan-300 vxlan vni 300
set vlans vlan-300 vxlan encapsulate-inner-vlan
```

Verification

Verify the IBGP sessions that make up the overlay:

```
jncie@vQFX1> show bgp summary group overlay

Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      22         22          0          0          0          0
bgp.evpn.0
                       7          7          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.2             65250         25         45       0       0        5:15 Establ
  bgp.evpn.0: 5/5/5/0
  default-switch.evpn.0: 4/4/4/0
  __default_evpn__.evpn.0: 1/1/1/0
172.16.1.3             65250         17         39       0       0        5:06 Establ
  bgp.evpn.0: 2/2/2/0
  default-switch.evpn.0: 2/2/2/0
  __default_evpn__.evpn.0: 0/0/0/0
```

Check the EVPN extensive information:

```
jncie@vQFX1> show evpn instance extensive

Instance: __default_evpn__
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Route Distinguisher: 172.16.1.1:0
  Number of bridge domains: 0
  Number of neighbors: 1
    172.16.1.2
      Received routes
        Ethernet Segment:                       1

Instance: default-switch
  Route Distinguisher: 172.16.1.1:65000
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         0       4
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
Number of local interfaces: 1 (1 up)
  Interface name  ESI                               Mode             Status
  ae0.0           00:12:12:12:12:12:12:12:12:12     all-active       Up
Number of IRB interfaces: 0 (0 up)
Number of bridge domains: 2
  VLAN  VNI    Intfs / up    IRB intf   Mode             MAC sync  IM route label
  200   200         1   1               Extended         Enabled   200
  300   300         1   1               Extended         Enabled   300
Number of neighbors: 2
  172.16.1.2
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    2
      Ethernet auto-discovery:                2
  172.16.1.3
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    2
      Ethernet auto-discovery:                0
Number of ethernet segments: 1
  ESI: 00:12:12:12:12:12:12:12:12:12
    Status: Resolved by IFL ae0.0
    Local interface: ae0.0, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.16.1.2       300        0               all-active
    Designated forwarder: 172.16.1.1
    Backup forwarder: 172.16.1.2
Router-ID: 172.16.1.1
Source VTEP interface IP: 172.16.1.1
```

Verify IP connectivity from vMX6:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping routing-instance S1 rapid 192.168.200.2

PING 192.168.200.2 (192.168.200.2): 56 data bytes
!!!!!
--- 192.168.200.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 71.274/101.050/125.136/20.324 ms

jncie@vMX6> ping routing-instance S1 rapid 192.168.201.2

PING 192.168.201.2 (192.168.201.2): 56 data bytes
!!!!!
--- 192.168.201.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 55.437/89.592/122.871/27.186 ms
```

Task 3.24: Expanding the Controllerless Overlay

- Incorporate vMX1 and vMX2 into the controllerless overlay and have the vMX routers peer with every leaf switch.
- Configure the following default-gateway on the vMX

	IPv4 Gateway	IPv6 Gateway	VLAN/VNI	Target
vMX1	192.168.220.254 (VIP)	2001:100:200::1 (VIP)	400	65534:1
vMX2	192.168.220.254 (VIP)	2001:100:200::1 (VIP)	400	65534:1

- Only import the VXLAN segment on the switches connected to S3.
- Ensure that S3 has IP connectivity to the IRB interfaces configured on the vMX devices.

Solution

vMX1:

```
set interfaces irb unit 400 family inet address 192.168.220.252/24 virtual-gateway-address
192.168.220.254

set routing-options route-distinguisher-id 172.16.1.10

set protocols ospf area 0.0.0.0 interface ae0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.16.1.10
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65250
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay neighbor 172.16.1.2
set protocols bgp group overlay neighbor 172.16.1.1
set protocols bgp group overlay neighbor 172.16.1.3
set protocols bgp group overlay neighbor 172.16.1.20

set policy-options policy-statement import-fabric-l3 term leaf from community leaf
set policy-options policy-statement import-fabric-l3 term leaf then accept
set policy-options policy-statement import-fabric-l3 term fabric-l3 from community fabric-l3
set policy-options policy-statement import-fabric-l3 term fabric-l3 then accept

set policy-options community fabric-l3 members target:65534:1

set policy-options community leaf members target:1:1

set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-import import-fabric-l3
set routing-instances fabric vrf-target target:65534:1
set routing-instances fabric protocols evpn encapsulation vxlan
set routing-instances fabric protocols evpn extended-vni-list 400
set routing-instances fabric protocols evpn multicast-mode ingress-replication
set routing-instances fabric bridge-domains bd-400 vlan-id 400
set routing-instances fabric bridge-domains bd-400 routing-interface irb.400
set routing-instances fabric bridge-domains bd-400 vxlan vni 400
```

vMX2:

```
set interfaces irb unit 400 family inet address 192.168.220.253/24 virtual-gateway-address
192.168.220.254

set routing-options route-distinguisher-id 172.16.1.20

set protocols ospf area 0.0.0.0 interface ae0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface lo0.0 passive

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.16.1.20
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65250
set protocols bgp group overlay neighbor 172.16.1.2
set protocols bgp group overlay neighbor 172.16.1.1
set protocols bgp group overlay neighbor 172.16.1.3
set protocols bgp group overlay neighbor 172.16.1.10

set policy-options policy-statement import-fabric-l3 term leaf from community leaf
set policy-options policy-statement import-fabric-l3 term leaf then accept
set policy-options policy-statement import-fabric-l3 term fabric-l3 from community fabric-l3
set policy-options policy-statement import-fabric-l3 term fabric-l3 then accept

set policy-options community fabric-l3 members target:65534:1
set policy-options community leaf members target:1:1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances fabric vtep-source-interface lo0.0
set routing-instances fabric instance-type virtual-switch
set routing-instances fabric vrf-import import-fabric-l3
set routing-instances fabric vrf-target target:65534:1
set routing-instances fabric protocols evpn encapsulation vxlan
set routing-instances fabric protocols evpn extended-vni-list 400
set routing-instances fabric protocols evpn multicast-mode ingress-replication
set routing-instances fabric bridge-domains bd-400 vlan-id 400
set routing-instances fabric bridge-domains bd-400 routing-interface irb.400
set routing-instances fabric bridge-domains bd-400 vxlan vni 400
```

vQFX1:

```
set protocols bgp group overlay neighbor 172.16.1.10
set protocols bgp group overlay neighbor 172.16.1.20

set protocols evpn extended-vni-list 400
set protocols evpn vni-options vni 400 vrf-target export target:65534:1

set policy-options policy-statement fabric-import term fabric-l3 from community fabric-l3
set policy-options policy-statement fabric-import term fabric-l3 then accept

set policy-options community fabric-l3 members target:65534:1

set switch-options vrf-import fabric-import

set vlans vlan-400 vlan-id 400
set vlans vlan-400 vxlan vni 400
```

vQFX2:

```
set protocols bgp group overlay neighbor 172.16.1.10
set protocols bgp group overlay neighbor 172.16.1.20

set protocols evpn extended-vni-list 400
set protocols evpn vni-options vni 400 vrf-target export target:65534:1

set policy-options policy-statement fabric-import term fabric-l3 from community fabric-l3
set policy-options policy-statement fabric-import term fabric-l3 then accept

set policy-options community fabric-l3 members target:65534:1

set switch-options vrf-import fabric-import

set vlans vlan-400 vlan-id 400
set vlans vlan-400 vxlan vni 400
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set protocols bgp group overlay neighbor 172.16.1.10
set protocols bgp group overlay neighbor 172.16.1.20
```

Verification

Verifying the OSPF neighbor relationship between vMX1 and vMX2:

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
192.168.0.1      ae0.0                  Full      172.16.1.20      128    35
```

Verifying the BGP sessions:

```
jncie@vMX1> show bgp summary group overlay

Groups: 2 Peers: 6 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      32         32          0          0          0          0
bgp.evpn.0
                      13         13          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.1            65250         17          8       0       0        4:02 Establ
  bgp.evpn.0: 3/3/3/0
  fabric.evpn.0: 3/3/3/0
  __default_evpn__.evpn.0: 0/0/0/0
172.16.1.2            65250         20         11       0       0        4:05 Establ
  bgp.evpn.0: 4/4/4/0
  fabric.evpn.0: 4/4/4/0
  __default_evpn__.evpn.0: 0/0/0/0
172.16.1.3            65250         16         11       0       0        4:08 Establ
  bgp.evpn.0: 0/0/0/0
  fabric.evpn.0: 0/0/0/0
  __default_evpn__.evpn.0: 0/0/0/0
172.16.1.20           65250          6          5       0       0          41 Establ
  bgp.evpn.0: 6/6/6/0
  fabric.evpn.0: 6/6/6/0
  __default_evpn__.evpn.0: 0/0/0/0
```

Check the EVPN instance on the vMX devices:

```
jncie@vMX1> show evpn instance fabric extensive

Instance: fabric
  Route Distinguisher: 172.16.1.10:8
  Encapsulation type: VXLAN
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
MAC database status                          Local  Remote
  MAC advertisements:                           1       3
  MAC+IP advertisements:                        2       2
  Default gateway MAC advertisements:           2       0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN   VNI    Status  L3 context
  irb.400                400     Up     master
Number of bridge domains: 1
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  400   400            0    0     irb.400    Extended        Enabled   400
Number of neighbors: 3
  172.16.1.1
    Received routes
      MAC address advertisement:              0
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
  172.16.1.2
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
  172.16.1.20
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           2
      Inclusive multicast:                    1
      Ethernet auto-discovery:                1
Number of ethernet segments: 2
  ESI: 00:12:12:12:12:12:12:12:12:12
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      172.16.1.2       400        0               all-active
      172.16.1.1       0          0               all-active
  ESI: 05:00:00:ff:fe:00:00:01:90:00
    Status: Resolved by IFL irb.400
    Local interface: irb.400, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.16.1.20      400        0               all-active
Router-ID: 172.16.1.10
Source VTEP interface IP: 172.16.1.10
```

Verify IP connectivity between S3 and the IRB interfaces on vMX1 and vMX2:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping routing-instance S3 rapid 192.168.220.252

PING 192.168.220.252 (192.168.220.252): 56 data bytes
!!!!!
--- 192.168.220.252 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 35.571/53.924/102.106/24.622 ms

jncie@vMX6> ping routing-instance S3 rapid 192.168.220.253

PING 192.168.220.253 (192.168.220.253): 56 data bytes
!!!!!
--- 192.168.220.253 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 43.166/47.386/50.228/2.511 ms
```

Check the MAC database on vQFX1 and vQFX2:

```
jncie@vQFX1> show ethernet-switching table vlan-id 400

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
   Vlan                 MAC                 MAC      Logical                  Active
   name                 address             flags    interface                source
   vlan-400             00:00:5e:00:01:01   DR,SD    esi.1756
05:00:00:ff:fe:00:00:01:90:00
   vlan-400             00:05:86:72:fb:f0   D        vtep.32772               172.16.1.10
   vlan-400             00:05:86:94:9b:f0   D        vtep.32771               172.16.1.20
```

Check the ARP entries on vMX router:

```
jncie@vMX1> show arp interface irb.400

MAC Address       Address         Name                      Interface               Flags
00:05:86:4d:34:c0 192.168.220.1   192.168.220.1             irb.400 [.local..8]     none
00:00:5e:00:01:01 192.168.220.254 192.168.220.254           irb.400                 permanent
published gateway
Total entries: 2
```

Task 3.25: IP-Fabric Routing Policies

- Create a policy called 'drain' on both vQFX4 as well as vQFX5. Do not apply the policy but ensure that it can be applied using only a single configuration command. The policy should ensure that, when applied, the vQFX will become less likely to receive traffic compared to other vQFX.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- vQFX3 should send traffic towards vMX2 via vQFX5 and use vQFX4 as backup. The vQFX3 leaf should continue to load share towards vMX1, vQFX1 and vQFX2.
- On vQFX3, you are only allowed to apply a policy at the BGP group level.
- Ensure that the vMX routers can use standard communities to differentiate between the routes that were received from the vQFX leaf switches when using show commands.
- Ensure IP connectivity between the servers with a layer 3 gateway on the vQFX and servers with a layer 3 gateway on the vMX.

Solution

vQFX4:

```
set groups drain protocols bgp group <*> export drain

set policy-options policy-statement drain term 1 then as-path-prepend "65500 65500"
```

vQFX5:

```
set groups drain protocols bgp group <*> export drain

set policy-options policy-statement drain term 1 then as-path-prepend "65500 65500"
```

vQFX1:

```
set policy-options policy-statement advertise-lo0 term lo0 then community add vqfx1

set policy-options policy-statement advertise-direct term xe-0/0/2 then community add vqfx1

set policy-options community qfx1 members 65000:1
set policy-options community qfx2 members 65000:2
set policy-options community qfx3 members 65000:3
```

vQFX2:

```
set policy-options policy-statement advertise-lo0 term lo0 then community add vqfx2

set policy-options community vqfx1 members 65000:1
set policy-options community vqfx2 members 65000:2
set policy-options community vqfx3 members 65000:3
```

vQFX3:

```
set policy-options policy-statement advertise-lo0 term lo0 then community add vqfx3

set policy-options policy-statement advertise-direct term xe-0/0/2 then community add vqfx3

set policy-options community vqfx1 members 65000:1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options community vqfx2 members 65000:2
set policy-options community vqfx3 members 65000:3

set protocols bgp group pod-leaf import vMX2-via-vQFX5

set policy-options policy-statement vMX2-via-vQFX5 term 1 from next-hop 192.168.6.1
set policy-options policy-statement vMX2-via-vQFX5 term 1 from route-filter 172.16.1.20/32
exact
set policy-options policy-statement vMX2-via-vQFX5 term 1 then local-preference 110
```

vMX1:

```
set protocols bgp group core-pod export advertise-direct

set policy-options prefix-list fabric-direct-irb apply-path "interfaces irb unit <*> family
inet address <*>"

set policy-options policy-statement advertise-direct term irb from protocol direct
set policy-options policy-statement advertise-direct term irb from prefix-list fabric-direct-
irb
set policy-options policy-statement advertise-direct term irb then accept

set policy-options community vqfx1 members 65000:1
set policy-options community vqfx2 members 65000:2
set policy-options community vqfx3 members 65000:3
```

vMX2:

```
set protocols bgp group core-pod export advertise-direct

set policy-options prefix-list fabric-direct-irb apply-path "interfaces irb unit <*> family
inet address <*>"

set policy-options policy-statement advertise-direct term irb from protocol direct
set policy-options policy-statement advertise-direct term irb from prefix-list fabric-direct-
irb
set policy-options policy-statement advertise-direct term irb then accept

set policy-options community vqfx1 members 65000:1
set policy-options community vqfx2 members 65000:2
set policy-options community vqfx3 members 65000:3
```

Verification

Testing the drain policy:

```
jncie@vQFX4> configure
Entering configuration mode
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4# set apply-groups drain

jncie@vQFX4# commit and-quit
configuration check succeeds
commit complete
Exiting configuration mode

jncie@vQFX4> show route advertising-protocol bgp 192.168.1.0

inet.0: 33 destinations, 34 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.2/32           Self                                    65500 65500 [65500] 65000 I
* 172.16.1.3/32           Self                                    65500 65500 [65500] 65000 I
* 172.16.1.4/32           Self                                    I
* 172.16.1.10/32          Self                                    65500 65500 [65500] 65534 I
* 172.16.1.20/32          Self                                    65500 65500 [65500] 65534 I
* 192.168.130.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.131.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.132.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.133.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.134.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.135.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.220.0/24        Self                                    65500 65500 [65500] 65534 I

jncie@vQFX4> show route advertising-protocol bgp 192.168.7.0

inet.0: 33 destinations, 34 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.1/32           Self                                    65500 65500 [65500] 65000 I
* 172.16.1.2/32           Self                                    65500 65500 [65500] 65000 I
* 172.16.1.3/32           Self                                    65500 65500 [65500] 65000 I
* 172.16.1.4/32           Self                                    I
* 192.168.110.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.111.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.112.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.113.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.114.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.115.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.130.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.131.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.132.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.133.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.134.0/24        Self                                    65500 65500 [65500] 65000 I
* 192.168.135.0/24        Self                                    65500 65500 [65500] 65000 I

jncie@vQFX4> configure
Entering configuration mode

jncie@vQFX4# rollback 1
load complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4# commit and-quit
configuration check succeeds
commit complete
Exiting configuration mode
```

Verifying the vQFX3 policy:

```
jncie@vQFX3> show policy vMX2-via-vQFX5

Policy vMX2-via-vQFX5:
    Term 1:
        from nexthop 192.168.6.1
        route filter:
            172.16.1.20/32 exact
        then localpref 110

{master:0}
jncie@vQFX3> show route 172.16.1.20

inet.0: 34 destinations, 45 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.20/32     *[BGP/170] 01:22:18, localpref 110
                      AS path: 65500 65534 I, validation-state: unverified
                    > to 192.168.6.1 via xe-0/0/6.6
                    [BGP/170] 00:05:46, localpref 100
                      AS path: 65500 65534 I, validation-state: unverified
                    > to 192.168.5.1 via xe-0/0/6.5
```

Verifying the communities on the vMX:

```
jncie@vMX1> show route 172.16.1.0/24 community-name vqfx1

inet.0: 30 destinations, 45 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32      *[BGP/170] 00:07:41, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
                      to 192.168.7.1 via xe-0/0/6.7
                    > to 192.168.9.1 via xe-0/0/6.9
                    [BGP/170] 00:07:41, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
                    > to 192.168.7.1 via xe-0/0/6.7

jncie@vMX1> show route 172.16.1.0/24 community-name vqfx2

inet.0: 30 destinations, 45 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.2/32      *[BGP/170] 00:07:43, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      to 192.168.7.1 via xe-0/0/6.7
                    > to 192.168.9.1 via xe-0/0/6.9
                    [BGP/170] 00:07:43, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
                    > to 192.168.7.1 via xe-0/0/6.7

jncie@vMX1> show route 172.16.1.0/24 community-name vqfx3

inet.0: 30 destinations, 45 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.3/32       *[BGP/170] 00:07:44, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
                      to 192.168.7.1 via xe-0/0/6.7
                    > to 192.168.9.1 via xe-0/0/6.9
                    [BGP/170] 00:07:44, localpref 100
                      AS path: 65500 65000 I, validation-state: unverified
                    > to 192.168.7.1 via xe-0/0/6.7
```

Appendix – Chapter 4: Data Center Interconnect

Part 1: MPLS and MPLS L3VPN

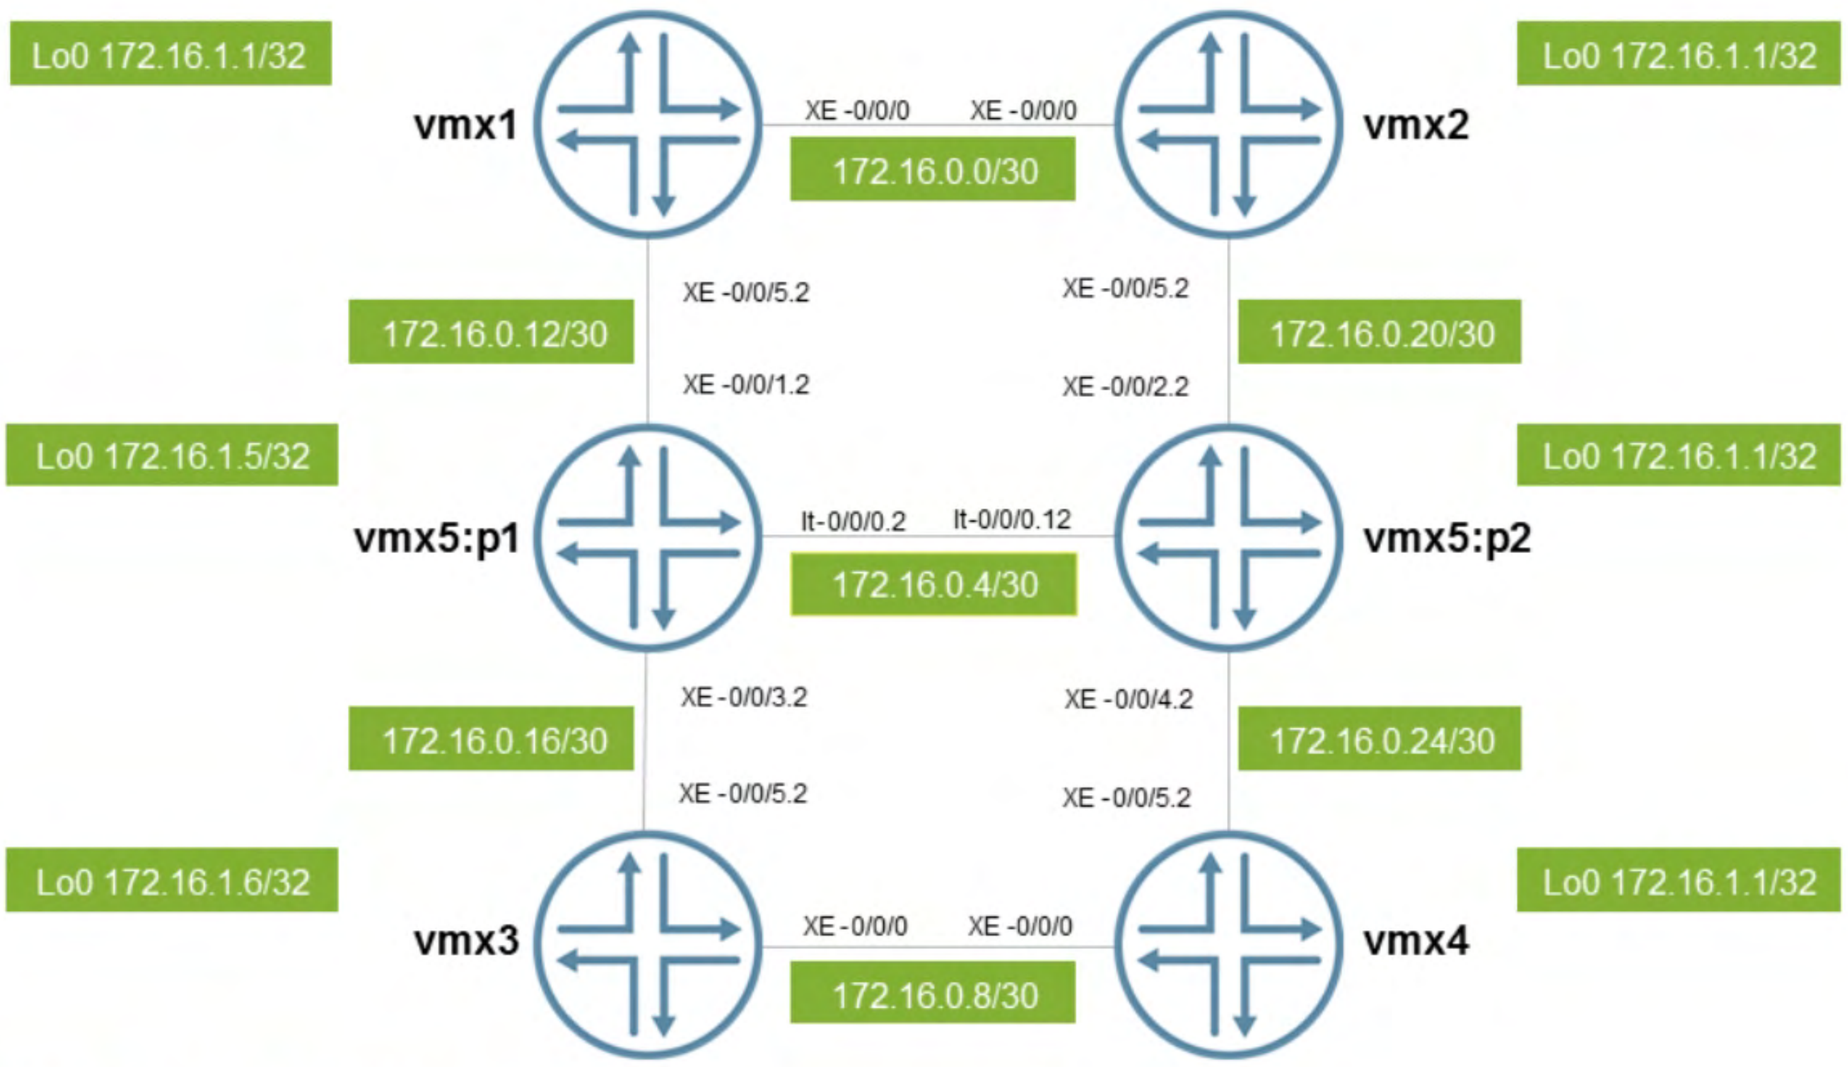

Figure 22 – MPLS network topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

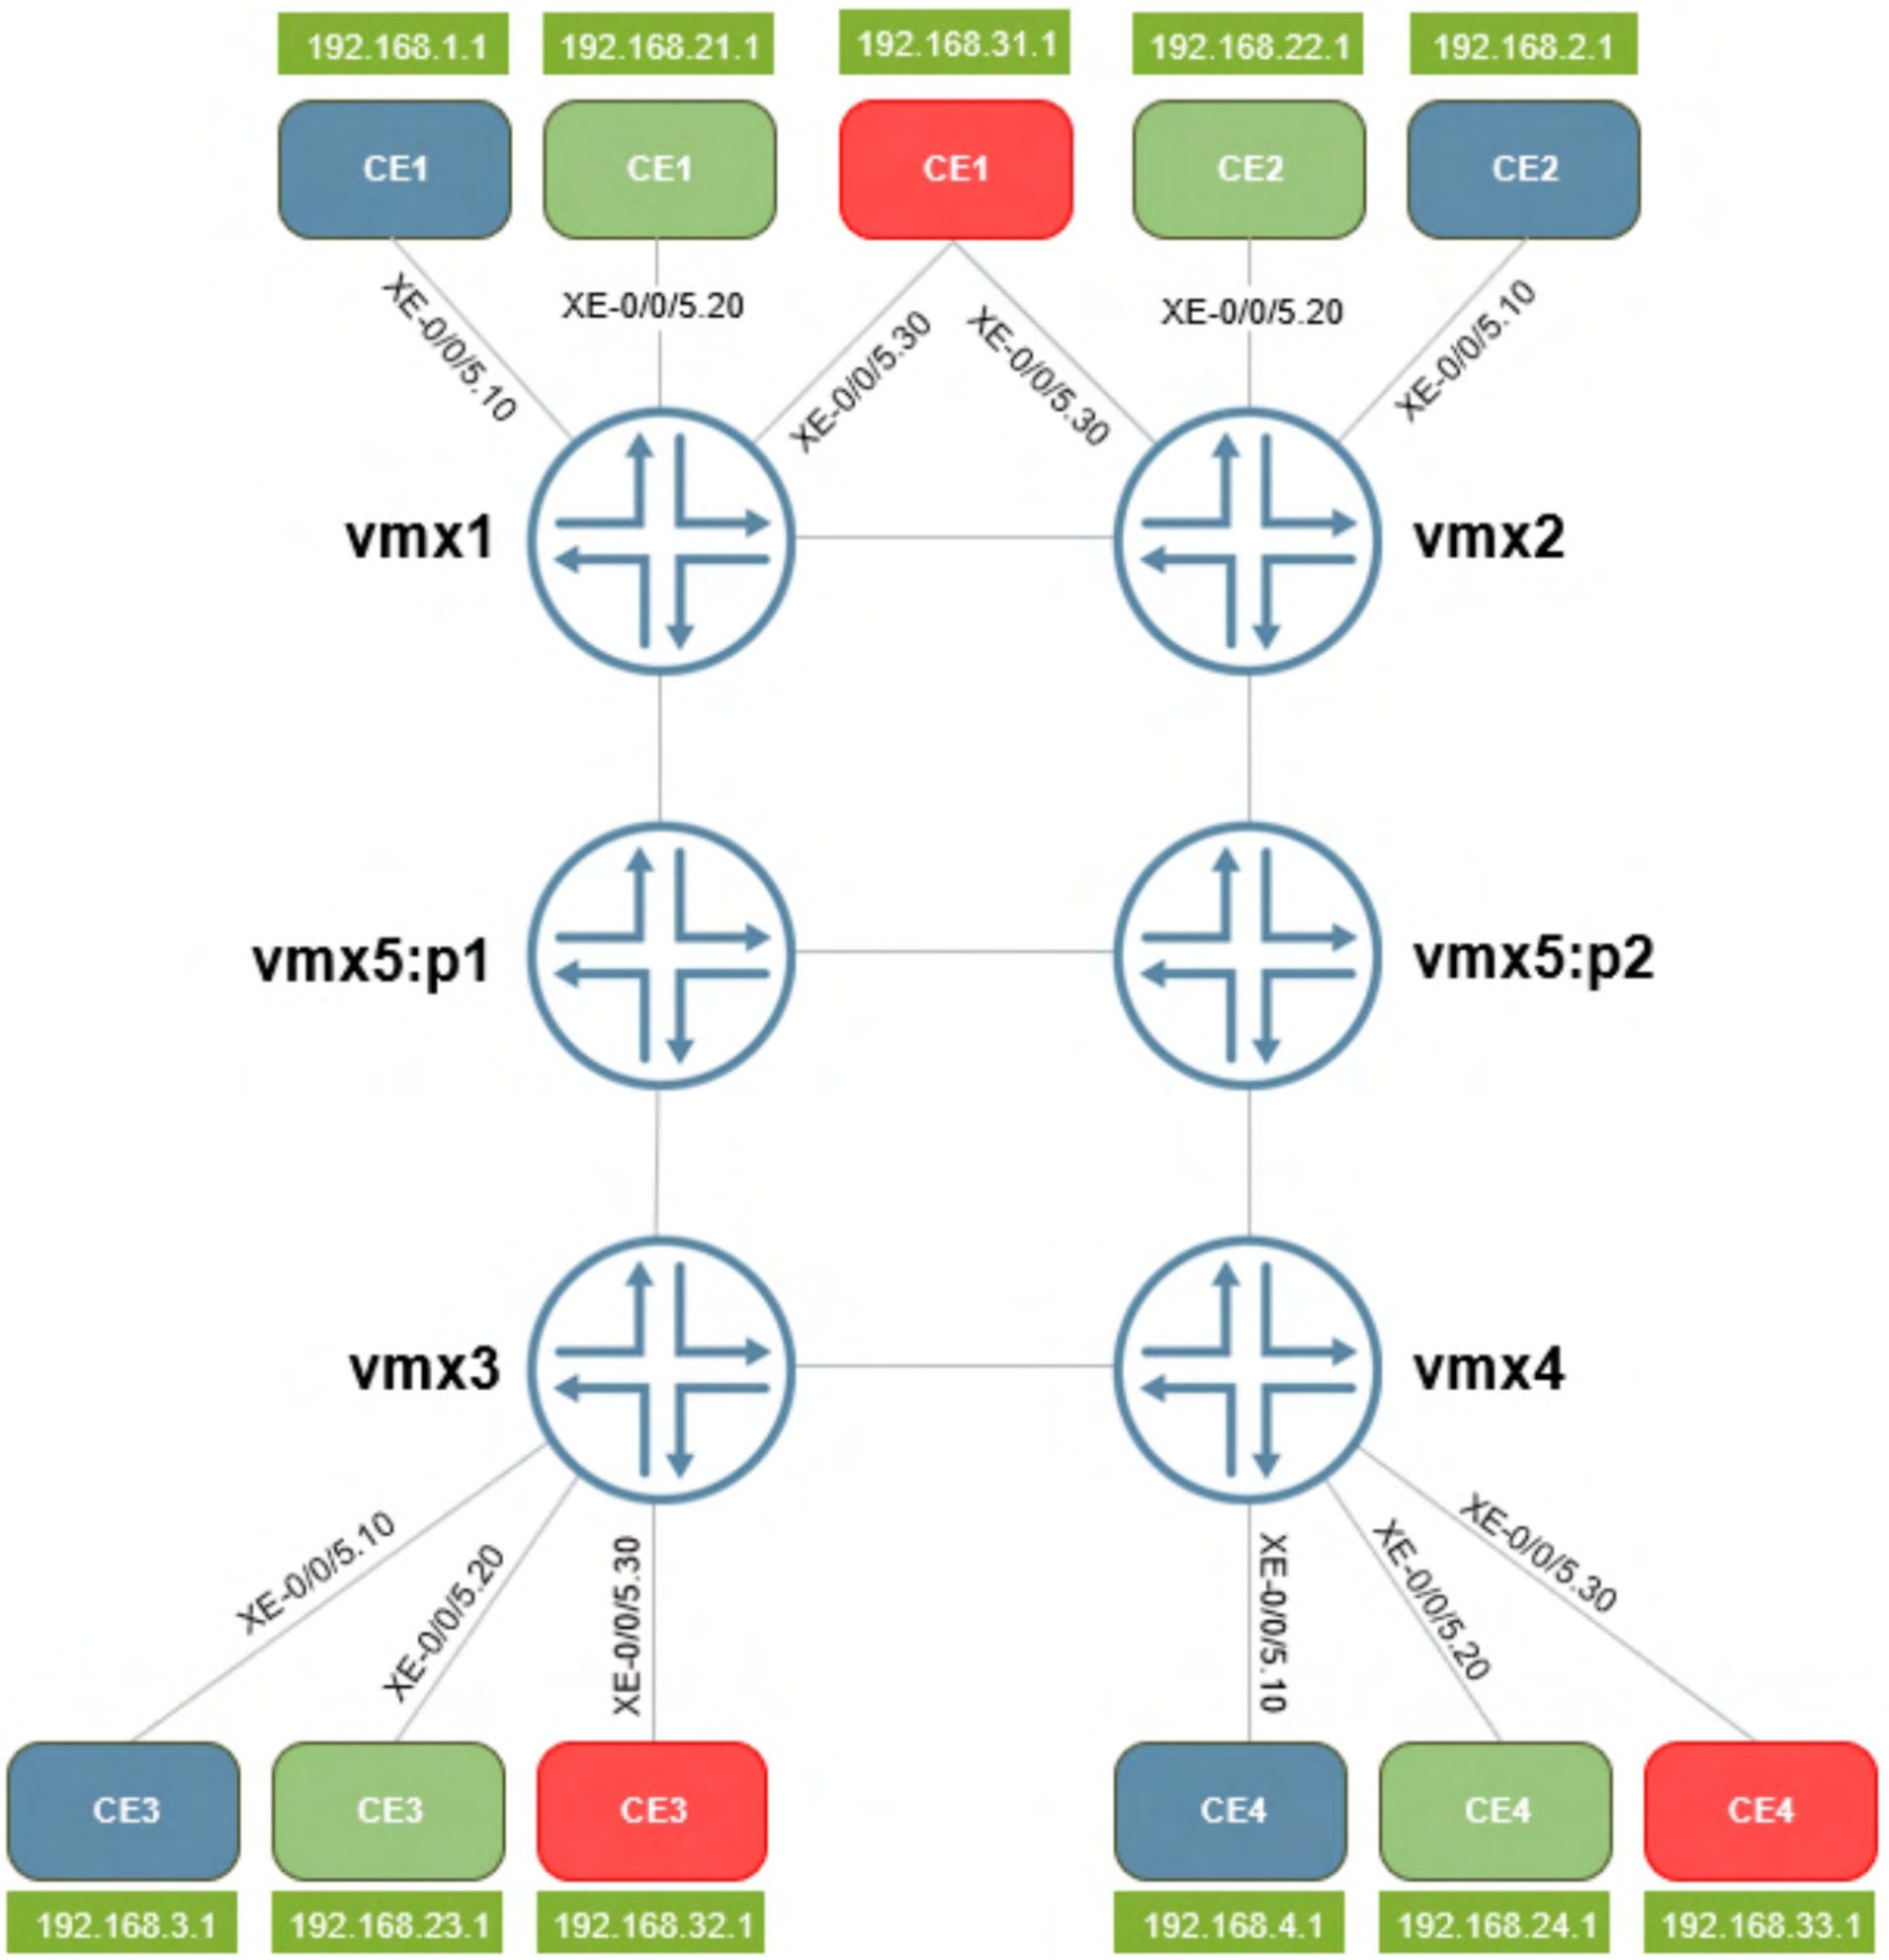

Figure 2 - VPN topology

Notes:

- Starting topology is preconfigured with IP-addressing only.
- On vMX5, three logical systems are configured. These are P1, P2 and CE

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.1: IGP Configuration

- Configure all vMX routers to be part of a single area OSPF network.
- Authenticate the OSPF neighbor relationship within the MD5 password jncie'
- Make sure OSPF will calculate path cost using 100Gbit as a reference.
- Advertise the loopback IP address into OSPF on every vMX router.
- Make sure the OSPF database only contains router LSAs.

Solution

vMX1:

```
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 authentication md5 1 key jncie
```

vMX2:

```
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 authentication md5 1 key jncie
```

vMX3:

```
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 authentication md5 1 key jncie
```

vMX4:

```
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 authentication md5 1 key jncie
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX5:P1:

```
edit logical-systems P1
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/1.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1.2 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/3.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/3.2 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface lt-0/0/0.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface lt-0/0/0.2 authentication md5 1 key jncie
```

vMX5:P2:

```
edit logical-systems P2
set protocols ospf reference-bandwidth 100g
set protocols ospf area 0.0.0.0 interface lo0.100 passive
set protocols ospf area 0.0.0.0 interface xe-0/0/2.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/2.2 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface xe-0/0/4.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/4.2 authentication md5 1 key jncie
set protocols ospf area 0.0.0.0 interface lt-0/0/0.12 interface-type p2p
set protocols ospf area 0.0.0.0 interface lt-0/0/0.12 authentication md5 1 key jncie
```

Verification

The verification commands for this task are the same on every router. In this example, we verify the IGP by examining command output on vMX1. First, we check the OSPF neighbor adjacencies:

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
172.16.0.2       xe-0/0/0.0             Full      172.16.1.2       128    37
172.16.0.14      xe-0/0/5.2             Full      172.16.1.5       128    32
```

After this, we check the OSPF interface type, the authentication type and the link cost:

```
jncie@vMX1> show ospf interface detail

Interface          State   Area            DR ID           BDR ID          Nbrs
lo0.0              DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 172.16.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  Adj count: 0, Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 0
xe-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 172.16.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 10
xe-0/0/5.2          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0             1
  Type: P2P, Address: 172.16.0.13, Mask: 255.255.255.252, MTU: 1500, Cost: 10
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 10
```

The cost for the 'xe-' links is 10. Without the reference bandwidth configuration command, Junos OS would have used the default reference bandwidth of 100Mbps. This means that all 100Mbps links, and all links with a bandwidth greater than 100Mbps, would have a cost of 1. The following output would have been shown without the reference bandwidth configuration option:

```
...
xe-0/0/5.2          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0             1
  Type: P2P, Address: 172.16.0.13, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 1

...
```

Another objective was to make sure that the OSPF area only contains Router LSAs. This is accomplished by configuring all links inside the OSPF area as point-to-point links. To verify whether or not this objective has been met, we check the OSPF database using the following command:

```
jncie@vMX1> show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1      0x80000006    62  0x22 0x4ed9  84
Router   172.16.1.2       172.16.1.2      0x80000005    61  0x22 0xa64a  84
Router   172.16.1.3       172.16.1.3      0x80000003   223  0x22 0x1e4   84
Router   172.16.1.4       172.16.1.4      0x80000003   232  0x22 0x975f  84
Router   172.16.1.5       172.16.1.5      0x80000004   222  0x22 0xb2c5 108
Router   172.16.1.6       172.16.1.6      0x80000004   222  0x22 0xba6b  96
```

After finishing task 1, the routing table should display OSPF routes towards the loopback IP addresses of the other routers:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show route protocol ospf

inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

172.16.0.4/30      *[OSPF/10] 00:07:18, metric 2
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.0.8/30      *[OSPF/10] 00:07:18, metric 13
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.0.16/30     *[OSPF/10] 00:07:18, metric 11
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.0.20/30     *[OSPF/10] 00:06:42, metric 11
                      > to 172.16.0.2 via xe-0/0/0.0
172.16.0.24/30     *[OSPF/10] 00:07:18, metric 12
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.1.2/32      *[OSPF/10] 00:06:42, metric 1
                      > to 172.16.0.2 via xe-0/0/0.0
172.16.1.3/32      *[OSPF/10] 00:07:18, metric 11
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.1.4/32      *[OSPF/10] 00:07:18, metric 12
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.1.5/32      *[OSPF/10] 00:07:18, metric 1
                      > to 172.16.0.14 via xe-0/0/5.2
172.16.1.6/32      *[OSPF/10] 00:00:01, metric 2
                      > to 172.16.0.14 via xe-0/0/5.2
224.0.0.5/32       *[OSPF/10] 00:10:57, metric 1
                         MultiRecv
```

Task 4.2: LDP Configuration

- Enable LDP neighbor relationships between the vMX routers.
- Authenticate LDP sessions using the MD5 hash ‘jncie’.
- LDP should use the same metric as the IGP.
- The IGP should prefer links with a working LDP session.

Solution

vMX1:

For LDP to run, the interfaces need to be enabled for MPLS. To enable interfaces for MPLS on Junos OS, we first configure the interfaces with the 'family mpls' statement:

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls
```

After this, we add the interfaces to the [protocols mpls] hierarchy:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols mpls interface xe-0/0/5.2
set protocols mpls interface xe-0/0/0.0
```

After making sure the interfaces are MPLS enabled, we are ready to proceed with the LDP configuration. First, we configure LDP to run on the MPLS-enabled interfaces:

```
set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/5.2
```

After this, we make sure that LDP will use the same metric as the IGP. The Junos OS default is to install every LDP route with a metric of 1. We can make LDP use the same metric as the IGP by issuing the following configuration command:

```
set protocols ldp track-igp-metric
```

The final LDP configuration task is to enable authentication for the LDP session. In older versions of Junos OS, authentication was configured inside the LDP session hierarchy. On newer Junos OS releases, we can configure authentication for the subnet in which we expect to establish the LDP sessions:

```
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie
```

To ensure that the IGP will always prefer links with a working LDP session, we turn to 'ldp-synchronization'. This will make the IGP automatically increase the link cost for interfaces when the LDP session is down. The feature is activated using the following configuration commands:

```
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 ldp-synchronization
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 ldp-synchronization
```

vMX2:

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2

set protocols ldp track-igp-metric
set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/5.2
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 ldp-synchronization
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 ldp-synchronization
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2

set protocols ldp track-igp-metric
set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/5.2
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 ldp-synchronization
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 ldp-synchronization
```

vMX4:

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2

set protocols ldp track-igp-metric
set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/5.2
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 ldp-synchronization
set protocols ospf area 0.0.0.0 interface xe-0/0/5.2 ldp-synchronization
```

vMX5:P1:

```
edit logical-systems P1
set interfaces lt-0/0/0 unit 2 family mpls
set interfaces xe-0/0/1 unit 2 family mpls
set interfaces xe-0/0/3 unit 2 family mpls

set protocols mpls interface xe-0/0/1.2
set protocols mpls interface xe-0/0/3.2
set protocols mpls interface lt-0/0/0.2

set protocols ldp track-igp-metric
set protocols ldp interface lt-0/0/0.2
set protocols ldp interface xe-0/0/1.2
set protocols ldp interface xe-0/0/3.2
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/1.2 ldp-synchronization
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols ospf area 0.0.0.0 interface xe-0/0/3.2 ldp-synchronization
set protocols ospf area 0.0.0.0 interface lt-0/0/0.2 ldp-synchronization
```

vMX5:P2:

```
edit logical-systems P2
set interfaces lt-0/0/0 unit 12 family mpls
set interfaces xe-0/0/2 unit 2 family mpls
set interfaces xe-0/0/4 unit 2 family mpls

set protocols mpls interface xe-0/0/2.2
set protocols mpls interface xe-0/0/4.2
set protocols mpls interface lt-0/0/0.12

set protocols ldp track-igp-metric
set protocols ldp interface lt-0/0/0.12
set protocols ldp interface xe-0/0/2.2
set protocols ldp interface xe-0/0/4.2
set protocols ldp session-group 172.16.1.0/24 authentication-key jncie

set protocols ospf area 0.0.0.0 interface xe-0/0/2.2 ldp-synchronization
set protocols ospf area 0.0.0.0 interface xe-0/0/4.2 ldp-synchronization
set protocols ospf area 0.0.0.0 interface lt-0/0/0.12 ldp-synchronization
```

Verification

Since LDP requires MPLS-enabled interface, we first check what interfaces have been enabled for MPLS:

```
jncie@vMX1> show mpls interface

Interface           State       Administrative groups (x: extended)
xe-0/0/0.0          Up          <none>
xe-0/0/5.2          Up          <none>
```

Those two interfaces should also be running LDP:

```
jncie@vMX1> show ldp interface

Interface          Address                          Label space ID    Nbr    Next
                                                                     count  hello
xe-0/0/0.0         172.16.0.1                       172.16.1.1:0        1      2
xe-0/0/5.2         172.16.0.13                      172.16.1.1:0        1      3
```

The output shows that vMX1 has discovered a neighbor on both links. This LDP neighbor relationship is discovered through the use of link local multicast messages. As soon as two neighbors have discovered each other as neighbors, they can start to form an LDP session. The LDP session is formed using a TCP connection. In our lab, the vMX1 router

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

is connected to vMX2 and P1. When the other routers are configured correctly, the vMX1 router should have been able to establish two LDP sessions:

```
jncie@vMX1> show ldp session

  Address                         State       Connection  Hold time  Adv. Mode
172.16.1.2                        Operational Open          20         DU
172.16.1.5                        Operational Open          22         DU
```

The output above shows us the vMX1 has established an LDP session with vMX2 and P1. The 'State' refers to the state that LDP is in. In this example, the 'Operational' state indicates LDP is working. The 'Connection' refers to the TCP connection state.

To verify that LDP sessions are authenticated:

```
jncie@vMX1> show ldp session detail

Address: 172.16.1.2, State: Operational, Connection: Open, Hold time: 23
  Session ID: 172.16.1.1:0--172.16.1.2:0
...
  Authentication type: MD5 (172.16.1.0/24)
...
```

Verifying the LDP synchronization feature can be done by examining the OSPF interface details:

```
jncie@vMX1> show ospf interface extensive

Interface          State   Area            DR ID           BDR ID          Nbrs
lo0.0              DRother 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 172.16.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  Adj count: 0, Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 0
xe-0/0/0.0           PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 172.16.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 1
  LDP sync state: in sync, for: 00:06:26, reason: LDP session up
           config holdtime: infinity
xe-0/0/5.2           PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 172.16.0.13, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 1
  LDP sync state: in sync, for: 00:00:16, reason: LDP session up
            config holdtime: infinity
```

To see how the feature works, we can temporarily disable LDP on vMX2:

```
jncie@vMX2> configure
Entering configuration mode

[edit]
jncie@vMX2# deactivate protocols ldp

[edit]
jncie@vMX2# commit
commit complete
```

We then move back to vMX1. When we examine the OSPF interface again, we can see that the 'LDP sync state' has changed and that the link cost was increased:

```
jncie@vMX1# show ospf interface extensive
...
xe-0/0/0.0           PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
  Type: P2P, Address: 172.16.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 65535
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 00:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 10
  LDP sync state: in hold-down, for: 00:00:02, reason: LDP session down
            config holdtime: infinity
```

Remember to enable LDP on vMX2 before moving on using rollback 1.

By default, Juniper routers enabled for LDP will only advertise a label mapping for the primary loopback interface. In addition to this, every router will also advertise a label mapping for every received route. The end result is a full-mesh of LDP-signaled LSPs between all of the routers.

Since we enabled all of the vMX routers for LDP, we should see 5 LDP routes (or LSPs) towards the 5 other routers in our lab:

```
jncie@vMX1> show route 172.16.1.0/24

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.1.1/32      *[Direct/0] 02:36:02
                     > via lo0.0
172.16.1.2/32      *[OSPF/10] 00:09:25, metric 1
                     > to 172.16.0.2 via xe-0/0/0.0
172.16.1.3/32      *[OSPF/10] 00:03:15, metric 11
                     > to 172.16.0.14 via xe-0/0/5.2
172.16.1.4/32      *[OSPF/10] 00:03:15, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2
172.16.1.5/32      *[OSPF/10] 00:03:15, metric 1
                     > to 172.16.0.14 via xe-0/0/5.2
172.16.1.6/32      *[OSPF/10] 00:03:15, metric 2
                     > to 172.16.0.14 via xe-0/0/5.2
```

```
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.2/32      *[LDP/9] 00:03:29, metric 1
                     > to 172.16.0.2 via xe-0/0/0.0
172.16.1.3/32      *[LDP/9] 00:03:15, metric 11
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300560
172.16.1.4/32      *[LDP/9] 00:03:15, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300576
172.16.1.5/32      *[LDP/9] 00:03:15, metric 1
                     > to 172.16.0.14 via xe-0/0/5.2
172.16.1.6/32      *[LDP/9] 00:03:15, metric 2
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300544
```

By default, all LDP (and RSVP) routes/LSPs are placed into the inet.3 table. The inet.3 table is used to store the egress addresses for all IPv4 MPLS LSPs. After examining the vMX1 routing table, we have verified that it has an LSP towards all the other routers. Remember that LSPs are unidirectional. The fact that the vMX1 router has LSPs towards all the other routers does not mean those other routers have an LSP towards the vMX1 router. Thorough verification requires you to check the routing table on every vMX.

After completing this task, the routers in the lab have been able to set up a full-mesh of LSPs. To better understand the way LDP works, let's follow a unidirectional LSP from vMX1 to vMX4.

Note: every time a new lab is started, LDP labels with different values are generated by the vMX routers. This means that if you follow along this walkthrough, the label values you see will differ from the values displayed here.

We start out examining the route towards vMX4 on vMX1:

```
jncie@vMX1> show route 172.16.1.4

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.1.4/32      *[OSPF/10] 00:06:14, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      *[LDP/9] 00:06:14, metric 12
                    > to 172.16.0.14 via xe-0/0/5.2, Push 300576
```

The IPv4 LSP towards the vMX4 router is stored in the inet.3 table. Traffic that is sent onto the LSP will exit interface 'xe-0/0/5.2' The label that vMX1 will use to reach vMX4 is '300576'. This means that if the vMX1 router wants to send traffic onto this LSP, it will push a label with this value onto the packet right before the IP header and switch the MPLS packet out interface 'xe-0/0/5.2'.

There are multiple ways to find out what the next router is. One of the possibilities is to examine the network diagram. Another quick way is to simply check what OSPF neighbor is found at the other end of the link:

```
jncie@vMX1> show ospf neighbor | match xe-0/0/5.2

172.16.0.14      xe-0/0/5.2            Full      172.16.1.5       128    36
```

The router ID of the OSPF neighbor that is active on this link corresponds to P1. So, we can conclude that the LDP-signaled LSP from vMX1 to vMX4 traverses P1. To see how P1 will handle this traffic, we move to the P1 router and issue the following command:

```
jncie@vMX5:P1> show route label 300576

mpls.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

300576             *[LDP/9] 00:00:19, metric 11
                    > to 172.16.0.6 via lt-0/0/0.2, Swap 299808
```

The P1 router is a transit router for the LSP from vMX1 to vMX4. P1 has a routing entry inside the 'mpls.0' table that tells us how it will handle traffic labeled 300576. The P1 router will swap the label with another label, 299808. This label value was learned from its downstream neighbor. The output also tells us via what interface P1 will switch the traffic. We can find out what router that is by checking the OSPF neighbor:

```
jncie@vMX5:P1> show ospf neighbor | match lt-0/0/0.2

172.16.0.6       lt-0/0/0.2            Full      172.16.1.6       128    31
```

The next router in the path is P2. Since P1 swapped the label again, on P2, we examine the entry for label 299808:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX5:P2> show route label 299808

mpls.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299808             *[LDP/9] 00:02:14, metric 10
                    > to 172.16.0.25 via xe-0/0/4.2, Pop
299808(S=0)        *[LDP/9] 00:02:14, metric 10
                    > to 172.16.0.25 via xe-0/0/4.2, Pop
```

The P2 router is a transit router in the LSP from vMX1 to vMX4, and by examining the command output, we can tell that this router is the second to last router for this LSP. This is because instead of swapping the label again, the P2 router will perform a pop operation. This means it will remove the MPLS label from the packet before switching it out the next interface.

The output also tells us the egress interface for traffic on the LSP is the xe-0/0/4.2 interface. Since this is the second to last router for the LSP, we should see the vMX4 router as an OSPF neighbor across this interface:

```
jncie@vMX5:P2> show ospf neighbor | match xe-0/0/4.2

172.16.0.25       xe-0/0/4.2            Full       172.16.1.4       128    37
```

If the LSP would be used for regular IPv4 traffic, the vMX4 router will receive an IP packet without any labels when vMX1 sends traffic onto this LDP signaled LSP.

Task 4.3: MPLS LSP Options Configuration

- Enable all protocols to use LDP signaled LSPs for forwarding.
- Tunnel ICMP messages inside MPLS.
- The MPLS network should be represented as a single hop to future VPN customers.

Solution

vMX1-4:

Enabling the vMX routers to forward all traffic using the LSPs:

```
set protocols mpls traffic-engineering mpls-forwarding
```

To hide the MPLS network from customers, we disable TTL decrementing on every router:

```
set protocols mpls no-propagate-ttl
```

To enable transit routers in our BGP free core to use LSPs when sending ICMP:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols mpls icmp-tunneling
```

vMX5:P1:

```
edit logical-systems P1
set protocols mpls traffic-engineering mpls-forwarding
set protocols mpls no-propagate-ttl
set protocols mpls icmp-tunneling
```

vMX5:P2:

```
edit logical-systems P2
set protocols mpls traffic-engineering mpls-forwarding
set protocols mpls no-propagate-ttl
set protocols mpls icmp-tunneling
```

Verification

The default Junos OS behavior is to install LSP next-hops into the inet.3 table. BGP can use this table to resolve a next-hop and have the router forward traffic towards a destination using the LSP. Under the [**protocols mpls traffic-engineering**] hierarchy, we find several knobs to alter this behavior. The configuration options we have are the following:

- traffic-engineering bgp-igp
    - move inet.3 routes to the main routing-table (inet.0)
- traffic-engineering bgp-igp-both-ribs
    - copy inet.3 routes to inet.0 (VPNs require a route in the inet.3 table)
- traffic-engineering mpls-forwarding
    - when routes are moved or copied to the inet.0 table, the MPLS routes can supersede the IGP routes and cause for a router to stop redistributing prefixes
    - with 'mpls-forwarding', LSPs are used to forward traffic but they are excluded from route selection
    - routes are added to the inet.0 and the inet.3 table
- traffic-engineering bgp
    - default, MPLS performs traffic engineering on BGP destinations only ingress routes are installed in the inet.3 table

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

For this task, we need to use the LSPs for forwarding only. For this reason, we configured the 'mpls-forwarding' option. We can verify our configuration by examining the routing table.

First, we check the routing table before configuring 'traffic-engineering mpls-forwarding':

```
jncie@vMX1> show route 172.16.1.4

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      *[OSPF/10] 00:00:01, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
172.16.1.4/32      *[LDP/9] 00:00:01, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300672
```

And now, after configuring 'traffic-engineering mpls-forwarding', we check the routing table again:

```
jncie@vMX1> show route 172.16.1.4

inet.0: 25 destinations, 30 routes (25 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      @[OSPF/10] 00:01:41, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2
                   #[LDP/9] 00:01:07, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300672

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      *[LDP/9] 00:01:07, metric 12
                     > to 172.16.0.14 via xe-0/0/5.2, Push 300672
```

When we check the routing entry in the inet.0 table, we can see the '@' character precede the OSPF route and we can see the '#' character precede the LDP route. This means the OSPF route is only used for routing and the LDP route is only used for forwarding.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Under the [protocols mpls] configuration hierarchy, we have two different options to hide the network topology from our customers. These two options are the following:

- no-decrement-ttl
    - no-decrement-ttl is proprietary and works for RSVP signaled LSPs only
    - existing LSPs need to be cleared
    - per-LSP or for all LSPs
    - configured on ingress router
- no-propagate-ttl
    - disables normal TTL decrementing on LSR
    - effective immediately, existing LSPs do not need to be cleared
    - needs to be configured on all routers
    - works for LDP and RSVP

Since we are using LDP in this lab, we configure the 'no-propagate-ttl' option. To better understand this option, lets first try to understand the default MPLS TTL handling on Juniper devices.

An ingress LSR will copy the value of the IP TTL, minus 1, to the MPLS header. The TTL value inside the IP header is left intact. Every LSR along the path will decrement the MPLS TTL by one. The penultimate router along the LSP will pop the label. When this router pops the label, it will also write the MPLS TTL to the IP TTL. Before actually sending the IP packet onto the last router, the penultimate router decrements the IP TTL by 1. In this way, every router along the path of the LSP will have decremented the TTL by 1.

When we apply the ‘no-propagate-ttl’ configuration command, TTL handling changes. The IP TTL is no longer copied to the MPLS header by the ingress LSR. Instead, the ingress LSR decrements the IP TTL by 1 and sets the TTL of the MPLS header to 255.

Every LSR along the path will decrement the MPLS TTL and leaves the IP TTL intact.

When the penultimate router along the path of the LSP pops the label, it will not write the MPLS TTL back to the IP TTL. Instead, the egress LSR of the LSP will decrement the IP TTL by 1 before sending the packet out of its egress interface.

In short, IP packets transiting the MPLS network will have their TTL decremented by 2. Once by the ingress LSR and a second time by the egress LSR. This effectively hides the internal MPLS topology to the outside world.

We can verify our configuration by issuing a traceroute. First, we issue a traceroute before applying the 'no-propagate-ttl' configuration on every router:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1# traceroute 172.16.1.4 source 172.16.1.1

traceroute to 172.16.1.4 (172.16.1.4) from 172.16.1.1, 30 hops max, 52 byte packets
 1  172.16.0.14 (172.16.0.14)  189.084 ms  1044.605 ms  1791.002 ms
 2  172.16.0.6 (172.16.0.6)  289.598 ms  3.652 ms  5.811 ms
 3  172.16.1.4 (172.16.1.4)  573.530 ms  357.596 ms  424.758 ms
```

This traceroute reveals every hop in the path. What happens here is that the ingress router for the LSP copies the TTL from the IP header into the MPLS header. Every router along the path will decrement the TTL by 1 and the penultimate hop will rewrite the MPLS TTL into the IP header.

Next, we configure 'no-propagate-ttl' and repeat the same traceroute:

```
jncie@vMX1> traceroute 172.16.1.4 source 172.16.1.1

traceroute to 172.16.1.4 (172.16.1.4) from 172.16.1.1, 30 hops max, 52 byte packets
 1  172.16.1.4 (172.16.1.4)  113.267 ms  60.198 ms  9.085 ms
```

We have effectively 'hidden' the MPLS network. When traffic is sent onto an LSP, the TTL from the IP header is no longer copied into the MPLS header. Instead, the ingress LSR will set the MPLS TTL to 255. Another thing that happens is that the penultimate router no longer writes the MPLS TTL into the IP header. Instead, it pops the label and simply forwards the IP packet with the same TTL it had when it entered the LSP.

Task 4.4: BGP Configuration

- Enable a BGP full-mesh between vMX1, vMX2, vMX3 and vMX4
- Authenticate all BGP sessions using the password ‘jncie’.
- The BGP session should allow for the exchange of MPLS L3 VPN addresses only.
- Use autonomous system 65000.

Solution

vMX1:

We are preparing the network to support MPLS L3VPN. We configure a full-mesh of BGP sessions and make sure that the BGP sessions supports the 'inet-vpn unicast' family:

```
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.1
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.3
set protocols bgp group ibgp neighbor 172.16.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-options autonomous-system 65000
```

vMX2:

```
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.2
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.1
set protocols bgp group ibgp neighbor 172.16.1.3
set protocols bgp group ibgp neighbor 172.16.1.4
set routing-options autonomous-system 65000
```

vMX3:

```
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.3
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.4
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.1

set routing-options autonomous-system 65000
```

vMX4:

```
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.3
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.1

set routing-options autonomous-system 65000
```

Verification

The easiest, and quickest, way to spot whether or not the BGP sessions successfully established with the proper address family is by issuing the following command:

```
jncie@vMX1> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       0          0          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.2            65000          6          5       0       0        1:44 Establ
  bgp.l3vpn.0: 0/0/0/0
172.16.1.3            65000          6          5       0       0        1:27 Establ
  bgp.l3vpn.0: 0/0/0/0
172.16.1.4            65000          5          4       0       0        1:19 Establ
  bgp.l3vpn.0: 0/0/0/0
```

More detailed output that can be useful for troubleshooting is gathered using the following command:

```
jncie@vMX1> show bgp neighbor 172.16.1.2

Peer: 172.16.1.2+52835 AS 65000 Local: 172.16.1.1+179 AS 65000
  Group: ibgp                    Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AuthKey AddressFamily Rib-group Refresh>
  Authentication key is configured
  Address families configured: inet-vpn-unicast
  Local Address: 172.16.1.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.16.1.2      Local ID: 172.16.1.1        Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65000)
  Peer does not support Addpath
  Table bgp.l3vpn.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
  Last traffic (seconds): Received 10684 Sent 134 Checked 10684
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Input messages:  Total 7      Updates 1       Refreshes 0     Octets 188
Output messages: Total 6      Updates 0       Refreshes 0     Octets 125
```

The lab now has a full mesh of LSPs to forward MPLS L3VPN traffic. It also has a full mesh of BGP sessions that are capable of exchanging the addresses for MPLS L3VPNs. No routes are exchanged yet because no MPLS L3VPNs have been configured.

In this scenario, we have a so-called BGP free core. During the configuration and verification of future tasks, we will learn that the P1 and P2 routers can forward the MPLS L3VPN traffic when though they have no knowledge of the VPNs whatsoever. They will only forward MPLS traffic.

Task 4.5: MPLS L3VPN JNCIE-BLUE

- PE-CE interfaces configured with VLAN-ID 10 should be placed into a VPN called 'JNCIE-BLUE'.

- Enable the use of the following subnets on CEs located in the MPLS L3VPN:

CE device	**Subnet**
CE1	192.168.1.0/24
CE2	192.168.2.0/24
CE3	192.168.3.0/24
CE4	192.168.4.0/24

- Do not use any dynamic routing protocols between the PE and the CE device. You are allowed to use static routes.
- The PE interface routes should be reachable for remote sites.
- This VPN, as well as future VPNs, should be active on vMX1, vMX2, vMX3 and vMX4.
- Choose any route-target you want.

Solution

vMX1:

As an alternative to configuring every MPLS L3VPN with a unique a Route-Distinguisher (RD) value, we configure a 'route-distinguisher-id' on every PE. This will make the PE automatically assign an RD when we configure a routing-instance, simplifying the configuration of VPNs. To have PEs generate unique RD, we use the loopback IP as route-distinguisher-id value :

```
set routing-options route-distinguisher-id 172.16.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

All further MPLS L3VPN configuration takes place in the routing-instances hierarchy. First, we specify what type of routing-instance we are configuring. There are over ten different types of routing-instances supported on an MX. We configure our MPLS L3VPN instances as a VRF (VPN Routing and Forwarding):

```
set routing-instances JNCIE-BLUE instance-type vrf
```

Note: you may have used the instance-type virtual-router before in the past. The virtual-router can be thought of as a VRF-lite routing-instance. It is configured without RD or route-target (RT) and it can be used for non-VPN applications. For instance, you can use it to carve out a dedicated routing-table on a single device.

For your convenience, the interface for the VPN was included in the starting configuration of this lab. When configuring the VPN, we only need to place the interface into the proper table. In this case, we place the xe-0/0/5.10 interface into the newly created VRF:

```
set routing-instances JNCIE-BLUE interface xe-0/0/5.10
```

As we will see during the verification, every MPLS L3VPN route will have an RT attached to it. This value determines the VPN membership. We configure the same RT for the JNCIE-BLUE VPN on every PE. By using the 'vrf-target' statement, we ensure both import as well as export of the RT:

```
set routing-instances JNCIE-BLUE vrf-target target:65000:1
```

The command 'vrf-table-label' eases the VPN configuration. This configuration knob will have the router allocate a single label for all the routes in the VRF and by using this option, we do not have to configure any additional vrf-export policies for the directly connected and static routes:

```
set routing-instances JNCIE-BLUE vrf-table-label
```

Static routes for MPLS L3VPN are configured in the routing-instances hierarchy:

```
set routing-instances JNCIE-BLUE routing-options static route 192.168.1.0/24 next-hop
10.0.10.2
```

vMX2:

```
set routing-options route-distinguisher-id 172.16.1.2

set routing-instances JNCIE-BLUE instance-type vrf
set routing-instances JNCIE-BLUE interface xe-0/0/5.10
set routing-instances JNCIE-BLUE vrf-target target:65000:1
set routing-instances JNCIE-BLUE vrf-table-label
set routing-instances JNCIE-BLUE routing-options static route 192.168.2.0/24 next-hop
10.0.10.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set routing-options route-distinguisher-id 172.16.1.3

set routing-instances JNCIE-BLUE instance-type vrf
set routing-instances JNCIE-BLUE interface xe-0/0/5.10
set routing-instances JNCIE-BLUE vrf-target target:65000:1
set routing-instances JNCIE-BLUE vrf-table-label
set routing-instances JNCIE-BLUE routing-options static route 192.168.3.0/24 next-hop
10.0.10.10
```

vMX4:

```
set routing-options route-distinguisher-id 172.16.1.4

set routing-instances JNCIE-BLUE instance-type vrf
set routing-instances JNCIE-BLUE interface xe-0/0/5.10
set routing-instances JNCIE-BLUE vrf-target target:65000:1
set routing-instances JNCIE-BLUE vrf-table-label
set routing-instances JNCIE-BLUE routing-options static route 192.168.4.0/24 next-hop
10.0.10.14
```

Verification

When we issue the 'show bgp summary' command on vMX1, we can see that every other PE has advertised two 'bgp.l3vpn.0' routes:

```
jncie@vMX1> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       6          6          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.2               65000         25         25       0       0        8:05 Establ
  bgp.l3vpn.0: 2/2/2/0
  JNCIE-BLUE.inet.0: 2/2/2/0
172.16.1.3               65000         23         24       0       0        7:48 Establ
  bgp.l3vpn.0: 2/2/2/0
  JNCIE-BLUE.inet.0: 2/2/2/0
172.16.1.4               65000         22         23       0       0        7:40 Establ
  bgp.l3vpn.0: 2/2/2/0
  JNCIE-BLUE.inet.0: 2/2/2/0
```

The output also tells us how many of the received routes are associated with the 'JNCIE-BLUE' VPN. Currently, we only have 1 VPN. When a lot of VPNs are active on a PE, this can be a useful parameter to examine.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To see what prefixes vMX1 is advertising to other PE routers, we can issue the following command:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.2

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.10.0/30            Self                         100        I
* 192.168.1.0/24          Self                         100        I
```

We can see that the directly connected route as well as the static route is advertised to the remote PE. In order to examine the details for this route, we add the detail keyword to the command. As far as MLS L3VPNs go, there are three additional parameters worth mentioning. The RD, the VPN label and the RT:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.2 detail

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
* 10.0.10.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.16.1.1:11
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:1

* 192.168.1.0/24 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.16.1.1:11
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:65000:1
```

The RD is prepended to the VPN prefix to make sure that it is unique. Imagine two VPNs using the same prefix. Without an RD prepended to the prefixes, the prefixes would not be unique. If the prefixes are not unique, customers inside the different VPNs would not be able to use overlapping subnets.

The RD is prepended to the prefix of the route and the RT is an extended BGP community attribute that is added to the route. This community value is used to determine the VPN membership. Effectively, it is the way the vMX1 router is signaling to the other routers that these routes belong to the JNCIE-BLUE VPN. The only thing different is that it is not signaling the value 'JNCIE-BLUE', it is signaling the value 'target:65000:1'. PE routers only

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

look at this attribute when they determine what VPN they should associate the routes with. It is possible to use a different routing-instance name on every PE.

VPN traffic that needs to traverse the MPLS network is send with two labels. One of these labels is the inner (or VPN) label. The PE routers use BGP to signal this label value.

To verify what the vMX1 router is receiving from other PE routers, which can be useful when troubleshooting, we can issue the following command:

```
jncie@vMX1> show route receive-protocol bgp 172.16.1.4 detail

inet.0: 23 destinations, 28 routes (23 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
* 10.0.10.12/30 (1 entry, 1 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.4:11
     VPN Label: 16
     Nexthop: 172.16.1.4
     Localpref: 100
     AS path: I
     Communities: target:65000:1

* 192.168.4.0/24 (1 entry, 1 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.4:11
     VPN Label: 16
     Nexthop: 172.16.1.4
     Localpref: 100
     AS path: I
     Communities: target:65000:1

mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

* 172.16.1.4:11:10.0.10.12/30 (1 entry, 0 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.4:11
     VPN Label: 16
     Nexthop: 172.16.1.4
     Localpref: 100
     AS path: I
     Communities: target:65000:1

* 172.16.1.4:11:192.168.4.0/24 (1 entry, 0 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.4:11
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    VPN Label: 16
    Nexthop: 172.16.1.4
    Localpref: 100
    AS path: I
    Communities: target:65000:1
```

Next up is the verification of the routing table for the instance JNCIE-BLUE:

```
jncie@vMX1> show route table JNCIE-BLUE.inet.0

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.10.0/30       *[Direct/0] 00:04:03
                    > via xe-0/0/5.10
10.0.10.1/32       *[Local/0] 00:04:03
                      Local via xe-0/0/5.10
10.0.10.4/30       *[BGP/170] 00:02:18, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.2 via xe-0/0/0.0, Push 16
10.0.10.8/30       *[BGP/170] 00:02:14, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 16, Push 300624(top)
10.0.10.12/30      *[BGP/170] 00:03:13, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 16, Push 300672(top)
192.168.1.0/24     *[Static/5] 00:04:03
                    > to 10.0.10.2 via xe-0/0/5.10
192.168.2.0/24     *[BGP/170] 00:02:18, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.2 via xe-0/0/0.0, Push 16
192.168.3.0/24     *[BGP/170] 00:02:14, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 16, Push 300624(top)
192.168.4.0/24     *[BGP/170] 00:03:13, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 16, Push 300672(top)
```

This overview tells us about the direct and static route that we configured. In addition to this, we can see all of the subnets reachable via remote PEs. These subnets were learned via BGP. Note how the routing table entry also lists the PE from which the route was learned.

The previous command displayed the routing table entries. Now, let's examine a single entry in greater detail. The following command displays how the 192.168.4.0/24 route is programmed into the forwarding table:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show route forwarding-table vpn JNCIE-BLUE destination 192.168.4.0/24

Routing table: JNCIE-BLUE.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
192.168.4.0/24     user     0                    indr  1048576     3
                              172.16.0.14       Push 16, Push 300672(top)      604     2 xe-
0/0/5.2
```

The entry lists two labels. The first label is the inner, or VPN, label. This label is signaled to the local PE (vMX1) by the remote PE (vMX4). If we go back to previous commands, we can see that this label value (16) corresponds the label value that was learned from a BGP routing update.

The second label is the outer, or transport, label. This label is used to reach the remote PE. As the following command output reveals, it is learned via LDP:

```
jncie@vMX1> show route 172.16.1.4 table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      *[LDP/9] 00:03:19, metric 12
                    > to 172.16.0.14 via xe-0/0/5.2, Push 300672
```

VPN traffic from vMX1 to vMX4 is send with both outer as well as inner label. The second to last router of the LSP will perform a pop operation. During this pop operation, the outer label is popped and the packet that is sent onto vMX4 will only contain the VPN label. When we move over to vMX4, we can check what entry the vMX router has for label 16:

```
jncie@vMX4> show route label 16

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

16                 *[VPN/0] 00:07:55
                    > via lsi.512 (JNCIE-BLUE), Pop
```

On the left, we see the entry for VPN label 16. The vMX4 router is configured with vrf-table-label. Therefore, the router has instantiated a Label Switched Interface (LSI). This logical interface is mapped to the VRF. When MPLS packets destined to the VPN arrive on vMX4, the packet is decapsulated using this interface.

Note: when vrf-table-label is not used, the VPN label is still visible in the routing table. The operation associated with the label will also be a pop-operation, but the next-hop will be the outgoing interface

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The last remaining thing to verify for this task is the verification of the IP-connectivity between the CE devices inside VPN JNCIE-BLUE. To verify IP connectivity, we issue a ping from the CE devices. The CE devices are configured as routing-instances on vMX5 inside the logical system 'CE':

```
jncie@vMX5:CE> ping routing-instance CE1-BLUE rapid 192.168.1.1

PING 192.168.1.1 (192.168.1.1): 56 data bytes
!!!!!
--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.238/0.309/0.354/0.040 ms

jncie@vMX5:CE> ping routing-instance CE1-BLUE rapid 192.168.2.1

PING 192.168.2.1 (192.168.2.1): 56 data bytes
!!!!!
--- 192.168.2.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 218.721/258.466/306.963/30.277 ms

jncie@vMX5:CE> ping routing-instance CE1-BLUE rapid 192.168.3.1
PING 192.168.3.1 (192.168.3.1): 56 data bytes
!!!!!
--- 192.168.3.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.103/39.458/75.450/28.738 ms

jncie@vMX5:CE> ping routing-instance CE1-BLUE rapid 192.168.4.1

PING 192.168.4.1 (192.168.4.1): 56 data bytes
!!!!!
--- 192.168.4.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.008/18.167/37.675/9.825 ms
```

Task 4.6: MPLS L3VPN JNCIE-GREEN

- PE-CE interfaces configured with VLAN-ID 20 should be placed into a VPN called 'JNCIE-GREEN'.
- Configure OSPF as PE-CE routing protocol.
- Choose any route-target you want.
- The PE interface routes should be reachable for remote sites.

Solution

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX1:

We specify the correct instance type:

```
set routing-instances JNCIE-GREEN instance-type vrf
```

We assign the interface to the JNCIE-GREEN instance:

```
set routing-instances JNCIE-GREEN interface xe-0/0/5.20
```

We configure the instance with a unique route-target:

```
set routing-instances JNCIE-GREEN vrf-target target:65000:2
```

We use the vrf-table-label option to ensure reachable of the directly connected interfaces:

```
set routing-instances JNCIE-GREEN vrf-table-label
```

We configure the PE to run OSPF on the interfaces we placed into the VPN:

```
set routing-instances JNCIE-GREEN protocols ospf area 0.0.0.0 interface all
```

Since routes received from remote PE routers will be BGP routes, we need to configure a routing-policy that advertises these routes into OSPF. This policy can be configured as follows:

```
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
```

We need to apply the policy under the ospf hierarchy in the routing-instance:

```
set routing-instances JNCIE-GREEN protocols ospf export bgp-to-ospf
```

vMX2:

```
set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface xe-0/0/5.20
set routing-instances JNCIE-GREEN vrf-target target:65000:2
set routing-instances JNCIE-GREEN vrf-table-label
set routing-instances JNCIE-GREEN protocols ospf export bgp-to-ospf
set routing-instances JNCIE-GREEN protocols ospf area 0.0.0.0 interface all

set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface xe-0/0/5.20
set routing-instances JNCIE-GREEN vrf-target target:65000:2
set routing-instances JNCIE-GREEN vrf-table-label
set routing-instances JNCIE-GREEN protocols ospf export bgp-to-ospf
set routing-instances JNCIE-GREEN protocols ospf area 0.0.0.0 interface all

set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
```

vMX4:

```
set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface xe-0/0/5.20
set routing-instances JNCIE-GREEN vrf-target target:65000:2
set routing-instances JNCIE-GREEN vrf-table-label
set routing-instances JNCIE-GREEN protocols ospf export bgp-to-ospf
set routing-instances JNCIE-GREEN protocols ospf area 0.0.0.0 interface all

set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
```

Verification

We move from one PE to the next and check both advertised and received routes:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.4

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.10.0/30            Self                         100        I
* 192.168.1.0/24          Self                         100        I

JNCIE-GREEN.inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.20.0/30            Self                         100        I
* 192.168.21.0/24         Self                 1       100        I
* 192.168.21.1/32         Self                 1       100        I

jncie@vMX1> show route receive-protocol bgp 172.16.1.4

inet.0: 21 destinations, 26 routes (21 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

JNCIE-BLUE.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.10.12/30           172.16.1.4                   100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 192.168.4.0/24          172.16.1.4                    100        I

JNCIE-GREEN.inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.20.12/30           172.16.1.4                   100        I
* 192.168.24.0/24         172.16.1.4           1       100        I
* 192.168.24.1/32         172.16.1.4           1       100        I

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.1.4:11:10.0.10.12/30
*                         172.16.1.4                   100        I
  172.16.1.4:11:192.168.4.0/24
*                         172.16.1.4                   100        I
  172.16.1.4:12:10.0.20.12/30
*                         172.16.1.4                   100        I
  172.16.1.4:12:192.168.24.0/24
*                         172.16.1.4           1       100        I
  172.16.1.4:12:192.168.24.1/32
*                         172.16.1.4           1       100        I
```

On every PE, we should see one OSPF neighbor relationship:

```
jncie@vMX1> show ospf neighbor instance JNCIE-GREEN

Address          Interface              State     ID               Pri  Dead
10.0.20.2        xe-0/0/5.20            Full      192.168.21.1     128    36
```

When we check the routing table for the JNCIE-GREEN instance, we notice direct, OSPF and BGP routes:

```
jncie@vMX1> show route table JNCIE-GREEN

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

10.0.20.0/30       *[Direct/0] 00:00:39
                    > via xe-0/0/5.20
10.0.20.1/32       *[Local/0] 00:00:39
                      Local via xe-0/0/5.20
10.0.20.4/30       *[BGP/170] 00:00:29, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.2 via xe-0/0/0.0, Push 17
10.0.20.8/30       *[BGP/170] 00:00:20, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300624(top)
10.0.20.12/30      *[BGP/170] 00:00:12, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300672(top)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
192.168.21.0/24    *[OSPF/10] 00:00:25, metric 1
                    > to 10.0.20.2 via xe-0/0/5.20
192.168.21.1/32    *[OSPF/10] 00:00:25, metric 1
                    > to 10.0.20.2 via xe-0/0/5.20
192.168.22.0/24    *[BGP/170] 00:00:15, MED 1, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.2 via xe-0/0/0.0, Push 17
192.168.22.1/32    *[BGP/170] 00:00:15, MED 1, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.2 via xe-0/0/0.0, Push 17
192.168.23.0/24    *[BGP/170] 00:00:11, MED 1, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300624(top)
192.168.23.1/32    *[BGP/170] 00:00:11, MED 1, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300624(top)
192.168.24.0/24    *[BGP/170] 00:00:03, MED 1, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300672(top)
192.168.24.1/32    *[BGP/170] 00:00:03, MED 1, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, Push 17, Push 300672(top)
224.0.0.5/32       *[OSPF/10] 00:00:40, metric 1
                      MultiRecv
```

The BGP routes are received from remote PE routers. The check whether or not the local PE is advertising the received BGP routes into OSPF, we issue the following command:

```
jncie@vMX1> show ospf database advertising-router self instance JNCIE-GREEN

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *10.0.20.1       10.0.20.1       0x80000048    55  0x22 0xa0ca  36
Summary *192.168.22.0    10.0.20.1       0x80000001    41  0xa2 0x7aa2  28
Summary *192.168.22.1    10.0.20.1       0x80000001    41  0xa2 0x70ab  28
Summary *192.168.23.0    10.0.20.1       0x80000001    37  0xa2 0x6fac  28
Summary *192.168.23.1    10.0.20.1       0x80000001    37  0xa2 0x65b5  28
Summary *192.168.24.0    10.0.20.1       0x80000001    29  0xa2 0x64b6  28
Summary *192.168.24.1    10.0.20.1       0x80000001    29  0xa2 0x5abf  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *10.0.20.4       10.0.20.1       0x80000008    55  0xa2 0x5fd5  36
Extern  *10.0.20.8       10.0.20.1       0x80000001    46  0xa2 0x45f2  36
Extern  *10.0.20.12      10.0.20.1       0x80000001    38  0xa2 0x1d17  36
```

The last step is the verification of the IP-connectivity. To this end, we issue several pings from the CE logical-system on vMX5:

```
jncie@vMX5:CE> ping routing-instance CE1-GREEN source 192.168.21.1 192.168.22.1 rapid
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
PING 192.168.22.1 (192.168.22.1): 56 data bytes
!!!!!
--- 192.168.22.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 64.777/89.606/122.371/19.081 ms

jncie@vMX5:CE> ping routing-instance CE1-GREEN source 192.168.21.1 192.168.23.1 rapid

PING 192.168.23.1 (192.168.23.1): 56 data bytes
!!!!!
--- 192.168.23.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 147.259/283.395/425.847/98.590 ms

jncie@vMX5:CE> ping routing-instance CE1-GREEN source 192.168.21.1 192.168.24.1 rapid

PING 192.168.24.1 (192.168.24.1): 56 data bytes
!!!!!
--- 192.168.24.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 15.070/21.333/26.648/4.549 ms
```

Task 4.7: MPLS L3VPN JNCIE-RED

- PE-CE interfaces configured with VLAN-ID 30 should be placed into a VPN called ‘JNCIE-RED’.
- Configure BGP as PE-CE routing protocol. CE devices in this VPN are using AS65500.
- The site connected to vMX1 and vMX2 is multihomed. Make sure traffic to and from that site traverses vMX1.
- Choose any route-target you want.
- The PE interface routes should be reachable for remote sites.

Solution

vMX1:

We configure a routing-instance of the type VRF and place the interface into the instance:

```
set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface xe-0/0/5.30
```

We assign a unique RT to the instance:

```
set routing-instances JNCIE-RED vrf-target target:65000:3
```

The option vrf-table-label is used to make sure the direct routes in the VPN are reachable to remote sites:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances JNCIE-RED vrf-table-label
```

The BGP configuration for this VPN is applied under the [routing-instances JNCIE-RED protocols bgp] hierarchy.

The CE device is using AS65500 to source the BGP session:

```
set routing-instances JNCIE-RED protocols bgp group CE peer-as 65500
```

Since all of the CE devices in the VPN are using the same AS, we use the AS-override knob to make sure that there is IP-connectivity between the sites. The AS-override command will make the PE replace AS65500 in the AS-path attribute with its own AS. This way, the BGP loop prevention mechanism is not triggered:

```
set routing-instances JNCIE-RED protocols bgp group CE as-override
```

Finally, we specify the IP address of the BGP peer:

```
set routing-instances JNCIE-RED protocols bgp group CE neighbor 10.0.30.2
```

The task specifies we should use the link between the CE1 and the vMX1 PE. There are a lot of different ways that we can achieve this goal. In this example, we will raise the local preference of the routes on vMX1. The highest local-preference value wins. This means we can achieve our goal by setting the value to anything above the100:

```
set routing-instances JNCIE-RED protocols bgp group CE import CE1-import
set policy-options policy-statement CE1-import term local-preference then local-preference 150
```

vMX2:

```
set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface xe-0/0/5.30
set routing-instances JNCIE-RED vrf-target target:65000:3
set routing-instances JNCIE-RED vrf-table-label
set routing-instances JNCIE-RED protocols bgp group CE peer-as 65500
set routing-instances JNCIE-RED protocols bgp group CE as-override
set routing-instances JNCIE-RED protocols bgp group CE neighbor 10.0.30.6
```

In order to make routes advertised from the vMX2 PE appear less desirable to the CE device, we use AS-path prepending to increase the AS-path length:

```
set routing-instances JNCIE-RED protocols bgp group CE export CE1-export
set policy-options policy-statement CE1-export term as-prepend then as-path-prepend "65000
65000"
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface xe-0/0/5.30
set routing-instances JNCIE-RED vrf-target target:65000:3
set routing-instances JNCIE-RED vrf-table-label
set routing-instances JNCIE-RED protocols bgp group CE peer-as 65500
set routing-instances JNCIE-RED protocols bgp group CE as-override
set routing-instances JNCIE-RED protocols bgp group CE neighbor 10.0.30.10
```

vMX4:

```
set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface xe-0/0/5.30
set routing-instances JNCIE-RED vrf-target target:65000:3
set routing-instances JNCIE-RED vrf-table-label
set routing-instances JNCIE-RED protocols bgp group CE peer-as 65500
set routing-instances JNCIE-RED protocols bgp group CE as-override
set routing-instances JNCIE-RED protocols bgp group CE neighbor 10.0.30.14
```

Verification

The 'show bgp summary' command displays BGP sessions for all routing-instances:

```
jncie@vMX1> show bgp summary

Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                      20         20          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.30.2             65500          8         12       0       0        2:39 Establ
  JNCIE-RED.inet.0: 2/2/2/0
172.16.1.2            65000         67         72       0       0       19:33 Establ
  bgp.l3vpn.0: 6/6/6/0
  JNCIE-BLUE.inet.0: 2/2/2/0
  JNCIE-GREEN.inet.0: 3/3/3/0
  JNCIE-RED.inet.0: 1/1/1/0
172.16.1.3            65000         64         73       0       0       19:16 Establ
  bgp.l3vpn.0: 7/7/7/0
  JNCIE-BLUE.inet.0: 2/2/2/0
  JNCIE-GREEN.inet.0: 3/3/3/0
  JNCIE-RED.inet.0: 2/2/2/0
172.16.1.4            65000         68         72       0       0       19:08 Establ
  bgp.l3vpn.0: 7/7/7/0
  JNCIE-BLUE.inet.0: 2/2/2/0
  JNCIE-GREEN.inet.0: 3/3/3/0
  JNCIE-RED.inet.0: 2/2/2/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Checking advertised and received routes is done in the same way that BGP sessions sourced from addresses in the main table are checked:

```
jncie@vMX1> show route receive-protocol bgp 10.0.30.2 table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               10.0.30.2                               65500 I
* 192.168.31.0/24         10.0.30.2                               65500 I

jncie@vMX1> show route advertising-protocol bgp 10.0.30.2 table JNCIE-RED.inet.0
JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.30.4/30            Self                                    I
* 10.0.30.8/30            Self                                    I
* 10.0.30.12/30           Self                                    I
* 192.168.32.0/24         Self                                    65000 I
* 192.168.33.0/24         Self                                    65000 II
```

Notice how the received routes display the routes in the way that they were received from the CE device. When routes are advertised across an EBGP session, the routes do not have a local-preference attached to it. The local preference is added to the routes after importing them into the routing table. So, to check whether or not the local preference has been increased, we issue the following command:

```
jncie@vMX1> show route protocol bgp table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0               *[BGP/170] 00:07:02, localpref 150
                          AS path: 65500 I, validation-state: unverified
                        > to 10.0.30.2 via xe-0/0/5.30
10.0.30.4/30            *[BGP/170] 00:05:52, localpref 100, from 172.16.1.2
                          AS path: I, validation-state: unverified
                        > to 172.16.0.2 via xe-0/0/0.0, Push 18
10.0.30.8/30            *[BGP/170] 00:06:34, localpref 100, from 172.16.1.3
                          AS path: I, validation-state: unverified
                        > to 172.16.0.14 via xe-0/0/5.2, Push 18, Push 300624(top)
10.0.30.12/30           *[BGP/170] 00:06:20, localpref 100, from 172.16.1.4
                          AS path: I, validation-state: unverified
                        > to 172.16.0.14 via xe-0/0/5.2, Push 18, Push 300672(top)
192.168.31.0/24         *[BGP/170] 00:07:02, localpref 150
                          AS path: 65500 I, validation-state: unverified
                        > to 10.0.30.2 via xe-0/0/5.30
192.168.32.0/24         *[BGP/170] 00:06:29, localpref 100, from 172.16.1.3
                          AS path: 65500 I, validation-state: unverified
                        > to 172.16.0.14 via xe-0/0/5.2, Push 18, Push 300624(top)
192.168.33.0/24         *[BGP/170] 00:06:16, localpref 100, from 172.16.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                  AS path: 65500 I, validation-state: unverified
                > to 172.16.0.14 via xe-0/0/5.2, Push 18, Push 300672(top)
```

To verify the export policy on vMX2:

```
jncie@vMX2> show route advertising-protocol bgp 10.0.30.6 table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 9 destinations, 11 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    65000 65000 [65000] 65000 I
* 10.0.30.0/30            Self                                    65000 65000 [65000] I
* 10.0.30.8/30            Self                                    65000 65000 [65000] I
* 10.0.30.12/30           Self                                    65000 65000 [65000] I
* 192.168.32.0/24         Self                                    65000 65000 [65000] 65000 I
* 192.168.33.0/24         Self                                    65000 65000 [65000] 65000 I
```

We see that we have correctly prepended the AS path to make vMX2 appear less as a less attractive path.

When we check the default route for this VPN on another PE, we should see that it was received from vMX1:

```
jncie@vMX4> show route table JNCIE-RED.inet.0 0.0.0.0

JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0               *[BGP/170] 00:02:00, localpref 150, from 172.16.1.1
                          AS path: 65500 I, validation-state: unverified
                        > to 172.16.0.26 via xe-0/0/5.2, Push 18, Push 299856(top)
```

We can only see 1 route on the vMX4 router. This is due to the fact that the PE routers will only advertise the best route. We can see that vMX1 advertises the default route to vMX4 with a local preference value of 150:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.4 table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                            150        65500 I
* 10.0.30.0/30            Self                            100        I
* 192.168.31.0/24         Self                            150        65500 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The vMX2 router is not advertising the default route:

```
jncie@vMX2> show route advertising-protocol bgp 172.16.1.4 table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 9 destinations, 11 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.0.30.4/30            Self                         100        I
```

This is due to the fact that the default route in use on vMX2 is the default route that was received from the vMX1 router:

```
jncie@vMX2> show route table JNCIE-RED.inet.0 0.0.0.0

JNCIE-RED.inet.0: 9 destinations, 11 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[BGP/170] 00:09:43, localpref 150, from 172.16.1.1
                      AS path: 65500 I, validation-state: unverified
                    > to 172.16.0.1 via xe-0/0/0.0, Push 18
                    [BGP/170] 00:04:53, localpref 100
                      AS path: 65500 I, validation-state: unverified
                    > to 10.0.30.6 via xe-0/0/5.30
```

Next is the verification of IP connectivity between the various sites:

```
jncie@vMX5:CE> ping routing-instance CE1-RED source 192.168.31.1 192.168.32.1 rapid

PING 192.168.32.1 (192.168.32.1): 56 data bytes
!!!!!
--- 192.168.32.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 55.992/275.717/485.537/161.237 ms

jncie@vMX5:CE> ping routing-instance CE1-RED source 192.168.31.1 192.168.33.1 rapid

PING 192.168.33.1 (192.168.33.1): 56 data bytes
!!!!!
--- 192.168.33.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.983/39.967/90.642/31.412 ms
```

The last thing we verify is the path packets towards CE1 take. Observe the following traceroute from CE3 and CE4 towards CE1:

```
jncie@vMX5:CE> traceroute routing-instance CE3-RED  192.168.31.1

traceroute to 192.168.31.1 (192.168.31.1), 30 hops max, 52 byte packets
 1  10.0.30.9 (10.0.30.9)  7.270 ms  13.042 ms  8.041 ms
 2  10.0.30.1 (10.0.30.1)  333.275 ms  426.667 ms  39.167 ms
 3  192.168.31.1 (192.168.31.1)  792.213 ms  549.786 ms  697.618 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX5:CE> traceroute routing-instance CE4-RED  192.168.31.1

traceroute to 192.168.31.1 (192.168.31.1), 30 hops max, 52 byte packets
 1  10.0.30.13 (10.0.30.13)  9.380 ms  20.112 ms  13.003 ms
 2  10.0.30.1 (10.0.30.1)  18.006 ms  19.769 ms  20.066 ms
 3  192.168.31.1 (192.168.31.1)  19.921 ms  18.950 ms *
```

In both cases, the second to last hop was the link between vMX1 and CE1:

```
jncie@vMX1> show route table JNCIE-RED.inet.0 protocol direct

JNCIE-RED.inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.30.0/30           *[Direct/0] 00:37:43
                        > via xe-0/0/5.30
```

To verify redundancy, we temporarily deactivate BGP on the vMX1 router:

```
jncie@vMX1> configure
Entering configuration mode

[edit]
jncie@vMX1# deactivate protocols bgp

[edit]
jncie@vMX1# commit
commit complete

[edit]
jncie@vMX1#
```

When we issue the same traceroutes, we can see that the second to last hop has changed:

```
jncie@vMX5:CE> traceroute routing-instance CE3-RED 192.168.31.1

traceroute to 192.168.31.1 (192.168.31.1), 30 hops max, 52 byte packets
 1  10.0.30.9 (10.0.30.9)  1539.560 ms  607.598 ms  11.984 ms
 2  10.0.30.5 (10.0.30.5)  13.488 ms  21.168 ms  14.057 ms
 3  192.168.31.1 (192.168.31.1)  13.028 ms  12.587 ms  11.804 ms

jncie@vMX5:CE> traceroute routing-instance CE4-RED 192.168.31.1

traceroute to 192.168.31.1 (192.168.31.1), 30 hops max, 52 byte packets
 1  10.0.30.13 (10.0.30.13)  26.834 ms  7.866 ms  7.710 ms
 2  10.0.30.5 (10.0.30.5)  13.585 ms  24.295 ms  11.821 ms
 3  192.168.31.1 (192.168.31.1)  14.880 ms  15.742 ms  16.880 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

This subnet is in use between CE1 and vMX2:

```
jncie@vMX2> show route table JNCIE-RED.inet.0 protocol direct

JNCIE-RED.inet.0: 9 destinations, 11 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.30.4/30       *[Direct/0] 00:39:45
                    > via xe-0/0/5.30
```

Before moving on, return to vMX1 and use the rollback command to enable the BGP protocol.

Since we can also log into the CE, we can also verify things from the CE point-of-view. Observe the following:

```
jncie@vMX5:CE> show route table CE1-RED protocol bgp

CE1-RED.inet.0: 11 destinations, 18 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

...
192.168.32.0/24    *[BGP/170] 00:03:16, localpref 100
                      AS path: 65000 65000 I, validation-state: unverified
                    > to 10.0.30.1 via xe-0/0/1.30
                    [BGP/170] 00:12:32, localpref 100
                      AS path: 65000 65000 65000 65000 I, validation-state: unverified
                    > to 10.0.30.5 via xe-0/0/2.30
...
```

We can see that the CE is receiving routes from both PE routers. It is preferring the route via vMX1 due to the shorter AS path.

An additional thing we notice here is the effect of the as-override knob. The AS in use on the other VPN sites does not appear in the AS path.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 2: MPLS, MPLS L3VPN and EVPN.

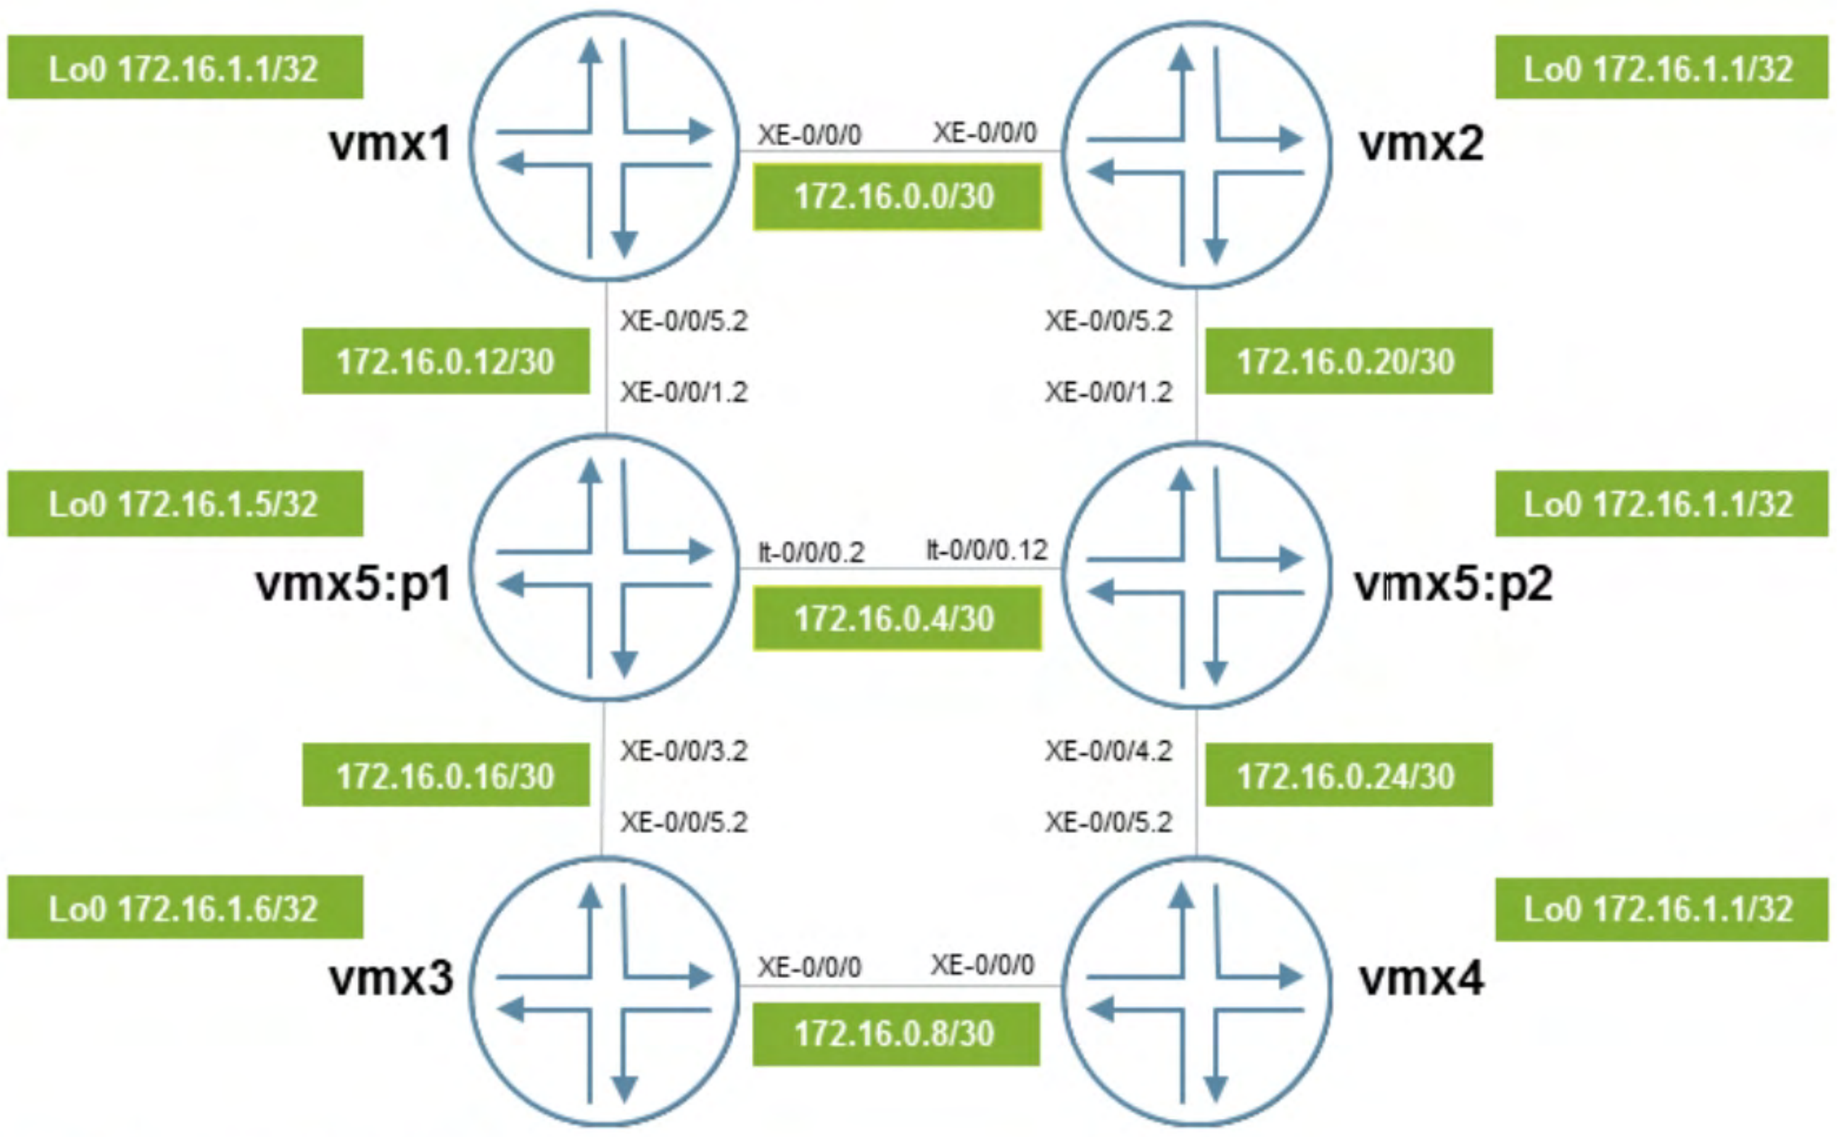

Figure 23 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

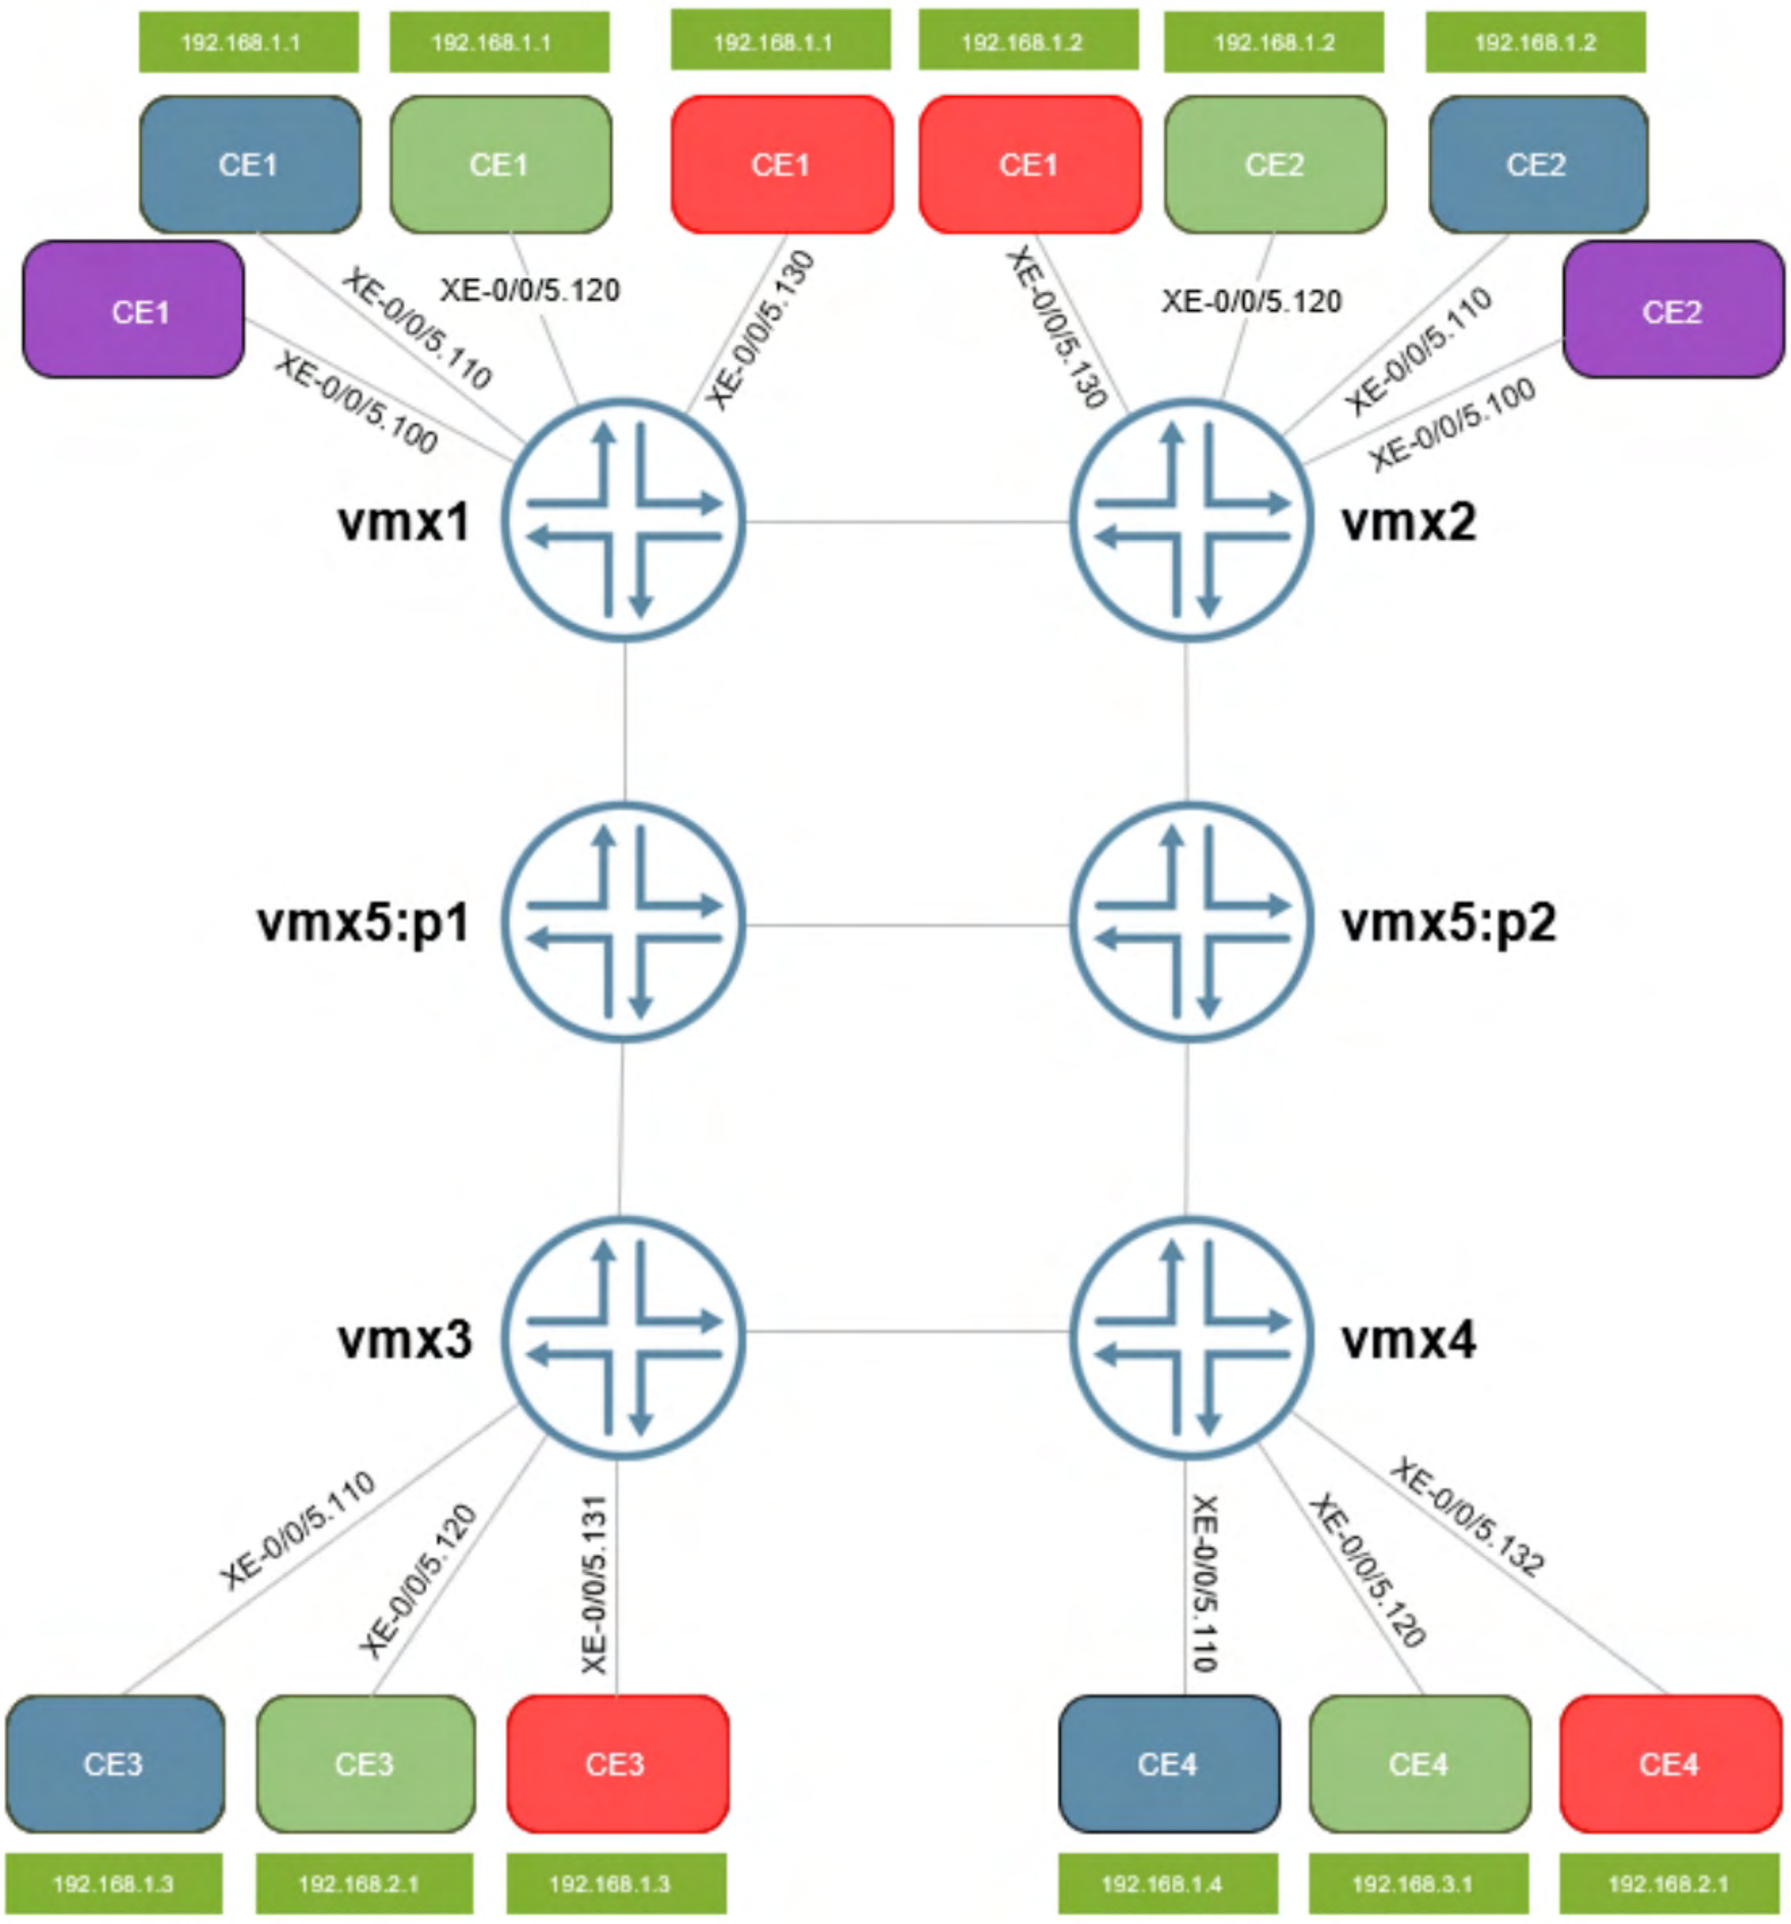

Figure 2 – VPN topology

Notes:

- Starting topology is configured with IP-addresses.
- OSPF is already running on the blue links.
- On vMX5, three logical systems are configured. These are P1, P2 and CE.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.8: RSVP Configuration

- Configure OSPF to support traffic engineering.
- Enable RSVP on all core interfaces.

Solution

vMX1:

Enabling traffic engineering for OSPF:

```
set protocols ospf traffic-engineering
```

Enabling interfaces for RSVP:

```
set protocols rsvp interface xe-0/0/0.0
set protocols rsvp interface xe-0/0/5.2
```

Interfaces for RSVP need to be enabled for MPLS as well. Enabling interfaces for MPLS requires us to configure the MPLS family under the [interfaces] stanza as well as activating the interface under the [protocols mpls interface] stanza:

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2
```

vMX2:

```
set protocols ospf traffic-engineering

set protocols rsvp interface xe-0/0/0.0
set protocols rsvp interface xe-0/0/5.2

set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2
```

vMX3:

```
set protocols ospf traffic-engineering

set protocols rsvp interface xe-0/0/0.0
set protocols rsvp interface xe-0/0/5.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2
```

vMX4:

```
set protocols ospf traffic-engineering

set protocols rsvp interface xe-0/0/0.0
set protocols rsvp interface xe-0/0/5.2
```

```
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/5 unit 2 family mpls

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/5.2
```

vMX5:P1:

```
edit logical-systems P1
set protocols ospf traffic-engineering

set protocols rsvp interface lt-0/0/0.2
set protocols rsvp interface xe-0/0/1.2
set protocols rsvp interface xe-0/0/3.2

set interfaces lt-0/0/0 unit 2 family mpls
set interfaces xe-0/0/1 unit 2 family mpls
set interfaces xe-0/0/3 unit 2 family mpls

set protocols mpls interface xe-0/0/3.2
set protocols mpls interface xe-0/0/1.2
set protocols mpls interface lt-0/0/0.2
```

vMX5:P2:

```
edit logical-systems P2
set protocols ospf traffic-engineering

set protocols rsvp interface lt-0/0/0.12
set protocols rsvp interface xe-0/0/2.2
set protocols rsvp interface xe-0/0/4.2

set interfaces lt-0/0/0 unit 12 family mpls
set interfaces xe-0/0/2 unit 2 family mpls
set interfaces xe-0/0/4 unit 2 family mpls

set protocols mpls interface lt-0/0/0.12
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols mpls interface xe-0/0/2.2
set protocols mpls interface xe-0/0/4.2
```

Verification

To verify that OSPF supports traffic engineering:

```
jncie@vMX1> show ospf overview

Instance: master
  Router ID: 172.16.1.1
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 2
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 5850
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

Verify that interfaces have been enabled for MPLS and RSVP:

```
jncie@vMX1> show mpls interface

Interface          State       Administrative groups (x: extended)
xe-0/0/0.0         Up          <none>
xe-0/0/5.2         Up          <none>
```

```
jncie@vMX1> show rsvp interface

RSVP interface: 2 active
                      Active  Subscr- Static      Available   Reserved  Highwater
Interface        State  resv    iption  BW              BW          mark
xe-0/0/0.0           Up       0   100%  10Gbps      10Gbps      0bps         0bps
xe-0/0/5.2           Up       0   100%  10Gbps      10Gbps      0bps         0bps
```

Task 4.9: RSVP LSP Configuration

- Enable a full-mesh of RSVP signaled LSPs between vMX1, vMX2, vMX3 and vMX4.
- Enable a backup path for all the LSPs that is ready to take over traffic in case of any failures.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Solution

vMX1:

Unlike with LDP, simply activating RSVP will not provide you with any LSPs. For RSVP to start signaling LSPs, you will need to configure them first under the [protocols mpls label-switched-path] hierarchy. Since we need to establish a full-mesh of RSVP signaled LSPs, we need to configure three LSPs on every PE router:

```
set protocols mpls label-switched-path vmx1-to-vmx2 to 172.16.1.2
set protocols mpls label-switched-path vmx1-to-vmx3 to 172.16.1.3
set protocols mpls label-switched-path vmx1-to-vmx4 to 172.16.1.4
```

We conveniently name them in such a way that their names reflect the start end the end of the LSP. This can be helpful during troubleshooting.

To enable a backup path for an LSP, we need to configure a primary and secondary path for the LSPs under the [protocols mpls path] configuration stanza:

```
set protocols mpls path pri
set protocols mpls path sec
```

We now have two empty paths that we can apply to every LSP. However, instead of applying them to every LSP individually, we configure and apply a group:

```
set groups mpls-lsp protocols mpls label-switched-path <*> primary pri
set groups mpls-lsp protocols mpls label-switched-path <*> secondary sec standby

set apply-groups mpls-lsp
```

Notice the 'standby' keyword that is configured for the secondary path. This keyword will ensure that the PE signals the standby path before losing the primary path.

vMX2:

```
set groups mpls-lsp protocols mpls label-switched-path <*> primary pri
set groups mpls-lsp protocols mpls label-switched-path <*> secondary sec standby
set apply-groups mpls-lsp

set protocols mpls label-switched-path vmx2-to-vmx1 to 172.16.1.1
set protocols mpls label-switched-path vmx2-to-vmx3 to 172.16.1.3
set protocols mpls label-switched-path vmx2-to-vmx4 to 172.16.1.4

set protocols mpls path pri
set protocols mpls path sec
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set groups mpls-lsp protocols mpls label-switched-path <*> primary pri
set groups mpls-lsp protocols mpls label-switched-path <*> secondary sec standby
set apply-groups mpls-lsp

set protocols mpls label-switched-path vmx3-to-vmx1 to 172.16.1.1
set protocols mpls label-switched-path vmx3-to-vmx2 to 172.16.1.2
set protocols mpls label-switched-path vmx3-to-vmx4 to 172.16.1.4

set protocols mpls path pri
set protocols mpls path sec
```

vMX4:

```
set groups mpls-lsp protocols mpls label-switched-path <*> primary pri
set groups mpls-lsp protocols mpls label-switched-path <*> secondary sec standby
set apply-groups mpls-lsp

set protocols mpls label-switched-path vmx4-to-vmx1 to 172.16.1.1
set protocols mpls label-switched-path vmx4-to-vmx2 to 172.16.1.2
set protocols mpls label-switched-path vmx4-to-vmx3 to 172.16.1.3

set protocols mpls path pri
set protocols mpls path sec
```

Verification

Verifying whether or not the PE was able to stand up the LSPs can be done by issuing the following command:

```
jncie@vMX1> show mpls lsp

Ingress LSP: 3 sessions
To              From            State Rt P    ActivePath      LSPname
172.16.1.2      172.16.1.1      Up     0 *    pri             vmx1-to-vmx2
172.16.1.3      172.16.1.1      Up     0 *    pri             vmx1-to-vmx3
172.16.1.4      172.16.1.1      Up     0 *    pri             vmx1-to-vmx4
Total 3 displayed, Up 3, Down 0

Egress LSP: 4 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.16.1.1      172.16.1.4      Up       0  1 FF       3         - vmx4-to-vmx1
172.16.1.1      172.16.1.2      Up       0  1 FF       3         - vmx2-to-vmx1
172.16.1.1      172.16.1.2      Up       0  1 FF       3         - vmx2-to-vmx1
172.16.1.1      172.16.1.3      Up       0  1 FF       3         - vmx3-to-vmx1
Total 4 displayed, Up 4, Down 0

Transit LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.1.3      172.16.1.2      Dn       0  0  -       -         - vmx2-to-vmx3
172.16.1.4      172.16.1.2      Up       0  1 FF  445376    465744 vmx2-to-vmx4
Total 2 displayed, Up 1, Down 1
```

This command tells us that the local PE was able to signal the LSPs towards the other PE routers. In addition to this, we can see that the PE is an egress LSR for LSPs signaled by the other PEs as well.

To verify that the LSP has a primary and a backup path, we need to add the 'detail' keyword:

```
jncie@vMX1> show mpls lsp ingress detail

Ingress LSP: 3 sessions

172.16.1.2
  From: 172.16.1.1, State: Up, ActiveRoute: 0, LSPname: vmx1-to-vmx2
  ActivePath: pri (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   pri              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 172.16.0.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          172.16.0.2
  Standby   sec              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 21)
 172.16.0.14 S 172.16.0.6 S 172.16.0.21 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          172.16.0.14 172.16.0.6 172.16.0.21

  172.16.1.3
    From: 172.16.1.1, State: Up, ActiveRoute: 0, LSPname: vmx1-to-vmx3
    ActivePath: pri (primary)
    LSPtype: Static Configured, Penultimate hop popping
    LoadBalance: Random
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
   *Primary   pri              State: Up
      Priorities: 7 0
      SmartOptimizeTimer: 180
      Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
   172.16.0.14 S 172.16.0.18 S
      Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
            172.16.0.14 172.16.0.18
    Standby   sec              State: Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 31)
 172.16.0.2 S 172.16.0.22 S 172.16.0.5 S 172.16.0.18 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          172.16.0.2 172.16.0.22 172.16.0.5 172.16.0.18

172.16.1.4
  From: 172.16.1.1, State: Up, ActiveRoute: 0, LSPname: vmx1-to-vmx4
  ActivePath: pri (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   pri              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 21)
 172.16.0.14 S 172.16.0.6 S 172.16.0.25 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          172.16.0.14 172.16.0.6 172.16.0.25
  Standby   sec              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 172.16.0.2 S 172.16.0.22 S 172.16.0.25 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
          172.16.0.2 172.16.0.22 172.16.0.25
Total 3 displayed, Up 3, Down 0
```

An ingress LSR precomputes the path before signaling an LSP. This path is the 'Computed ERO'. In the previous command output, we can see this computed ERO (Explicit Route Object) for every path.

By examining the IP addresses listed in the ERO, we know exactly what path the LSP takes. Also, notice how the LSR has two different EROs for the primary and the secondary LSP. By default, when possible, the primary and backup LSP will be signaled across different links.

We can also use the 'show rsvp session' command to verify the establishment of the RSVP signaled LSPs:

```
jncie@vMX1> show rsvp session detail

Ingress RSVP: 6 sessions

172.16.1.2
  From: 172.16.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: vmx1-to-vmx2, LSPpath: Primary
  LSPtype: Static Configured
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 465968
Resv style: 1 FF, Label in: -, Label out: 465968
Time left:    -, Since: Wed Apr 26 06:20:04 2017
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 3 receiver 56436 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 172.16.0.14 (xe-0/0/5.2) 2 pkts
RESV rcvfrom: 172.16.0.14 (xe-0/0/5.2) 2 pkts, Entropy label: Yes
Explct route: 172.16.0.14 172.16.0.6 172.16.0.21
Record route: <self> 172.16.0.14 172.16.0.6 172.16.0.21
```

By default, the Junos OS will place the egress address for every LSP into the inet.3 table:

```
jncie@vMX1> show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.2/32      *[RSVP/7/1] 00:00:22, metric 10
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
172.16.1.3/32      *[RSVP/7/1] 00:01:22, metric 20
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
172.16.1.4/32      *[RSVP/7/1] 00:01:22, metric 21
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
```

Task 4.10: BGP Configuration

- Configure a full mesh of IBGP sessions between vMX1, vMX2, vMX3 and vMX4.
- Use autonomous system 65000.
- Authenticate BGP sessions using the MD5 hash ‘jncie’.
- Enable the sessions to support EVPN as well as L3VPN.

Solution

In this chapter, we will configure several EVPNs. Some of these EVPNs will be integrated into an MPLS L3VPN. To this end, we need a full mesh of IBGP sessions that support ‘evpn signaling’ as well as ‘inet-vpn unicast’.

vMX1:

```
set routing-options autonomous-system 65000

set protocols bgp group ibgp type internal
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group ibgp local-address 172.16.1.1
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.3
set protocols bgp group ibgp neighbor 172.16.1.4
```

vMX2:

```
set routing-options autonomous-system 65000

set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.2
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.1
set protocols bgp group ibgp neighbor 172.16.1.3
set protocols bgp group ibgp neighbor 172.16.1.4
```

vMX3:

```
set routing-options autonomous-system 65000

set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.3
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.1
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.4
```

vMX4:

```
set routing-options autonomous-system 65000

set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 172.16.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family evpn signaling
set protocols bgp group ibgp authentication-key jncie
set protocols bgp group ibgp neighbor 172.16.1.1
set protocols bgp group ibgp neighbor 172.16.1.2
set protocols bgp group ibgp neighbor 172.16.1.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

The following command verifies the full mesh of IBGP sessions that support MPLS L3VPN as well as EVPN signaling:

```
jncie@vMX1> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       0          0          0          0          0          0
bgp.evpn.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.1.2            65000          6          5       0       0        1:09 Establ
  bgp.l3vpn.0: 0/0/0/0
  bgp.evpn.0: 0/0/0/0
172.16.1.3            65000          6          5       0       0          58 Establ
  bgp.l3vpn.0: 0/0/0/0
  bgp.evpn.0: 0/0/0/0
172.16.1.4            65000          5          4       0       0          53 Establ
  bgp.l3vpn.0: 0/0/0/0
  bgp.evpn.0: 0/0/0/0
```

Task 4.11: EVPN JNCIE-PURPLE

- Purple CEs should be placed into an EVPN called ‘JNCIE-PURPLE-EVPN’.
- The PE-CE interfaces must be configured for VLAN 100.
- The EVPN is used for layer 2 services only.

Solution

vMX1:

Only interfaces with IP addresses were preconfigured for this task. Therefore, we start out with the configuration of the PE-CE interface:

```
set interfaces xe-0/0/5 unit 100 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 100 vlan-id 100
```

The xe-0/0/5 interface is being used for different purposes in this lab. Configuration wise, this is made possible by enabling the ‘encapsulation flexible-ethernet-services’ under the main interface.

```
set routing-options route-distinguisher-id 172.16.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Since the EVPN does not need to be VLAN-aware, we can conclude that we should configure a VLAN-based EVPN service. For this reason, we set the instance-type to EVPN:

```
set routing-instances JNCIE-PURPLE-EVPN instance-type evpn
```

We need to configure the routing-instance with a VLAN-ID:

```
set routing-instances JNCIE-PURPLE-EVPN vlan-id 100
```

The PE-CE interface needs to be placed into the routing instance:

```
set routing-instances JNCIE-PURPLE-EVPN interface xe-0/0/5.100
```

Using the vrf-target statement, we have the PE generate an import and export policy for this EVPN:

```
set routing-instances JNCIE-PURPLE-EVPN vrf-target target:100:100
```

We do not need to configure a route-distinguisher for this (or any other) VPN. This is because the PE routers have been configured with a route-distinguisher-id in the [routing-options route-distinguisher-id] hierarchy.

There is only one configuration item left for the EVPN. We need to enable Ethernet VPN for the routing instance:

```
set routing-instances JNCIE-PURPLE-EVPN protocols evpn
```

vMX2:

```
set interfaces xe-0/0/5 unit 100 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 100 vlan-id 100

set routing-options route-distinguisher-id 172.16.1.2

set routing-instances JNCIE-PURPLE-EVPN instance-type evpn
set routing-instances JNCIE-PURPLE-EVPN vlan-id 100
set routing-instances JNCIE-PURPLE-EVPN interface xe-0/0/5.100
set routing-instances JNCIE-PURPLE-EVPN vrf-target target:100:100
set routing-instances JNCIE-PURPLE-EVPN protocols evpn
```

Verification

To obtain a complete overview of the state the EVPN is in, we issue the ‘show evpn instance JNCIE-PURPLE-EVPN extensive’ command on both PE routers:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show evpn instance JNCIE-PURPLE-EVPN extensive

Instance: JNCIE-PURPLE-EVPN
  Route Distinguisher: 172.16.1.1:11
  VLAN ID: 100
  Per-instance MAC route label: 445520
  Per-instance multicast route label: 445504
  MAC database status                         Local  Remote
    MAC advertisements:                           1       1
    MAC+IP advertisements:                        0       0
    Default gateway MAC advertisements:           0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.100    00:00:00:00:00:00:00:00:00:00  single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100                  1    1               Extended        Enabled   445504
  Number of neighbors: 1
    172.16.1.2
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          0
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
  Number of ethernet segments: 0

jncie@vMX2> show evpn instance JNCIE-PURPLE-EVPN extensive

Instance: JNCIE-PURPLE-EVPN
  Route Distinguisher: 172.16.1.2:11
  VLAN ID: 100
  Per-instance MAC route label: 452864
  Per-instance multicast route label: 452848
  MAC database status                         Local  Remote
    MAC advertisements:                           1       1
    MAC+IP advertisements:                        0       0
    Default gateway MAC advertisements:           0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.100    00:00:00:00:00:00:00:00:00:00  single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100                  1    1               Extended        Enabled   452848
  Number of neighbors: 1
    172.16.1.1
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          0
        Inclusive multicast:                   1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
        Ethernet auto-discovery:                  0
  Number of ethernet segments: 0
```

The output is rather information-packed. The most relevant parts of the output have been highlighted. The local interface for the EPVN is up, MAC addresses are learned locally and remote and both PE routers have 'discovered' 1 neighbor for this EVPN.

To get a more detailed view on the MAC table for this EPVN:

```
jncie@vMX1> show evpn mac-table instance JNCIE-PURPLE-EVPN

MAC flags       (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
    O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : JNCIE-PURPLE-EVPN
 Bridging domain : __JNCIE-PURPLE-EVPN__, VLAN : 100
   MAC                 MAC      Logical          NH     RTR
   addresssss          flags    interface        Index  ID
   00:0c:29:b9:e9:4b   D        xe-0/0/5.100
   00:0c:29:b9:e9:55   DC                        1048577 1048577
```

As an alternative to the previous command, we can also examine the EVPN database. Notice how this command also outputs the time on which the PE added the MAC address to the EVPN MAC table:

```
jncie@vMX1> show evpn database

Instance: JNCIE-PURPLE-EVPN
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
100             00:0c:29:b9:e9:4b  xe-0/0/5.100                   Apr 26 06:27:54
100             00:0c:29:b9:e9:55  172.16.1.2                     Apr 26 06:27:56
```

To see what routes we are advertising to the EVPN neighbor:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.2 extensive

JNCIE-PURPLE-EVPN.evpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 2:172.16.1.1:11::100::00:0c:29:b9:e9:4b/304 MAC/IP (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.16.1.1:11
     Route Label: 445520
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:100:100

* 3:172.16.1.1:11::100::172.16.1.1/304 IM (1 entry, 1 announced)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
BGP group ibgp type Internal
    Route Distinguisher: 172.16.1.1:11
    Route Label: 445504
    PMSI: Flags 0x0: Label 445504: Type INGRESS-REPLICATION 172.16.1.1
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:100:100
    PMSI: Flags 0x0: Label 445504: Type INGRESS-REPLICATION 172.16.1.1
```

To see what routes we are receiving from the EVPN neighbor:

```
jncie@vMX1> show route receive-protocol bgp 172.16.1.2 extensive
...

JNCIE-PURPLE-EVPN.evpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

* 2:172.16.1.2:11::100::00:0c:29:b9:e9:55/304 MAC/IP (1 entry, 1 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.2:11
     Route Label: 525264
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: 172.16.1.2
     Localpref: 100
     AS path: I
     Communities: target:100:100

* 3:172.16.1.2:11::100::172.16.1.2/304 IM (1 entry, 1 announced)
     Import Accepted
     Route Distinguisher: 172.16.1.2:11
     Nexthop: 172.16.1.2
     Localpref: 100
     AS path: I
     Communities: target:100:100
     PMSI: Flags 0x0: Label 525248: Type INGRESS-REPLICATION 172.16.1.2
```

Finally, we can check whether or not the CE devices inside the EVPN were able to form an OSPF neighbor relationship:

```
jncie@vMX5:CE> show ospf neighbor instance CE1-PURPLE-EVPN

Address          Interface              State     ID               Pri  Dead
192.168.1.2      xe-0/0/1.100           Full      192.168.1.2      128    35

jncie@vMX5:CE> show ospf neighbor instance CE2-PURPLE-EVPN

Address          Interface              State     ID               Pri  Dead
192.168.1.1      xe-0/0/2.100           Full      192.168.1.1      128    34
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.12: EVPN JNCIE-BLUE

- Blue CEs should be placed into an EVPN called 'JNCIE-BLUE-EVPN'.
- The EVPN is used for layer 2 and layer 3 services.
- Configure a layer 3 gateway on every vMX.
- Use IP address 192.168.1.254/24 as the default gateway address.
- Make sure P1 and P2 can exchange traffic with hosts inside the EPVN.

Solution

vMX1:

```
set interfaces xe-0/0/5 unit 110 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 110 vlan-id 110
```

To offer a default gateway for hosts inside the EVPN, we configure an IRB interface with the same IP address on every vMX:

```
set interfaces irb.110 family inet address 192.168.1.254/24
```

The subnet that the hosts inside the EVPN will use needs to be advertised to other routers in the network. In this example, we inject the prefix into OSPF. The use of the passive keyword makes sure that no OSPF neighbor relationship can be formed across the interface:

```
set protocols ospf area 0.0.0.0 interface irb.110 passive
```

We use a VLAN-based EVPN service here because it is enough to meet the tasks requirements:

```
set routing-instances JNCIE-BLUE-EVPN instance-type evpn
set routing-instances JNCIE-BLUE-EVPN vlan-id 110
set routing-instances JNCIE-BLUE-EVPN interface xe-0/0/5.110
```

We associate the previously configured IRB interface with the EVPN instance:

```
set routing-instances JNCIE-BLUE-EVPN routing-interface irb.110
```

Finally, we configure a unique route-target for this routing instance and we enable EVPN for this instance:

```
set routing-instances JNCIE-BLUE-EVPN vrf-target target:110:110
set routing-instances JNCIE-BLUE-EVPN protocols evpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX2:

```
set interfaces xe-0/0/5 unit 110 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 110 vlan-id 110

set interfaces irb.110 family inet address 192.168.1.254/24

set protocols ospf area 0.0.0.0 interface irb.110 passive

set routing-instances JNCIE-BLUE-EVPN instance-type evpn
set routing-instances JNCIE-BLUE-EVPN vlan-id 110
set routing-instances JNCIE-BLUE-EVPN interface xe-0/0/5.110
set routing-instances JNCIE-BLUE-EVPN routing-interface irb.110
set routing-instances JNCIE-BLUE-EVPN vrf-target target:110:110
set routing-instances JNCIE-BLUE-EVPN protocols evpn
```

vMX3:

```
set interfaces xe-0/0/5 unit 110 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 110 vlan-id 110

set interfaces irb.110 family inet address 192.168.1.254/24

set protocols ospf area 0.0.0.0 interface irb.110 passive

set routing-options route-distinguisher-id 172.16.1.3

set routing-instances JNCIE-BLUE-EVPN instance-type evpn
set routing-instances JNCIE-BLUE-EVPN vlan-id 110
set routing-instances JNCIE-BLUE-EVPN interface xe-0/0/5.110
set routing-instances JNCIE-BLUE-EVPN routing-interface irb.110
set routing-instances JNCIE-BLUE-EVPN vrf-target target:110:110
set routing-instances JNCIE-BLUE-EVPN protocols evpn
```

vMX4:

```
set interfaces xe-0/0/5 unit 110 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 110 vlan-id 110

set interfaces irb.110 family inet address 192.168.1.254/24

set protocols ospf area 0.0.0.0 interface irb.110 passive

set routing-options route-distinguisher-id 172.16.1.4

set routing-instances JNCIE-BLUE-EVPN instance-type evpn
set routing-instances JNCIE-BLUE-EVPN vlan-id 110
set routing-instances JNCIE-BLUE-EVPN interface xe-0/0/5.110
set routing-instances JNCIE-BLUE-EVPN routing-interface irb.110
set routing-instances JNCIE-BLUE-EVPN vrf-target target:110:110
set routing-instances JNCIE-BLUE-EVPN protocols evpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

On vMX1, we can see that interface xe-0/0/5.110 and irb.110 are active. We also see that three neighbors have been discovered:

```
jncie@vMX1> show evpn instance JNCIE-BLUE-EVPN extensive

Instance: JNCIE-BLUE-EVPN
  Route Distinguisher: 172.16.1.1:12
  VLAN ID: 110
  Per-instance MAC route label: 445712
  Per-instance multicast route label: 445696
MAC database status                         Local  Remote
  MAC advertisements:                           2       6
  MAC+IP advertisements:                        2       3
  Default gateway MAC advertisements:           1       3
Number of local interfaces: 1 (1 up)
  Interface name  ESI                                Mode             Status
  xe-0/0/5.110    00:00:00:00:00:00:00:00:00:00      single-homed     Up
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN   VNI    Status  L3 context
  irb.110         110             Up     master
Number of bridge domains: 1
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  110                  1    1     irb.110    Extended        Enabled   445696
Number of neighbors: 3
  172.16.1.2
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           1
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
  172.16.1.3
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           1
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
  172.16.1.4
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           1
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
Number of ethernet segments: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We verify IP connectivity between the hosts inside the EVPN:

```
jncie@vMX5:CE> ping rapid routing-instance CE1-BLUE-EVPN 192.168.1.2

PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.154/35.933/60.474/18.091 ms

jncie@vMX5:CE> ping rapid routing-instance CE1-BLUE-EVPN 192.168.1.3

PING 192.168.1.3 (192.168.1.3): 56 data bytes
!!!!!
--- 192.168.1.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 49.788/172.111/464.277/158.618 ms

jncie@vMX5:CE> ping rapid routing-instance CE1-BLUE-EVPN 192.168.1.4

PING 192.168.1.4 (192.168.1.4): 56 data bytes
!!!!!
--- 192.168.1.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.860/25.187/57.824/16.593 ms
```

We check whether or not P1 and P2 are reachable:

```
jncie@vMX5:CE> ping rapid routing-instance CE1-BLUE-EVPN 172.16.1.5

PING 172.16.1.5 (172.16.1.5): 56 data bytes
!!!!!
--- 172.16.1.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 32.614/99.382/267.055/86.670 ms

jncie@vMX5:CE> ping rapid routing-instance CE1-BLUE-EVPN 172.16.1.6

PING 172.16.1.6 (172.16.1.6): 56 data bytes
!!!!!
--- 172.16.1.6 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.708/15.998/23.170/4.225 ms
```

After this, we check the EVPN database:

```
jncie@vMX1> show evpn database instance JNCIE-BLUE-EVPN

Instance: JNCIE-BLUE-EVPN
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
110             00:05:86:4d:34:f0  172.16.1.3                Apr 05 07:45:30
192.168.1.254
110             00:05:86:72:fb:f0  irb.110                   Apr 05 07:32:12
192.168.1.254
110             00:05:86:82:68:f0  172.16.1.4                Apr 05 07:45:30
192.168.1.254
110             00:05:86:94:9b:f0  172.16.1.2                Apr 05 07:45:30
192.168.1.254
110             00:0c:29:b9:e9:4b  xe-0/0/5.110              Apr 05 07:45:28  192.168.1.1
110             00:0c:29:b9:e9:55  172.16.1.2                Apr 05 07:45:30  192.168.1.2
110             00:0c:29:b9:e9:5f  172.16.1.3                Apr 05 07:45:30  192.168.1.3
110             00:0c:29:b9:e9:69  172.16.1.4                Apr 05 07:45:30  192.168.1.4
```

We can see that we have received a route towards every other IRB interface that is configured on the other PEs. In addition to this, we can see that we have received both MAC and IP address for every host located inside the EVPN.

The vMX routers can use these EVPN host routes to forward traffic directly to the appropriate remote PE. The following output displays the EVPN routes on vMX1. Here, we can see that the route has installed a /32 host route towards every other PE:

```
jncie@vMX1> show route protocol evpn

inet.0: 30 destinations, 33 routes (30 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:02:31
                      > via irb.110
192.168.1.2/32     *[EVPN/7] 00:00:06
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
192.168.1.3/32     *[EVPN/7] 00:02:29
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
192.168.1.4/32     *[EVPN/7] 00:02:29
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
192.168.1.254/32    [EVPN/7] 00:02:29
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                    [EVPN/7] 00:00:06
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                    [EVPN/7] 00:02:29
                      > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
                        to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Notice that the vMX router does not have any ARP entries for the hosts connected to the remote PE routers:

```
jncie@vMX1> show arp interface irb.110

MAC Address       Address         Name                    Interface               Flags
00:0c:29:b9:e9:4b 192.168.1.1     192.168.1.1             irb.110 [xe-0/0/5.110]  none
```

When we check the vMX3 router, we can see that router does have an ARP entry for its connected host:

```
jncie@vMX3> show arp interface irb.110

MAC Address       Address         Name                    Interface               Flags
00:0c:29:b9:e9:5f 192.168.1.3     192.168.1.3             irb.110 [xe-0/0/5.110]  none
```

This ARP entry is advertised by the vMX3 router as an EVPN MAC/IP route:

```
jncie@vMX3> show route advertising-protocol bgp 172.16.1.1 detail

JNCIE-BLUE-EVPN.evpn.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
...

* 2:172.16.1.3:11::110::00:0c:29:b9:e9:5f::192.168.1.3/304 MAC/IP (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.16.1.3:11
     Route Label: 374768
     ESI: 00:00:00:00:00:00:00:00:00:00
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [65000] I
     Communities: target:110:110

...
```

Every PE router advertises MAC/IP or MAC routes for connected hosts. This reduces ARP flooding and optimizes traffic by preventing traffic instabilities.

Task 4.13: EVPN JNCIE-GREEN

- The green CEs connected to vMX1 and vMX2 should be placed into an EVPN called ‘JNCIE-GREEN-EVPN’.
- Configure an MPLS L3VPN called ‘JNCIE-GREEN’ for this customer.
- Configure the following gateways on the PEs:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

PE	Gateway
vMX1	192.168.1.254/24
vMX2	192.168.1.254/24
vMX3	192.168.2.254/24
vMX4	192.168.3.254/24

Solution

vMX1:

```
set interfaces xe-0/0/5 unit 120 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 120 vlan-id 120

set interfaces irb unit 120 family inet address 192.168.1.254/24
```

We configure an MPLS L3VPN in the same way we would have without the presence of any EVPN:

```
set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface irb.120
set routing-instances JNCIE-GREEN vrf-target target:1000:120
set routing-instances JNCIE-GREEN vrf-table-label
```

The same goes for the EVPN. We can configure it as though the MPLS L3VPN is absent. In this example, we stick to the same VLAN-based EVPN we have been configuring throughout this chapter:

```
set routing-instances JNCIE-GREEN-EVPN instance-type evpn
set routing-instances JNCIE-GREEN-EVPN vlan-id 120
set routing-instances JNCIE-GREEN-EVPN interface xe-0/0/5.120
set routing-instances JNCIE-GREEN-EVPN routing-interface irb.120
set routing-instances JNCIE-GREEN-EVPN vrf-target target:120:120
set routing-instances JNCIE-GREEN-EVPN protocols evpn
```

vMX2:

```
set interfaces xe-0/0/5 unit 120 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 120 vlan-id 120

set interfaces irb unit 120 family inet address 192.168.1.254/24

set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface irb.120
set routing-instances JNCIE-GREEN vrf-target target:1000:120
set routing-instances JNCIE-GREEN vrf-table-label

set routing-instances JNCIE-GREEN-EVPN instance-type evpn
set routing-instances JNCIE-GREEN-EVPN vlan-id 120
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances JNCIE-GREEN-EVPN interface xe-0/0/5.120
set routing-instances JNCIE-GREEN-EVPN routing-interface irb.120
set routing-instances JNCIE-GREEN-EVPN vrf-target target:120:120
set routing-instances JNCIE-GREEN-EVPN protocols evpn
```

vMX3:

```
set interfaces xe-0/0/5 unit 120 vlan-id 120
set interfaces xe-0/0/5 unit 120 family inet address 192.168.2.254/24

set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface xe-0/0/5.120
set routing-instances JNCIE-GREEN vrf-target target:1000:120
set routing-instances JNCIE-GREEN vrf-table-label
```

vMX4:

```
set interfaces xe-0/0/5 unit 120 vlan-id 120
set interfaces xe-0/0/5 unit 120 family inet address 192.168.3.254/24

set routing-instances JNCIE-GREEN instance-type vrf
set routing-instances JNCIE-GREEN interface xe-0/0/5.120
set routing-instances JNCIE-GREEN vrf-target target:1000:120
set routing-instances JNCIE-GREEN vrf-table-label
```

Verification

First, we verify the EVPN. On vMX1, we should see 1 neighbor, the vMX2 router:

```
jncie@vMX1> show evpn instance JNCIE-GREEN-EVPN extensive

Instance: JNCIE-GREEN-EVPN
  Route Distinguisher: 172.16.1.1:13
  VLAN ID: 120
  Per-instance MAC route label: 446688
  Per-instance multicast route label: 446672
  MAC database status                       Local  Remote
    MAC advertisements:                         2       2
    MAC+IP advertisements:                      2       2
    Default gateway MAC advertisements:         1       1
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.120    00:00:00:00:00:00:00:00:00:00      single-homed     Up
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.120         120             Up     JNCIE-GREEN
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    120                  1    1     irb.120    Extended        Enabled   446672
  Number of neighbors: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    172.16.1.2
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           2
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
  Number of ethernet segments: 0

jncie@vMX1> show route table JNCIE-GREEN-EVPN.evpn.0

JNCIE-GREEN-EVPN.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:172.16.1.1:13::120::00:05:86:72:fb:f0/304 MAC/IP
                   *[EVPN/170] 00:09:25
                      Indirect
2:172.16.1.1:13::120::00:0c:29:b9:e9:4b/304 MAC/IP
                   *[EVPN/170] 00:09:22
                      Indirect
2:172.16.1.2:13::120::00:05:86:94:9b:f0/304 MAC/IP
                   *[BGP/170] 00:01:30, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
2:172.16.1.2:13::120::00:0c:29:b9:e9:55/304 MAC/IP
                   *[BGP/170] 00:01:30, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
2:172.16.1.1:13::120::00:05:86:72:fb:f0::192.168.1.254/304 MAC/IP
                   *[EVPN/170] 00:09:25
                      Indirect
2:172.16.1.1:13::120::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
                   *[EVPN/170] 00:09:22
                      Indirect
2:172.16.1.2:13::120::00:05:86:94:9b:f0::192.168.1.254/304 MAC/IP
                   *[BGP/170] 00:01:30, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
2:172.16.1.2:13::120::00:0c:29:b9:e9:55::192.168.1.2/304 MAC/IP
                   *[BGP/170] 00:01:30, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
3:172.16.1.1:13::120::172.16.1.1/304 IM
                   *[EVPN/170] 00:09:24
                      Indirect
3:172.16.1.2:13::120::172.16.1.2/304 IM
                   *[BGP/170] 00:01:30, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
```

Next, we check the routing table for the MPLS L3VPN:

```
jncie@vMX1> show route table JNCIE-GREEN.inet.0

JNCIE-GREEN.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.0/24     *[Direct/0] 00:09:15
                    > via irb.120
                    [BGP/170] 00:01:19, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
192.168.1.1/32     *[EVPN/7] 00:09:11
                    > via irb.120
192.168.1.2/32     *[EVPN/7] 00:01:18
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                    [BGP/170] 00:01:19, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
192.168.1.254/32   *[Local/0] 00:09:15
                      Local via irb.120
                    [EVPN/7] 00:01:18
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
192.168.2.0/24     *[BGP/170] 00:01:19, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
192.168.3.0/24     *[BGP/170] 00:01:19, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
```

Notice how the routing table is populated with both EVPN as well as 'normal' IPv4 routes towards the remote PEs.

The vMX router advertises prefixes for the EVPN as well as the MPLS L3VPN to other PE routers:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.2

JNCIE-GREEN.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                         100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 192.168.1.1/32          Self                          100        I

...

JNCIE-GREEN-EVPN.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:13::120::00:05:86:72:fb:f0/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:05:86:72:fb:f0::192.168.1.254/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                         Self                          100        I
  3:172.16.1.1:13::120::172.16.1.1/304 IM
*                         Self                          100        I

jncie@vMX1> show route advertising-protocol bgp 172.16.1.3

JNCIE-GREEN.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                          100        I
* 192.168.1.1/32          Self                          100        I

...

JNCIE-GREEN-EVPN.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:13::120::00:05:86:72:fb:f0/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:05:86:72:fb:f0::192.168.1.254/304 MAC/IP
*                         Self                          100        I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                         Self                          100        I
  3:172.16.1.1:13::120::172.16.1.1/304 IM
*                         Self                          100        I

jncie@vMX1> show route advertising-protocol bgp 172.16.1.4

JNCIE-GREEN.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                          100        I
* 192.168.1.1/32          Self                          100        I

...

JNCIE-GREEN-EVPN.evpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:13::120::00:05:86:72:fb:f0/304 MAC/IP
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*                          Self                                100          I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b/304 MAC/IP
*                          Self                                100          I
  2:172.16.1.1:13::120::00:05:86:72:fb:f0::192.168.1.254/304 MAC/IP
*                          Self                                100          I
  2:172.16.1.1:13::120::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                          Self                                100          I
  3:172.16.1.1:13::120::172.16.1.1/304 IM
*                          Self                                100          I
```

The last check is to verify IP connectivity between the different CE devices:

```
jncie@vMX5:CE> ping rapid routing-instance CE1-GREEN-EVPN 192.168.1.2

PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.738/15.083/17.895/1.866 ms

jncie@vMX5:CE> ping rapid routing-instance CE1-GREEN-EVPN 192.168.2.1

PING 192.168.2.1 (192.168.2.1): 56 data bytes
!!!!!
--- 192.168.2.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 24.784/54.443/162.588/54.159 ms

jncie@vMX5:CE> ping rapid routing-instance CE1-GREEN-EVPN 192.168.3.1

PING 192.168.3.1 (192.168.3.1): 56 data bytes
!!!!!
--- 192.168.3.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 32.852/219.683/334.483/101.076 ms
```

Task 4.14: EVPN JNCIE-RED

- The red CEs connected to vMX1, vMX2 and vMX3 should be placed into an EVPN called ‘JNCIE-RED-EVPN’.
- Configure a distributed layer 3 gateway for this EVPN on vMX1, vMX2 and vMX3:

Gateway	Real IP	Virtual IP
vMX1	192.168.1.251/24	192.168.1.254
vMX2	192.168.1.252/24	192.168.1.254
vMX3	192.168.1.253/24	192.168.1.254

- Configure a layer 3 gateway for this customer on vMX4. Use address 192.168.2.254/24 and VLAN-ID 132.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Configure an MPLS L3VPN called 'JNCIE-RED' for this customer.
- For the sites where distributed gateways are used, do not advertise the EVPN gateway addresses.

Solution

vMX1:

```
set interfaces xe-0/0/5 unit 130 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 130 vlan-id 130
```

To provide hosts in the EVPN with a distributed layer 3 gateway offered with the same MAC address, we configure the IRB interface with the 'virtual-gateway-address' knob. This knob will make the vMX offer the gateway using a VRRP-MAC address. This way, (virtual) hosts can migrate from one PE to the other without noticing a change in their gateway:

```
set interfaces irb unit 130 family inet address 192.168.1.251/24 virtual-gateway-address
192.168.1.254

set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface irb.130
set routing-instances JNCIE-RED vrf-target target:1000:130
set routing-instances JNCIE-RED vrf-table-label

set routing-instances JNCIE-RED-EVPN instance-type evpn
set routing-instances JNCIE-RED-EVPN vlan-id 130
set routing-instances JNCIE-RED-EVPN interface xe-0/0/5.130
set routing-instances JNCIE-RED-EVPN routing-interface irb.130
set routing-instances JNCIE-RED-EVPN vrf-target target:130:130
```

For this task, the gateway address should not be advertised to other PE routers. This knob is found under the [protocols evpn] stanza:

```
set routing-instances JNCIE-RED-EVPN protocols evpn default-gateway do-not-advertise
```

vMX2:

```
set interfaces xe-0/0/5 unit 130 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 130 vlan-id 130

set interfaces irb unit 130 family inet address 192.168.1.252/24 virtual-gateway-address
192.168.1.254

set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface irb.130
set routing-instances JNCIE-RED vrf-target target:1000:130
set routing-instances JNCIE-RED vrf-table-label
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances JNCIE-RED-EVPN instance-type evpn
set routing-instances JNCIE-RED-EVPN vlan-id 130
set routing-instances JNCIE-RED-EVPN interface xe-0/0/5.130
set routing-instances JNCIE-RED-EVPN routing-interface irb.130
set routing-instances JNCIE-RED-EVPN vrf-target target:130:130
set routing-instances JNCIE-RED-EVPN protocols evpn default-gateway do-not-advertise
```

vMX3:

```
set interfaces xe-0/0/5 unit 131 encapsulation vlan-bridge
set interfaces xe-0/0/5 unit 131 vlan-id 131
set interfaces irb unit 131 family inet address 192.168.1.253/24 virtual-gateway-address
192.168.1.254

set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface irb.131
set routing-instances JNCIE-RED vrf-target target:1000:130
set routing-instances JNCIE-RED vrf-table-label

set routing-instances JNCIE-RED-EVPN instance-type evpn
```

On vMX3, the PE-CE link is configured with a VLAN that differs from vMX1 and vMX2. We do need to establish layer 2 communications between hosts inside the EVPN. To this end, we simply configure the EVPN instance on vMX3 with the same VLAN-ID that was used on vMX1 and vMX2:

```
set routing-instances JNCIE-RED-EVPN vlan-id 130

set routing-instances JNCIE-RED-EVPN interface xe-0/0/5.131
set routing-instances JNCIE-RED-EVPN routing-interface irb.131
set routing-instances JNCIE-RED-EVPN vrf-target target:130:130
set routing-instances JNCIE-RED-EVPN protocols evpn default-gateway do-not-advertise
```

vMX4:

```
set interfaces xe-0/0/5 unit 132 vlan-id 132
set interfaces xe-0/0/5 unit 132 family inet address 192.168.2.254/24

set routing-instances JNCIE-RED instance-type vrf
set routing-instances JNCIE-RED interface xe-0/0/5.132
set routing-instances JNCIE-RED vrf-target target:1000:130
set routing-instances JNCIE-RED vrf-table-label
```

Verification

We verify that the EVPN is established across all three PEs:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX1> show evpn instance JNCIE-RED-EVPN extensive

Instance: JNCIE-RED-EVPN
  Route Distinguisher: 172.16.1.1:15
  VLAN ID: 130
  Per-instance MAC route label: 436688
  Per-instance multicast route label: 436672
  MAC database status                         Local  Remote
    MAC advertisements:                           2       2
    MAC+IP advertisements:                        2       2
    Default gateway MAC advertisements:           1       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.130    00:00:00:00:00:00:00:00:00:00      single-homed     Up
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.130         130             Up     JNCIE-RED
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    130                  1    1     irb.130    Extended        Enabled   436672
  Number of neighbors: 2
    172.16.1.2
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          1
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
    172.16.1.3
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          1
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
  Number of ethernet segments: 0
```

Notice how, in the following command output, the gateway address is absent from the EVPN routing table:

```
jncie@vMX1> show route table JNCIE-RED-EVPN.evpn.0

JNCIE-RED-EVPN.evpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:172.16.1.1:15::130::00:0c:29:b9:e9:4b/304 MAC/IP
                   *[EVPN/170] 00:07:13
                      Indirect
2:172.16.1.2:15::130::00:0c:29:b9:e9:55/304 MAC/IP
                   *[BGP/170] 00:00:06, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
2:172.16.1.3:15::130::00:0c:29:b9:e9:5f/304 MAC/IP
                   *[BGP/170] 00:01:00, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
2:172.16.1.1:15::130::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
                   *[EVPN/170] 00:07:13
                      Indirect
2:172.16.1.2:15::130::00:0c:29:b9:e9:55::192.168.1.2/304 MAC/IP
                   *[BGP/170] 00:00:06, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
2:172.16.1.3:15::130::00:0c:29:b9:e9:5f::192.168.1.3/304 MAC/IP
                   *[BGP/170] 00:01:00, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
3:172.16.1.1:15::130::172.16.1.1/304 IM
                   *[EVPN/170] 00:20:34
                      Indirect
3:172.16.1.2:15::130::172.16.1.2/304 IM
                   *[BGP/170] 00:00:06, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
3:172.16.1.3:15::130::172.16.1.3/304 IM
                   *[BGP/170] 00:01:03, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3

jncie@vMX1> show route table JNCIE-RED.inet.0

JNCIE-RED.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.0/24     *[Direct/0] 00:02:19
                    > via irb.130
                    [BGP/170] 00:02:09, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                    [BGP/170] 00:01:53, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
192.168.1.1/32     *[EVPN/7] 00:02:06
                    > via irb.130
192.168.1.2/32     *[EVPN/7] 00:02:01
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                    [BGP/170] 00:02:02, localpref 100, from 172.16.1.2
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx2
192.168.1.3/32      *[BGP/170] 00:01:32, localpref 100, from 172.16.1.3
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx3
192.168.1.251/32    *[Local/0] 00:02:19
                      Local via irb.130
192.168.2.0/24      *[BGP/170] 00:01:52, localpref 100, from 172.16.1.4
                      AS path: I, validation-state: unverified
                    > to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
                      to 172.16.0.14 via xe-0/0/5.2, label-switched-path vmx1-to-vmx4
```

Due to the '**default-gateway do-not-advertise**' knob, the vMX1 router is no longer advertising the IP address that is configured for the IRB interface. The prefixes advertised for the EVPN are the locally configured subnet as well as locally attached hosts:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.2
...
JNCIE-RED.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                         100        I
* 192.168.1.1/32          Self                         100        I
...
JNCIE-RED-EVPN.evpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:15::130::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::130::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                         Self                         100        I
  3:172.16.1.1:15::130::172.16.1.1/304 IM
*                         Self                         100        I
```

On vMX3, we changed the VLAN-ID through routing-instance configuration. We can observe this by examining the following command output:

```
jncie@vMX3> show interfaces xe-0/0/5.131

  Logical interface xe-0/0/5.131 (Index 352) (SNMP ifIndex 584)
    Flags: Up SNMP-Traps 0x20004000 VLAN-Tag [ 0x8100.131 ] In(swap .130) Out(swap .131)
Encapsulation: VLAN-Bridge
    Input packets : 10
    Output packets: 14
    Protocol bridge, MTU: 1522
      Flags: Is-Primary
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show evpn instance JNCIE-RED-EVPN extensive

Instance: JNCIE-RED-EVPN
  Route Distinguisher: 172.16.1.3:15
  VLAN ID: 130
  Per-instance MAC route label: 424640
  Per-instance multicast route label: 424624
  MAC database status                         Local  Remote
    MAC advertisements:                           2       2
    MAC+IP advertisements:                        2       2
    Default gateway MAC advertisements:           1       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.131    00:00:00:00:00:00:00:00:00:00  single-homed     Up
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.131         130             Up     JNCIE-RED
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    130                  1    1     irb.131    Extended        Enabled   424624
  Number of neighbors: 2
    172.16.1.1
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          1
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
    172.16.1.2
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          1
        Inclusive multicast:                   1
        Ethernet auto-discovery:               0
  Number of ethernet segments: 0
```

To verify IP-connectivity, we can issue a ping from one of the CE devices:

```
jncie@vMX5:CE> ping routing-instance CE1-RED-EVPN rapid 192.168.1.3

PING 192.168.1.3 (192.168.1.3): 56 data bytes
!!!!!
--- 192.168.1.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 29.832/64.617/79.486/18.241 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 3: MPLS L3VPN and VLAN-Aware EVPN.

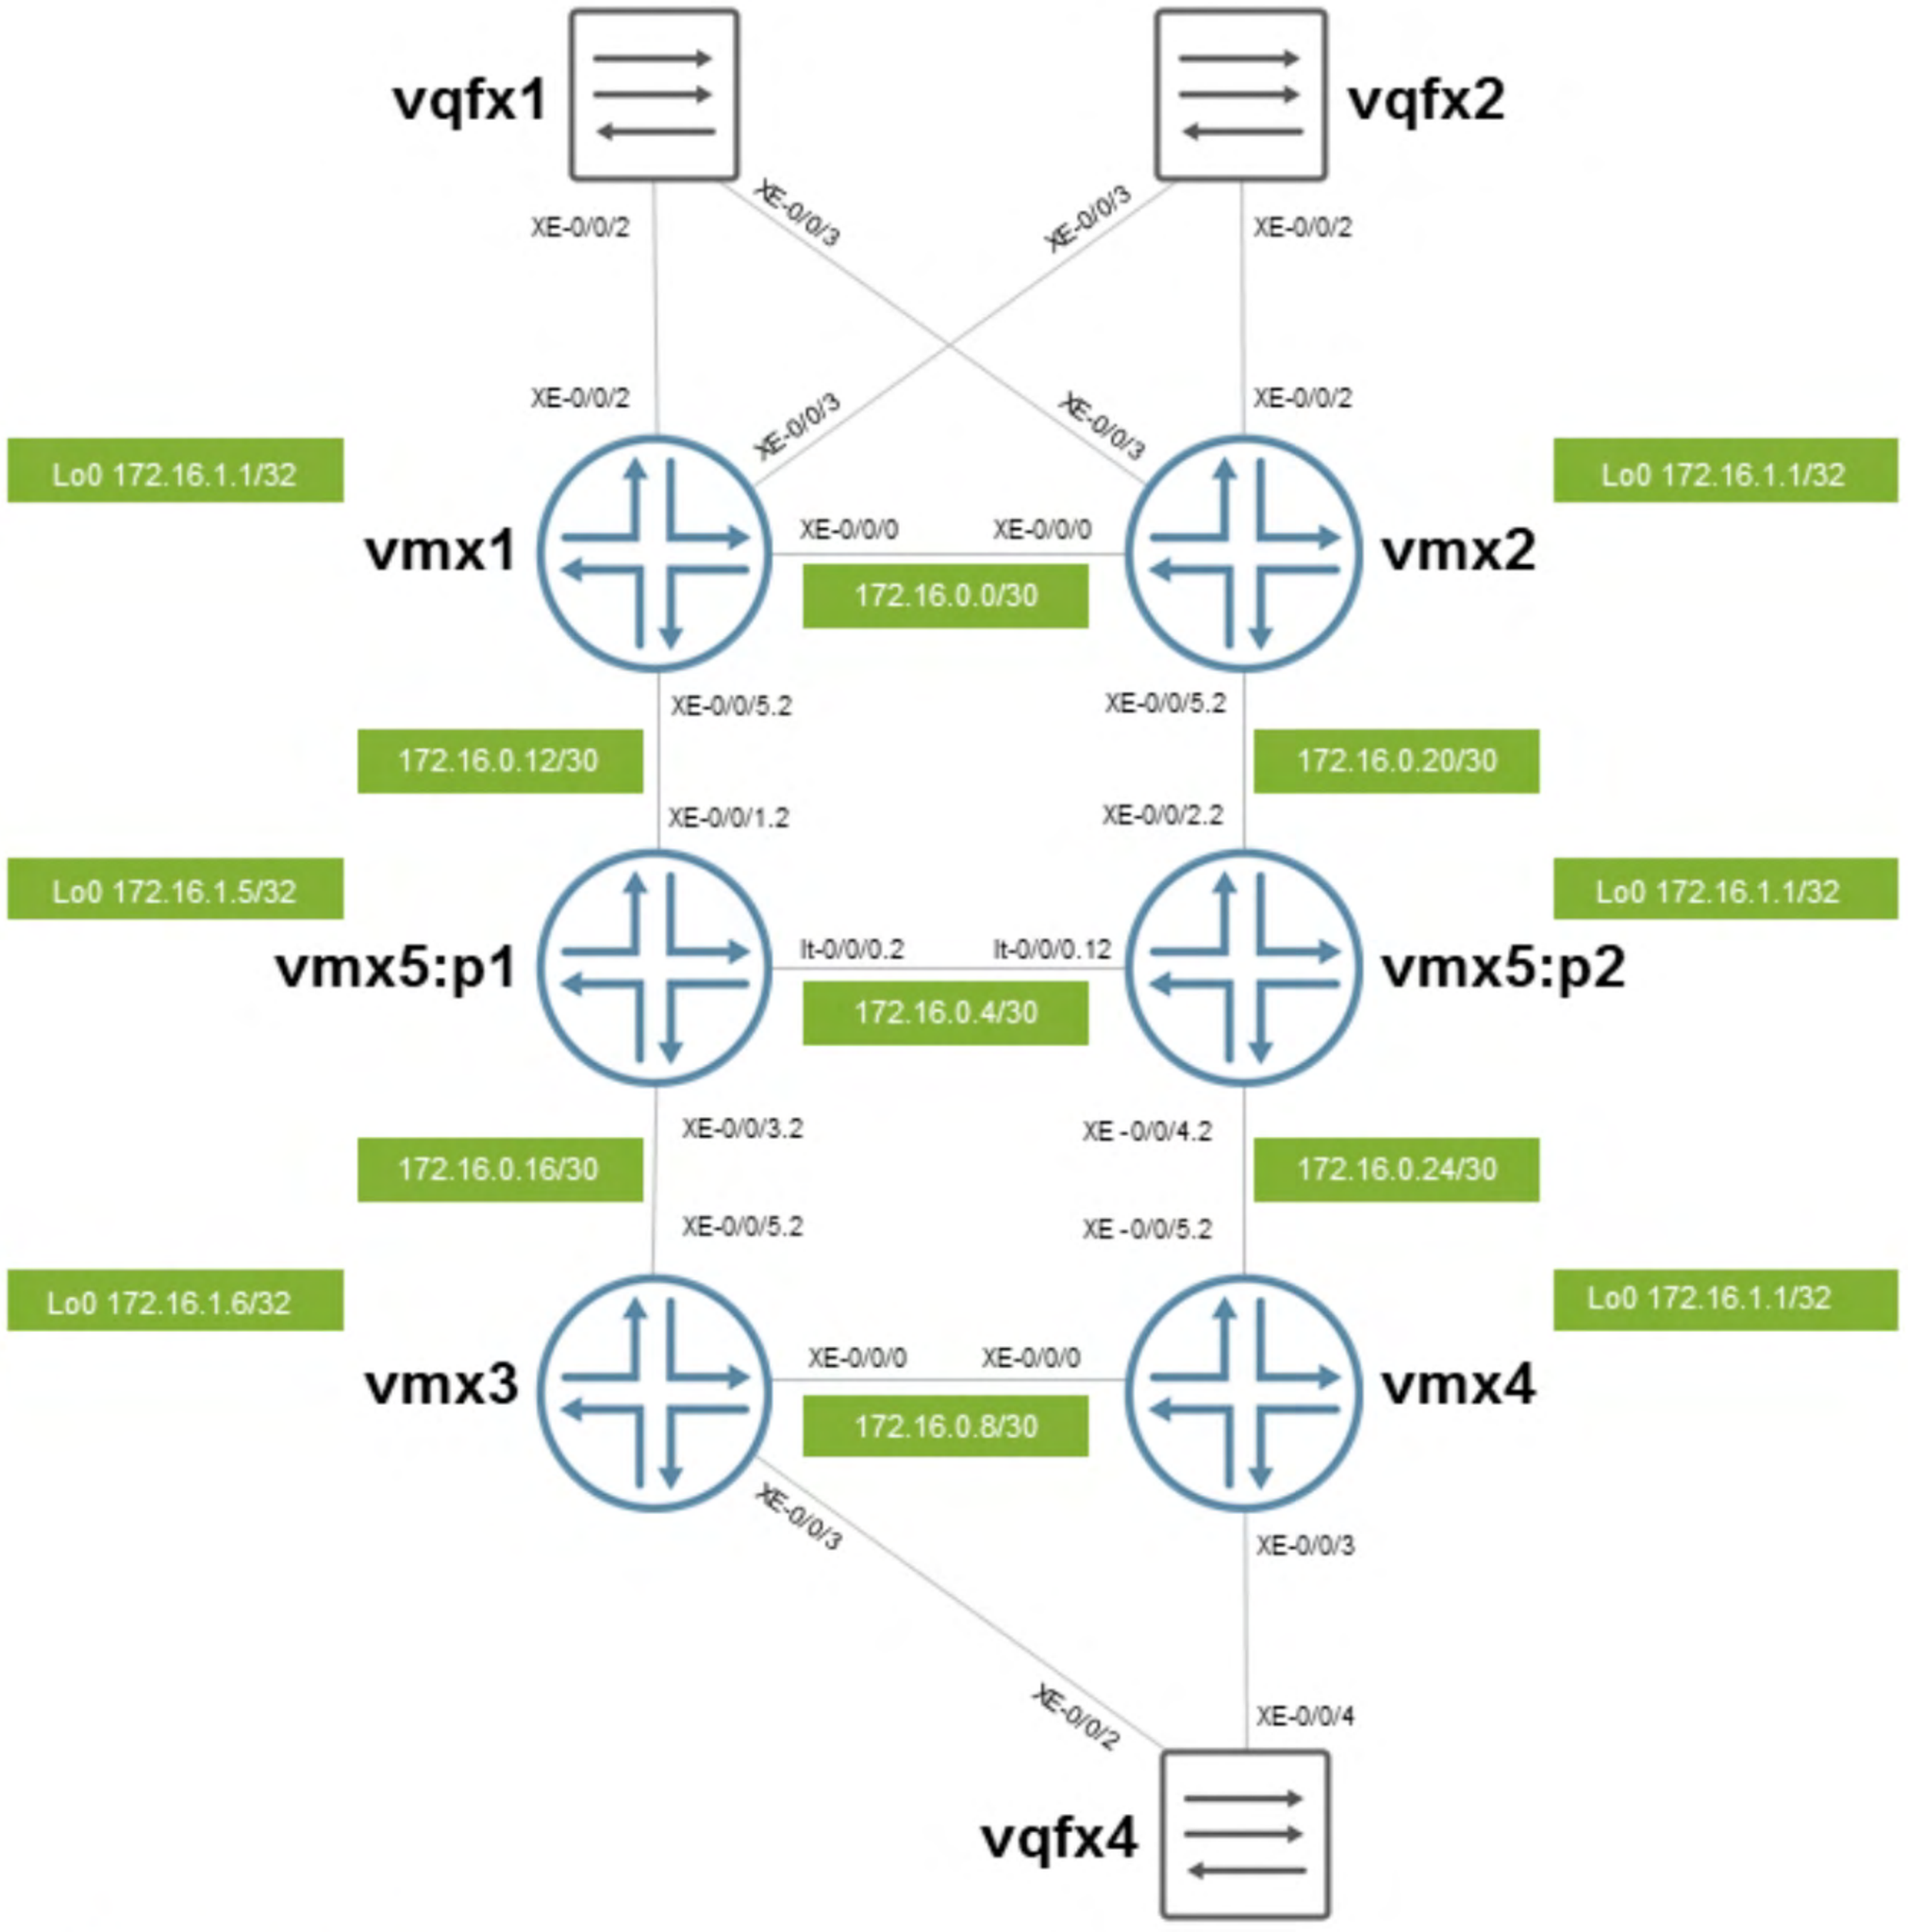

Figure 24 - Physical topology

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

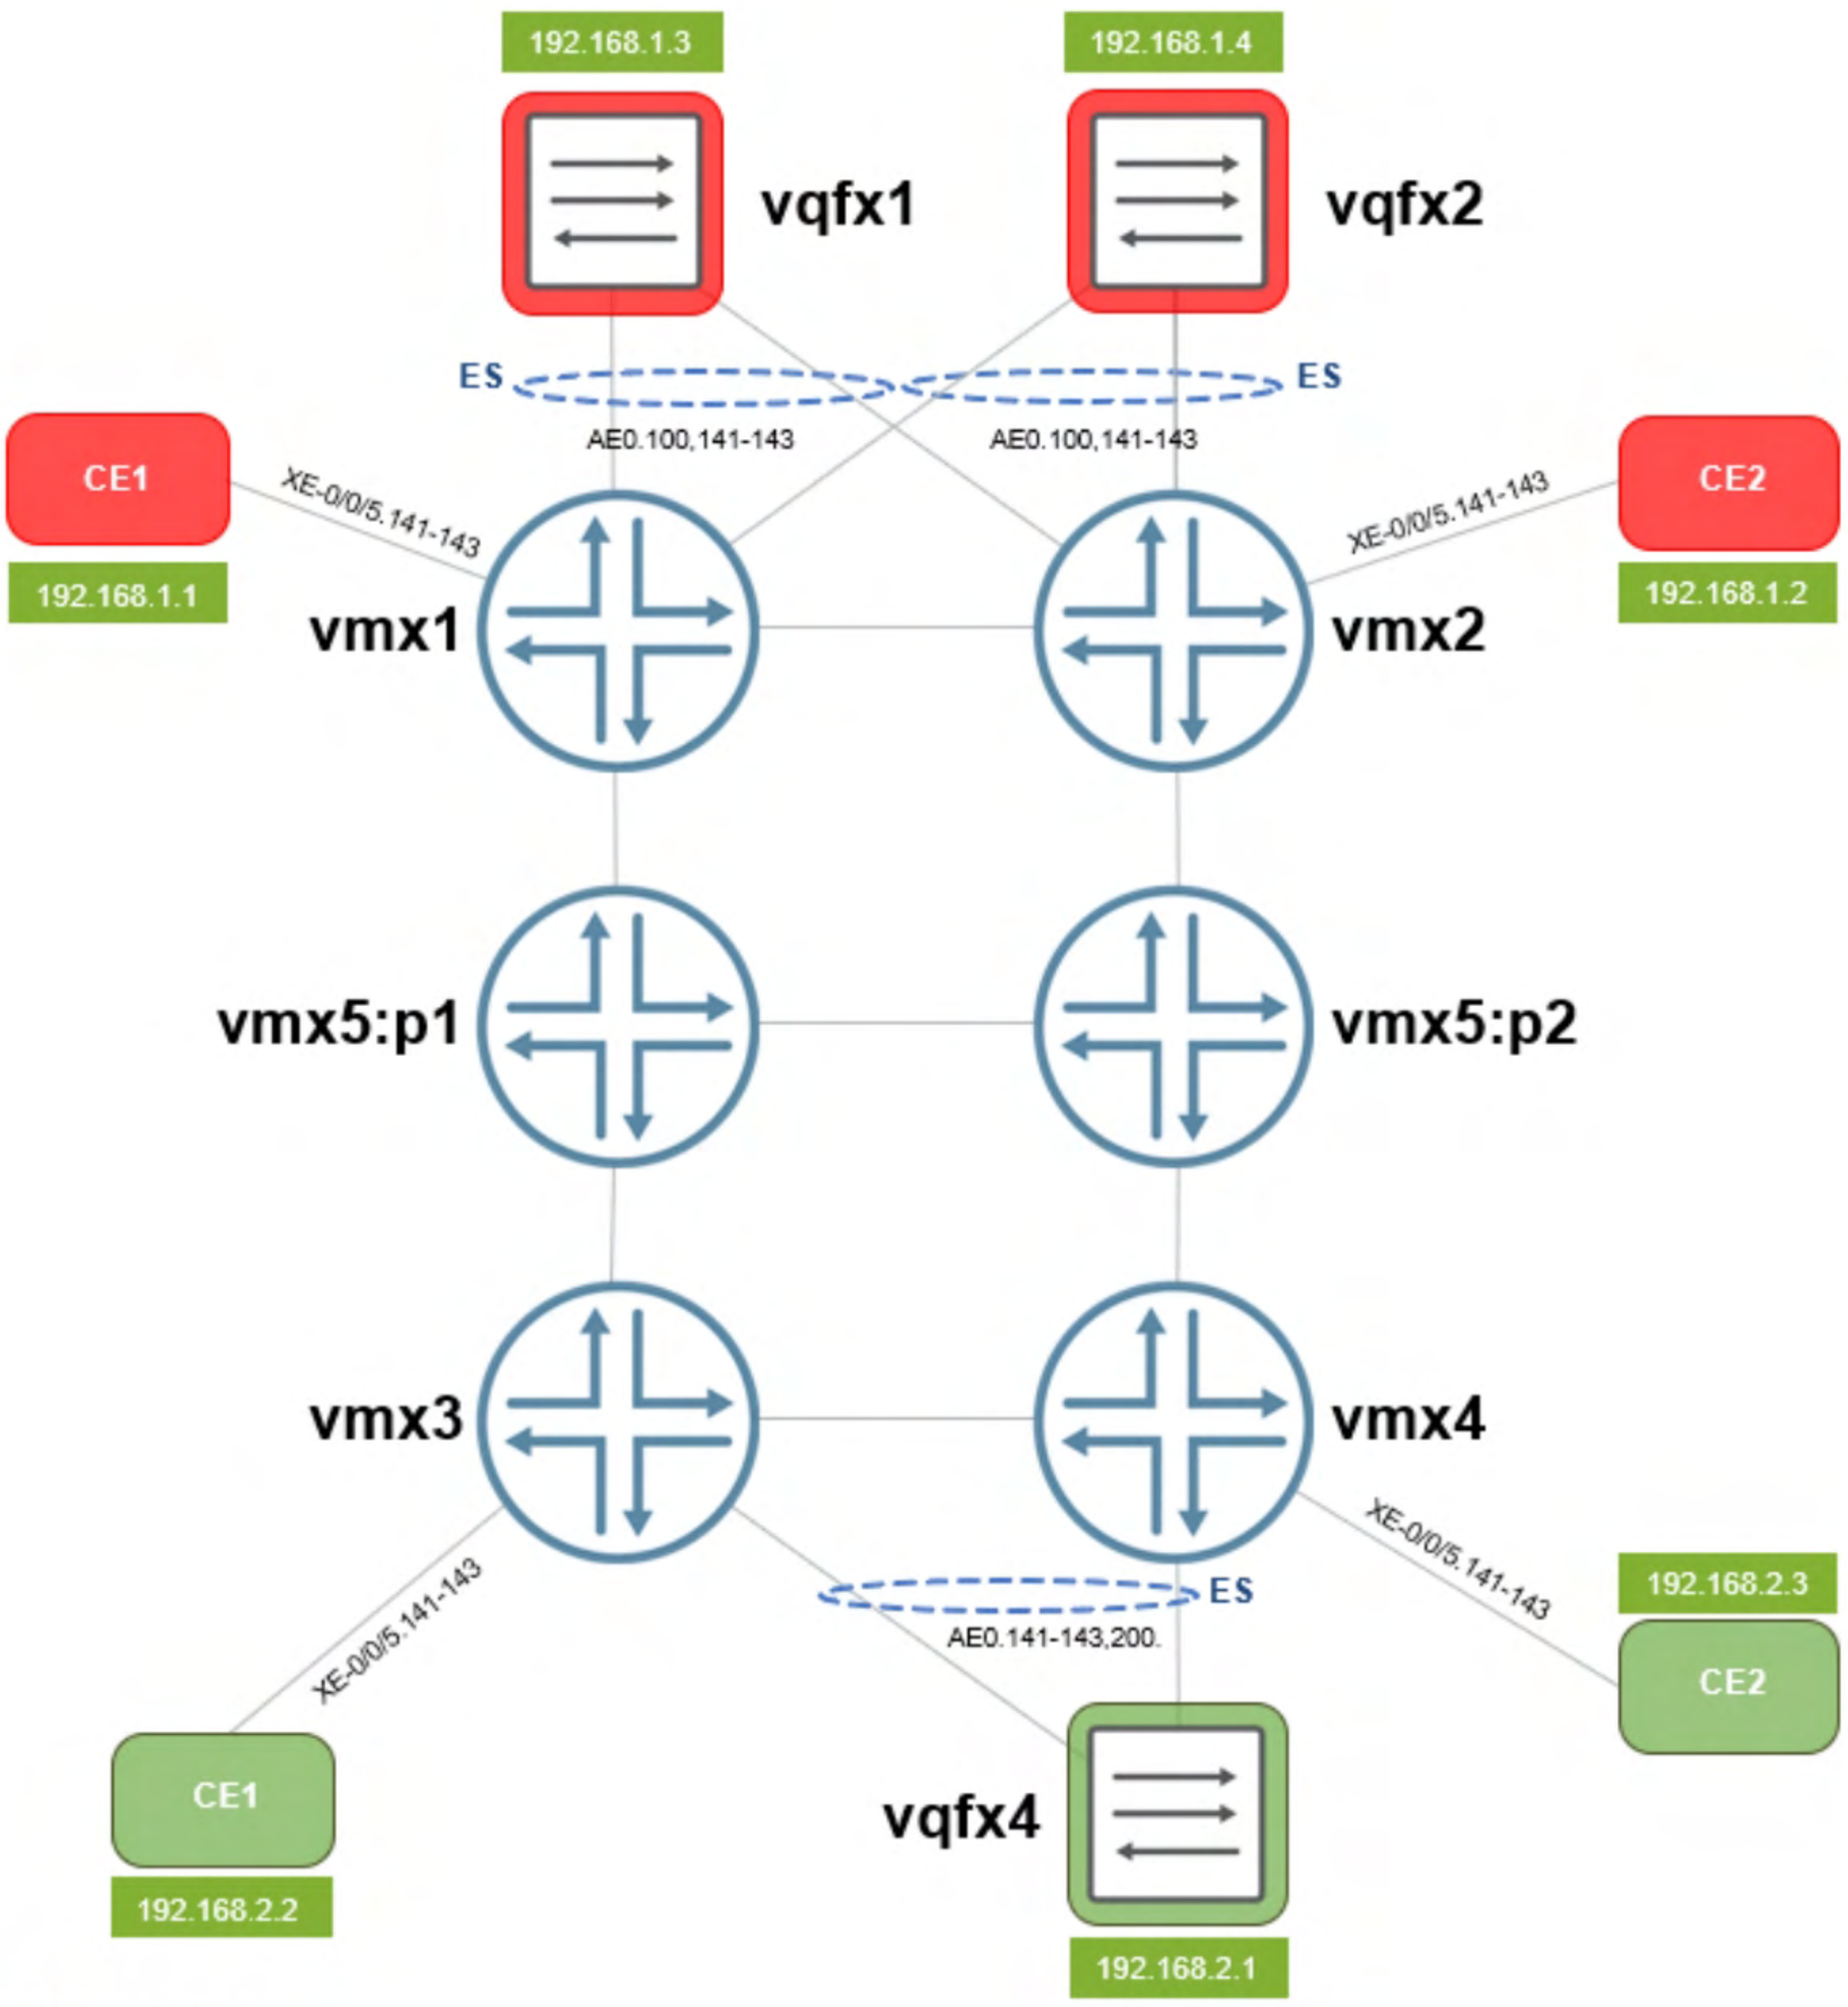

Figure 2 - VPNs

Note: starting topology is configured IP-addressing and LSPs between PE routers.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.16: EVPN North Configuration

- Configure an EVPN called 'jncie-north' on vMX1 and vMX2 for CE1 and CE2.
- Configure a VLAN-aware EVPN using the following parameters:

Gateway	IP address vMX1	IP address vMX2	Gateway address
VLAN 141	192.168.1.251/24	192.168.1.252/24	192.168.1.254
VLAN 142	192.168.1.251/24	192.168.1.252/24	192.168.1.254
VLAN 143			192.168.1.254

- Ensure IP connectivity between hosts inside VLAN 141 and all of the vMX routers.
- Separate layer 3 traffic for VLAN 142 and 143 through the use of MPLS L3VPN.

Solution

vMX1:

The configuration of the trunk interface:

```
set interfaces xe-0/0/5 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/5 unit 0 family bridge vlan-id-list 141-143
```

The requirements of the task were to configure a single IP on IRB.143. Therefore, we configure a virtual-gateway-address on IRB.141 and IRB.142:

```
set interfaces irb unit 141 family inet address 192.168.1.251/24 virtual-gateway-address
192.168.1.254
set interfaces irb unit 142 family inet address 192.168.1.251/24 virtual-gateway-address
192.168.1.254
```

On IRB.143, we configure the exact same IP address and a MAC address as we will configure on vMX2 later on:

```
set interfaces irb unit 143 family inet address 192.168.1.254/24
set interfaces irb unit 143 mac aa:aa:aa:aa:aa:aa
```

We need to advertise the subnet configured on IRB.141 to the rest of the network. In this example, we simply inject the prefix into OSPF:

```
set protocols ospf area 0.0.0.0 interface irb.141 passive
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The configuration of the VLAN-aware EVPN means we need to use 'instance-type virtual-switch':

```
set routing-instances jncie-north instance-type virtual-switch
set routing-instances jncie-north interface xe-0/0/5.0
set routing-instances jncie-north vrf-target target:65000:1
```

The VLAN-aware EVPN also needs to be configured with bridge-domains and protocols EVPN.

```
set routing-instances jncie-north protocols evpn extended-vlan-list 141-143
set routing-instances jncie-north bridge-domains bd-141 vlan-id 141
set routing-instances jncie-north bridge-domains bd-141 routing-interface irb.141
set routing-instances jncie-north bridge-domains bd-142 domain-type bridge
set routing-instances jncie-north bridge-domains bd-142 vlan-id 142
set routing-instances jncie-north bridge-domains bd-142 routing-interface irb.142
set routing-instances jncie-north bridge-domains bd-143 domain-type bridge
set routing-instances jncie-north bridge-domains bd-143 vlan-id 143
set routing-instances jncie-north bridge-domains bd-143 routing-interface irb.143
```

Another requirement for this task was to separate layer 3 traffic for VLAN 142 and 143 through the use of MPLS L3VPN. We configure two L3VPNs with a unique route-target each:

```
set routing-instances ie-142 instance-type vrf
set routing-instances ie-142 interface irb.142
set routing-instances ie-142 vrf-target target:65000:142
set routing-instances ie-142 vrf-table-label

set routing-instances ie-143 instance-type vrf
set routing-instances ie-143 interface irb.143
set routing-instances ie-143 vrf-target target:65000:143
set routing-instances ie-143 vrf-table-label
```

When we configure the other routers, we need to bear in mind what route-targets we choose. We use the same route-target for the EVPN spanning vMX1 and vMX2. Another route-target is used for the EVPN spanning vMX3 and vMX4. We use 1 route-target per MPLS L3VPN, but they have to be consistent on all four PE routers.

vMX2:

```
set interfaces xe-0/0/5 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/5 unit 0 family bridge vlan-id-list 141-143

set interfaces irb unit 141 family inet address 192.168.1.252/24 virtual-gateway-address
192.168.1.254
set interfaces irb unit 142 family inet address 192.168.1.252/24 virtual-gateway-address
192.168.1.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces irb unit 143 family inet address 192.168.1.254/24
set interfaces irb unit 143 mac aa:aa:aa:aa:aa:aa

set protocols ospf area 0.0.0.0 interface irb.141 passive

set routing-instances jncie-north instance-type virtual-switch
set routing-instances jncie-north interface xe-0/0/5.0
set routing-instances jncie-north vrf-target target:65000:1
set routing-instances jncie-north protocols evpn extended-vlan-list 141-143
set routing-instances jncie-north bridge-domains bd-141 vlan-id 141
set routing-instances jncie-north bridge-domains bd-141 routing-interface irb.141
set routing-instances jncie-north bridge-domains bd-142 domain-type bridge
set routing-instances jncie-north bridge-domains bd-142 vlan-id 142
set routing-instances jncie-north bridge-domains bd-142 routing-interface irb.142
set routing-instances jncie-north bridge-domains bd-143 domain-type bridge
set routing-instances jncie-north bridge-domains bd-143 vlan-id 143
set routing-instances jncie-north bridge-domains bd-143 routing-interface irb.143

set routing-instances ie-142 instance-type vrf
set routing-instances ie-142 interface irb.142
set routing-instances ie-142 vrf-target target:65000:142
set routing-instances ie-142 vrf-table-label

set routing-instances ie-143 instance-type vrf
set routing-instances ie-143 interface irb.143
set routing-instances ie-143 vrf-target target:65000:143
set routing-instances ie-143 vrf-table-label
```

Verification

To start the verification of the EVPN, we start out examining the following command output:

```
jncie@vMX1> show evpn instance jncie-north extensive

Instance: jncie-north
  Route Distinguisher: 172.16.1.1:15
  Per-instance MAC route label: 325424
  Per-instance multicast route label: 325408
  MAC database status                         Local  Remote
    MAC advertisements:                           5       4
    MAC+IP advertisements:                        4       4
    Default gateway MAC advertisements:           3       3
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                 Mode             Status
    xe-0/0/5.0      00:00:00:00:00:00:00:00:00:00       single-homed     Up
  Number of IRB interfaces: 3 (3 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.141         141             Up     master
    irb.142         142             Up     ie-142
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  irb.143             143              Up      ie-143
Number of bridge domains: 3
  VLAN  Domain ID    Intfs / up    IRB intf   Mode            MAC sync  IM route label
  141                   1    1     irb.141    Extended        Enabled   325408
  142                   1    1     irb.142    Extended        Enabled   325408
  143                   1    1     irb.143    Extended        Enabled   325408
Number of neighbors: 1
  172.16.1.2
    Received routes
      MAC address advertisement:                4
      MAC+IP address advertisement:             4
      Inclusive multicast:                      3
      Ethernet auto-discovery:                  0
    Number of ethernet segments: 0
```

We see that vMX1 has 1 active single-homed interface that is carrying 3 bridge-domains with a single IRB interface each. We also see that IRB.141 is associated with the main table. The other IRB interfaces are associated with VPNs ‘ie-142’ and ‘ie-143’. We also see that vMX1 has discovered 1 neighbor for this EVPN, vMX2.

The IRB interfaces for VLAN 142 and VLAN 143 were placed inside routing-instances that had ‘vrf-target’ configured. The configuration of ‘vrf-target’ will have the PE (automatically) create an import and export policy. As soon as we extend the VPN to other PE routers (as we will see in subsequent tasks) the subnets configured on the IRB interfaces will be advertised to the other PE routers. The IRB interface for VLAN 141 is not associated with any VPN. It is associated with the main table instead. We configured the IRB as a passive interface under the [protocols ospf area 0.0.0.0] stanza to have the PE router advertise the subnet to the rest of the network. To verify that that both PE1 as well as PE2 are advertising the subnet, we issue the following command:

```
jncie@vMX1> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000e5f    74  0x22 0xefcd  96
  bits 0x0, link count 6
  id 192.168.1.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 100
...

jncie@vMX2> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.2       172.16.1.2       0x80000da6   103  0x22 0xcc96  96
  bits 0x0, link count 6
  id 192.168.1.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 100
  ...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

To verify the IP-connectivity, we move to the logical-system called CE on vMX5:

```
jncie@vMX5:CE> ping routing-instance jncie-north-CE1-VLAN141 rapid 192.168.1.2

PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.026/12.458/16.570/2.786 ms

jncie@vMX5:CE> ping routing-instance jncie-north-CE1-VLAN142 rapid 192.168.1.2
PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.388/15.509/21.723/3.150 ms

jncie@vMX5:CE> ping routing-instance jncie-north-CE1-VLAN143 rapid 192.168.1.2

PING 192.168.1.2 (192.168.1.2): 56 data bytes
!!!!!
--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 13.558/15.112/17.316/1.380 ms
```

As soon as some traffic is exchanged, the PE routers will see what hosts are attached to the locally configured segment. This information is advertised as a MAC/IP route to the other PE routers. The following command will tell us which EPVN routes are active and which table they are active in. In addition to this, we can also see if the route is towards a locally attached host or to a host that is connected to a remote PE.

```
jncie@vMX1> show route protocol evpn

inet.0: 24 destinations, 29 routes (24 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:03:12
                      > via irb.141
192.168.1.2/32     *[EVPN/7] 00:00:09
                      > to 172.16.0.2 via xe-0/0/0.0, Push 321472
192.168.1.252/32   *[EVPN/7] 00:00:09
                      > to 172.16.0.2 via xe-0/0/0.0, Push 321472

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
192.168.1.1/32     *[EVPN/7] 00:03:09
                    > via irb.142
192.168.1.2/32     *[EVPN/7] 00:00:09
                    > to 172.16.0.2 via xe-0/0/0.0, Push 321472
192.168.1.252/32   *[EVPN/7] 00:00:09
                    > to 172.16.0.2 via xe-0/0/0.0, Push 321472

ie-143.inet.0: 4 destinations, 6 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:03:09
                    > via irb.143
192.168.1.2/32     *[EVPN/7] 00:00:09
                    > to 172.16.0.2 via xe-0/0/0.0, Push 321472
...

jncie@vMX2> show route protocol evpn

inet.0: 24 destinations, 29 routes (24 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:01:44
                      > to 172.16.0.1 via xe-0/0/0.0, Push 325424
192.168.1.2/32     *[EVPN/7] 00:04:45
                      > via irb.141
192.168.1.251/32   *[EVPN/7] 00:01:44
                      > to 172.16.0.1 via xe-0/0/0.0, Push 325424

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:01:44
                    > to 172.16.0.1 via xe-0/0/0.0, Push 325424
192.168.1.2/32     *[EVPN/7] 00:04:43
                    > via irb.142
192.168.1.251/32   *[EVPN/7] 00:01:44
                    > to 172.16.0.1 via xe-0/0/0.0, Push 325424

ie-143.inet.0: 4 destinations, 6 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.1/32     *[EVPN/7] 00:01:44
                    > to 172.16.0.1 via xe-0/0/0.0, Push 325424
192.168.1.2/32     *[EVPN/7] 00:04:43
                    > via irb.143
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Notice in the printout above that locally connected hosts have an EVPN routing entry with a next-hop of the type interface. The EVPN routes towards hosts connected to remote PEs are reachable via an LSP, hence the 'Push' operation.

There are other ways to examine the route exchange between the PE routers. One such way is to examine the advertised and received routes:

```
jncie@vMX1> show route receive-protocol bgp 172.16.1.2

inet.0: 24 destinations, 29 routes (24 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.1.0/24          172.16.1.2                   100        I
  192.168.1.2/32          172.16.1.2                   100        I
...

mpls.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.1.2:13:192.168.1.0/24
*                         172.16.1.2                   100        I
...

bgp.evpn.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.2:15::141::00:05:86:94:9b:f0/304 MAC/IP
*                         172.16.1.2                   100        I
  2:172.16.1.2:15::141::00:0c:29:b9:e9:55/304 MAC/IP
...
jncie-north.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.2:15::141::00:05:86:94:9b:f0/304 MAC/IP
*                         172.16.1.2                   100        I
  2:172.16.1.2:15::141::00:0c:29:b9:e9:55/304 MAC/IP
...

jncie@vMX1> show route advertising-protocol bgp 172.16.1.2

ie-142.inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                         100        I
* 192.168.1.1/32          Self                         100        I

...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie-north.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:15::141::00:05:86:72:fb:f0/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::141::00:0c:29:b9:e9:4b/304 MAC/IP
...
```

In addition to checking the advertised and received routes, we can also examine the EVPN database:

```
jncie@vMX1> show evpn database
Instance: jncie-north
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
141             00:05:86:72:fb:f0  irb.141                        May 09 12:00:55
192.168.1.251
141             00:05:86:94:9b:f0  172.16.1.2                     May 09 12:41:52
192.168.1.252
141             00:0c:29:b9:e9:4b  xe-0/0/5.0                     May 09 13:16:36
```

Task 4.17: EVPN South Configuration

- Configure an EVPN called jncie-south on vMX3 and vMX4 for CE1 and CE2.
- Configure the VLAN-aware EVPN using the following parameters:

Gateway	IP address vMX3	IP address vMX4	Gateway address
VLAN 141	192.168.2.251/24	192.168.2.252/24	192.168.2.254
VLAN 142	192.168.2.251/24	192.168.2.252/24	192.168.2.254
VLAN 143			192.168.2.254

- Ensure IP connectivity between hosts inside VLAN 141 and all of the vMX routers.
- Separate layer 3 traffic for VLAN 142 and 143 through the use of MPLS L3VPN.
- Ensure that hosts inside VLAN 142 have IP-connectivity with hosts inside VLAN 142 that was previously configured for the 'jncie-north' EVPN.
- Ensure that hosts inside VLAN 143 have IP-connectivity with hosts inside VLAN 143 that was previously configured for the 'jncie-north' EVPN.

Solution

vMX3:

```
set interfaces xe-0/0/5 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/5 unit 0 family bridge vlan-id-list 141-143

set interfaces irb unit 141 family inet address 192.168.2.251/24 virtual-gateway-address
192.168.2.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces irb unit 142 family inet address 192.168.2.251/24 virtual-gateway-address
192.168.2.254
set interfaces irb unit 143 family inet address 192.168.2.254/24
set interfaces irb unit 143 mac aa:aa:aa:aa:aa:aa

set protocols ospf area 0.0.0.0 interface irb.141 passive

set routing-instances ie-142 instance-type vrf
set routing-instances ie-142 interface irb.142
set routing-instances ie-142 vrf-target target:65000:142
set routing-instances ie-142 vrf-table-label

set routing-instances ie-143 instance-type vrf
set routing-instances ie-143 interface irb.143
set routing-instances ie-143 vrf-target target:65000:143
set routing-instances ie-143 vrf-table-label

set routing-instances jncie-south instance-type virtual-switch
set routing-instances jncie-south interface xe-0/0/5.0
set routing-instances jncie-south vrf-target target:65000:2
set routing-instances jncie-south protocols evpn extended-vlan-list 141-143

set routing-instances jncie-south bridge-domains bd-141 domain-type bridge
set routing-instances jncie-south bridge-domains bd-141 vlan-id 141
set routing-instances jncie-south bridge-domains bd-141 routing-interface irb.141
set routing-instances jncie-south bridge-domains bd-142 domain-type bridge
set routing-instances jncie-south bridge-domains bd-142 vlan-id 142
set routing-instances jncie-south bridge-domains bd-142 routing-interface irb.142
set routing-instances jncie-south bridge-domains bd-143 domain-type bridge
set routing-instances jncie-south bridge-domains bd-143 vlan-id 143
set routing-instances jncie-south bridge-domains bd-143 routing-interface irb.143
```

vMX4:

```
set interfaces xe-0/0/5 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/5 unit 0 family bridge vlan-id-list 141-143

set interfaces irb unit 141 family inet address 192.168.2.252/24 virtual-gateway-address
192.168.2.254
set interfaces irb unit 142 family inet address 192.168.2.252/24 virtual-gateway-address
192.168.2.254
set interfaces irb unit 143 family inet address 192.168.2.254/24
set interfaces irb unit 143 mac aa:aa:aa:aa:aa:aa

set protocols ospf area 0.0.0.0 interface irb.141 passive

set routing-instances ie-142 instance-type vrf
set routing-instances ie-142 interface irb.142
set routing-instances ie-142 vrf-target target:65000:142
set routing-instances ie-142 vrf-table-label
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances ie-143 instance-type vrf
set routing-instances ie-143 interface irb.143
set routing-instances ie-143 vrf-target target:65000:143
set routing-instances ie-143 vrf-table-label

set routing-instances jncie-south instance-type virtual-switch
set routing-instances jncie-south interface xe-0/0/5.0
set routing-instances jncie-south vrf-target target:65000:2
set routing-instances jncie-south protocols evpn extended-vlan-list 141-143

set routing-instances jncie-south bridge-domains bd-141 domain-type bridge
set routing-instances jncie-south bridge-domains bd-141 vlan-id 141
set routing-instances jncie-south bridge-domains bd-141 routing-interface irb.141
set routing-instances jncie-south bridge-domains bd-142 domain-type bridge
set routing-instances jncie-south bridge-domains bd-142 vlan-id 142
set routing-instances jncie-south bridge-domains bd-142 routing-interface irb.142
set routing-instances jncie-south bridge-domains bd-143 domain-type bridge
set routing-instances jncie-south bridge-domains bd-143 vlan-id 143
set routing-instances jncie-south bridge-domains bd-143 routing-interface irb.143
```

Verification

Since this task is so similar to the previous one, the verification is exactly the same as well. We start out Verifying the EVPN:

```
jncie@vMX3> show evpn instance jncie-south extensive

Instance: jncie-south
  Route Distinguisher: 172.16.1.3:15
  Per-instance MAC route label: 319520
  Per-instance multicast route label: 319504
  MAC database status                       Local  Remote
    MAC advertisements:                         6       6
    MAC+IP advertisements:                      6       6
    Default gateway MAC advertisements:         3       3
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    xe-0/0/5.0      00:00:00:00:00:00:00:00:00:00      single-homed     Up
  Number of IRB interfaces: 3 (3 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.141         141             Up     master
    irb.142         142             Up     ie-142
    irb.143         143             Up     ie-143
  Number of bridge domains: 3
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    141                  1    1     irb.141    Extended        Enabled   319504
    142                  1    1     irb.142    Extended        Enabled   319504
    143                  1    1     irb.143    Extended        Enabled   319504
  Number of neighbors: 1
    172.16.1.4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Received routes
        MAC address advertisement:              6
        MAC+IP address advertisement:           6
        Inclusive multicast:                    3
        Ethernet auto-discovery:                0
    Number of ethernet segments: 0

jncie@vMX4> show evpn instance jncie-south extensive

Instance: jncie-south
  Route Distinguisher: 172.16.1.4:15
  Per-instance MAC route label: 317520
  Per-instance multicast route label: 317504
MAC database status                         Local  Remote
  MAC advertisements:                           6       6
  MAC+IP advertisements:                        6       6
  Default gateway MAC advertisements:           3       3
Number of local interfaces: 1 (1 up)
  Interface name  ESI                                 Mode             Status
  xe-0/0/5.0      00:00:00:00:00:00:00:00:00:00       single-homed     Up
Number of IRB interfaces: 3 (3 up)
  Interface name  VLAN   VNI    Status  L3 context
  irb.141         141             Up     master
  irb.142         142             Up     ie-142
  irb.143         143             Up     ie-143
Number of bridge domains: 3
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  141                  1    1     irb.141    Extended        Enabled   317504
  142                  1    1     irb.142    Extended        Enabled   317504
  143                  1    1     irb.143    Extended        Enabled   317504
Number of neighbors: 1
  172.16.1.3
    Received routes
      MAC address advertisement:              6
      MAC+IP address advertisement:           6
      Inclusive multicast:                    3
      Ethernet auto-discovery:                0
  Number of ethernet segments: 0
```

The subnet configured for IRB.141 need to be advertised and in this case, we will be using OSPF. To verify that both vMX3 as well as vMX4 are advertising the subnet, we issue the following command:

```
jncie@vMX3> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.3       172.16.1.3       0x80000c85   265  0x22 0xf581  96
  bits 0x0, link count 6
  id 192.168.2.0, data 255.255.255.0, Type Stub (3)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Topology count: 0, Default metric: 100
...

jncie@vMX4> show ospf database advertising-router self detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.4       172.16.1.4       0x80000bfc   311  0x22 0x34b9  96
  bits 0x0, link count 6
  id 192.168.2.0, data 255.255.255.0, Type Stub (3)
    Topology count: 0, Default metric: 100
  id 172.16.1.4, data 255.255.255.255, Type Stub (3)
```

Verifying IP connectivity between hosts connected to interfaces placed inside the newly created EVPN:

```
jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN141 rapid 192.168.2.3

PING 192.168.2.3 (192.168.2.3): 56 data bytes
!!!!!
--- 192.168.2.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 9.814/13.228/18.437/3.021 ms

jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN142 rapid 192.168.2.3

PING 192.168.2.3 (192.168.2.3): 56 data bytes
!!!!!
--- 192.168.2.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 41.049/85.350/208.860/62.244 ms

jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN143 rapid 192.168.2.3
PING 192.168.2.3 (192.168.2.3): 56 data bytes
!!!!!
--- 192.168.2.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.719/11.737/22.838/5.733 ms
```

Verifying IP connectivity between hosts connected to 'jncie-north' and 'jncie-south':

```
jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN141 rapid 192.168.1.1

PING 192.168.1.1 (192.168.1.1): 56 data bytes
!!!!!
--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.436/16.305/18.616/1.493 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN142 rapid 192.168.1.1

PING 192.168.1.1 (192.168.1.1): 56 data bytes
!!!!!
--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 53.521/260.316/426.714/133.843 ms

jncie@vMX5:CE> ping routing-instance jncie-south-CE1-VLAN143 rapid 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
!!!!!
--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 12.216/16.328/24.928/4.635 ms
```

Checking the exchange of routes between vMX3 and vMX4:

```
jncie@vMX3> show route advertising-protocol bgp 172.16.1.4

ie-142.inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.2.0/24          Self                         100        I
* 192.168.2.2/32          Self                         100        I

ie-143.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.2.0/24          Self                         100        I
* 192.168.2.2/32          Self                         100        I

jncie-south.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.3:15::141::00:05:86:4d:34:f0/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::141::00:0c:29:b9:e9:5f/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::142::00:05:86:4d:34:f0/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::142::00:0c:29:b9:e9:5f/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::143::00:0c:29:b9:e9:5f/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::141::00:05:86:4d:34:f0::192.168.2.251/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::141::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::142::00:05:86:4d:34:f0::192.168.2.251/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.3:15::142::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*                           Self                                100        I
  2:172.16.1.3:15::143::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
*                           Self                                100        I
  2:172.16.1.3:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*                           Self                                100        I
  3:172.16.1.3:15::141::172.16.1.3/304 IM
*                           Self                                100        I
  3:172.16.1.3:15::142::172.16.1.3/304 IM
*                           Self                                100        I
  3:172.16.1.3:15::143::172.16.1.3/304 IM
*                           Self                                100        I

jncie@vMX3> show route receive-protocol bgp 172.16.1.4

inet.0: 24 destinations, 30 routes (24 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.2.0/24          172.16.1.4                   100        I
  192.168.2.3/32          172.16.1.4                   100        I

ie-143.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.2.0/24          172.16.1.4                   100        I
  192.168.2.3/32          172.16.1.4                   100        I

mpls.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.1.4:13:192.168.2.0/24
*                         172.16.1.4                   100        I
  172.16.1.4:13:192.168.2.3/32
*                         172.16.1.4                   100        I
  172.16.1.4:14:192.168.2.0/24
*                         172.16.1.4                   100        I
  172.16.1.4:14:192.168.2.3/32
*                         172.16.1.4                   100        I

bgp.evpn.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.4:15::141::00:05:86:82:68:f0/304 MAC/IP
*                         172.16.1.4                   100        I
  2:172.16.1.4:15::141::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                   100        I
  2:172.16.1.4:15::142::00:05:86:82:68:f0/304 MAC/IP
*                         172.16.1.4                   100        I
  2:172.16.1.4:15::142::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                   100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  2:172.16.1.4:15::143::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::141::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::141::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*                         172.16.1.4                    100        I
  3:172.16.1.4:15::141::172.16.1.4/304 IM
*                         172.16.1.4                    100        I
  3:172.16.1.4:15::142::172.16.1.4/304 IM
*                         172.16.1.4                    100        I
  3:172.16.1.4:15::143::172.16.1.4/304 IM
*                         172.16.1.4                    100        I

jncie-south.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.4:15::141::00:05:86:82:68:f0/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::141::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:05:86:82:68:f0/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::00:0c:29:b9:e9:69/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::141::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::141::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::142::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*                         172.16.1.4                    100        I
  2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*                         172.16.1.4                    100        I
  3:172.16.1.4:15::141::172.16.1.4/304 IM
*                         172.16.1.4                    100        I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  3:172.16.1.4:15::142::172.16.1.4/304 IM
*                         172.16.1.4                      100        I
  3:172.16.1.4:15::143::172.16.1.4/304 IM
*                         172.16.1.4                      100        I
```

Checking the EVPN database:

```
jncie@vMX3> show evpn database

Instance: jncie-south
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
141             00:05:86:4d:34:f0  irb.141                        May 09 18:27:54
192.168.2.251
141             00:05:86:82:68:f0  172.16.1.4                     May 09 18:28:12
192.168.2.251
141             00:0c:29:b9:e9:5f  xe-0/0/5.0                     May 09 18:40:36  192.168.2.2
141             00:0c:29:b9:e9:69  172.16.1.4                     May 09 18:28:24  192.168.2.3
142             00:05:86:4d:34:f0  irb.142                        May 09 18:27:54
192.168.2.251
142             00:05:86:82:68:f0  172.16.1.4                     May 09 18:28:12
192.168.2.251
142             00:0c:29:b9:e9:5f  xe-0/0/5.0                     May 09 18:40:48  192.168.2.2
142             00:0c:29:b9:e9:69  172.16.1.4                     May 09 18:28:22  192.168.2.3
143             00:0c:29:b9:e9:5f  xe-0/0/5.0                     May 09 18:40:51  192.168.2.2
143             00:0c:29:b9:e9:69  172.16.1.4                     May 09 18:28:17  192.168.2.3
143             aa:aa:aa:aa:aa:aa  irb.143                        May 09 18:27:54
192.168.2.254

jncie@vMX4> show evpn database

Instance: jncie-south
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
141             00:05:86:4d:34:f0  172.16.1.3                     May 09 18:28:08
192.168.2.251
141             00:05:86:82:68:f0  irb.141                        May 09 18:28:11
192.168.2.251
141             00:0c:29:b9:e9:5f  172.16.1.3                     May 09 18:28:08  192.168.2.2
141             00:0c:29:b9:e9:69  xe-0/0/5.0                     May 09 18:45:44  192.168.2.3
142             00:05:86:4d:34:f0  172.16.1.3                     May 09 18:28:08
192.168.2.251
142             00:05:86:82:68:f0  irb.142                        May 09 18:28:11
192.168.2.251
142             00:0c:29:b9:e9:5f  172.16.1.3                     May 09 18:40:49  192.168.2.2
142             00:0c:29:b9:e9:69  xe-0/0/5.0                     May 09 18:46:07  192.168.2.3
143             00:0c:29:b9:e9:5f  172.16.1.3                     May 09 18:28:08  192.168.2.2
143             00:0c:29:b9:e9:69  xe-0/0/5.0                     May 09 18:46:09  192.168.2.3
143             aa:aa:aa:aa:aa:aa  irb.143                        May 09 18:28:11
192.168.2.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The EVPN configured on vMX3 and vXM4 differs from the previously configured EVPN. The MPLS L3VPNs do not, they were configured with the same route-target. When we check the routes vMX3 received from vMX1 and vMX2, we can see the route from the 'ie-141' and the 'ie-142' VPNs:

```
jncie@vMX3> show route receive-protocol bgp 172.16.1.1

inet.0: 24 destinations, 30 routes (24 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          172.16.1.1                   100        I
* 192.168.1.1/32          172.16.1.1                   100        I

ie-143.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          172.16.1.1                   100        I
* 192.168.1.1/32          172.16.1.1                   100        I

mpls.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.1.1:13:192.168.1.0/24
*                         172.16.1.1                   100        I
  172.16.1.1:13:192.168.1.1/32
*                         172.16.1.1                   100        I
  172.16.1.1:14:192.168.1.0/24
*                         172.16.1.1                   100        I
  172.16.1.1:14:192.168.1.1/32
*                         172.16.1.1                   100        I

bgp.evpn.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)

jncie-south.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)

jncie@vMX3> show route receive-protocol bgp 172.16.1.2

inet.0: 24 destinations, 30 routes (24 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.1.0/24          172.16.1.2                   100        I
* 192.168.1.2/32          172.16.1.2                   100        I

ie-143.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Prefix                  Nexthop              MED     Lclpref    AS path
  192.168.1.0/24          172.16.1.2                   100        I
* 192.168.1.2/32          172.16.1.2                   100        I

mpls.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.1.2:13:192.168.1.0/24
*                         172.16.1.2                   100        I
  172.16.1.2:13:192.168.1.2/32
*                         172.16.1.2                   100        I
  172.16.1.2:14:192.168.1.0/24
*                         172.16.1.2                   100        I
  172.16.1.2:14:192.168.1.2/32
*                         172.16.1.2                   100        I
```

When we check vMX1, we can see that the vMX did in fact send the EVPN routes for the 'jncie-north' EVPN. However, the vMX3 router does not have any policy that accepts the routes into its RIB:

```
jncie@vMX1> show route advertising-protocol bgp 172.16.1.3

ie-142.inet.0: 8 destinations, 11 routes (8 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                         100        I
* 192.168.1.1/32          Self                         100        I

ie-143.inet.0: 7 destinations, 10 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 192.168.1.0/24          Self                         100        I
* 192.168.1.1/32          Self                         100        I

jncie-north.evpn.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  2:172.16.1.1:15::141::00:05:86:72:fb:f0/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::141::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::142::00:05:86:72:fb:f0/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::142::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::143::00:0c:29:b9:e9:4b/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::141::00:05:86:72:fb:f0::192.168.1.251/304 MAC/IP
*                         Self                         100        I
  2:172.16.1.1:15::141::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*                          Self                                100          I
  2:172.16.1.1:15::142::00:05:86:72:fb:f0::192.168.1.251/304 MAC/IP
*                          Self                                100          I
  2:172.16.1.1:15::142::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                          Self                                100          I
  2:172.16.1.1:15::143::00:0c:29:b9:e9:4b::192.168.1.1/304 MAC/IP
*                          Self                                100          I
  2:172.16.1.1:15::143::aa:aa:aa:aa:aa:aa::192.168.1.254/304 MAC/IP
*                          Self                                100          I
  3:172.16.1.1:15::141::172.16.1.1/304 IM
*                          Self                                100          I
  3:172.16.1.1:15::142::172.16.1.1/304 IM
*                          Self                                100          I
  3:172.16.1.1:15::143::172.16.1.1/304 IM
*                          Self                                100          I
```

Task 4.18: EVPN Multihoming

- Configure the following Ethernet segments:

	ES 1	ES 2	ES3
vMX1	xe-0/0/2	xe-0/0/3	
vMX2	xe-0/0/3	xe-0/0/2	
vMX3			xe-0/0/3
vMX4			xe-0/0/3

- On vMX1 and vMX2, place the interfaces into the 'jncie-north' EVPN. On vMX3 and vMX4, place the interfaces into the 'jncie-south' EVPN.
- Enable the newly created interfaces for any VLAN that is configured for the EVPN.

Solution

vMX1:

Enabling the device for the required amount of AE interfaces:

```
set chassis aggregated-devices ethernet device-count 2
```

Assigning physical interfaces to the relevant AE interface:

```
set interfaces xe-0/0/2 gigether-options 802.3ad ae0
set interfaces xe-0/0/3 gigether-options 802.3ad ae1
```

We configure the AE interface for ‘flexible-vlan-tagging’ and ‘flexible-ethernet-services’. This way, we remain flexible for future tasks that may or may not require different encapsulation methods:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
```

We need to assign an ESI value to every multihomed segment. This value has to be the same on the interfaces that are a part of the ES and it has to be unique in the network. This means that vMX1 and vMX2 need to have the same ESI value configured for the AE connected to vQFX1 and that value must be different from the values used for the Ethernet Segments that will be used to connect to vQFX2 and vQFX4.

Two other things worth noting: the ESI value must start with '00' and the value cannot be all 0s. This last requirement is because of the fact that non-multihomed Ethernet Segments use an ESI of all 0's.

```
set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
```

We set the mode to all-active:

```
set interfaces ae0 esi all-active
```

Next is the configuration of the aggregated link options. We specify the link speed and we set the LACP mode to active:

```
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
```

We need to have vMX1 and vMX2 present themselves as one and the same device to vQFX1. Therefore, we configure the AE interface on both vMX1 as well as vMX2 with the same LACP system-id value:

```
set interfaces ae0 aggregated-ether-options lacp system-id 11:11:11:11:11:11

set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 141-143
```

The second AE interface requires another ESI value and another LACP system ID:

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:22:22:22:22:22:22:22:22:22
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 22:22:22:22:22:22
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 141-143
```

Lastly, we should not forget to assign the AE interfaces to the EVPN instance:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances jncie-north interface ae0.0
set routing-instances jncie-north interface ae1.0
```

When configuring the other vMX routers with the multihomed Ethernet Segment, the most important consideration will be the ESI value and the LACP system ID.

vMX2:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/3 gigether-options 802.3ad ae0
set interfaces xe-0/0/2 gigether-options 802.3ad ae1

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:11:11:11:11:11:11:11:11:11
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 11:11:11:11:11:11
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 141-143

set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:22:22:22:22:22:22:22:22:22
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 22:22:22:22:22:22
set interfaces ae1 unit 0 family bridge interface-mode trunk
set interfaces ae1 unit 0 family bridge vlan-id-list 141-143

set routing-instances jncie-north interface ae0.0
set routing-instances jncie-north interface ae1.0
```

vMX3:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/3 gigether-options 802.3ad ae0

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:33:33:33:33:33:33:33:33:33
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 33:33:33:33:33:33
set interfaces ae0 unit 0 family bridge interface-mode trunk
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae0 unit 0 family bridge vlan-id-list 141-143

set routing-instances jncie-south interface ae0.0
```

vMX4:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/3 gigether-options 802.3ad ae0

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 esi 00:33:33:33:33:33:33:33:33:33
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 33:33:33:33:33:33
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 141-143

set routing-instances jncie-south interface ae0.0
```

Verification

During the verification of this task, the main focus will be the EVPN multihomed interface that is active on vMX3 and vMX4. Verifying the EVPN multihoming feature can be broken down into smaller pieces.

First, we verify the AE interface configuration:

```
jncie@vMX3> show lacp interfaces

Aggregated interface: ae0
  LACP state:    Role  Exp  Def Dist Col Syn Aggr Timeout Activity
   xe-0/0/3    Actor  No  No  Yes Yes Yes  Yes   Fast  Active
   xe-0/0/3   Partner  No  No  Yes Yes Yes  Yes   Fast  Active
  LACP protocol:    Receive State Transmit State     Mux State
   xe-0/0/3         Current  Fast periodic Collecting distributing

jncie@vMX3> show lacp statistics interfaces

Aggregated interface: ae0
  LACP Statistics:    LACP Rx   LACP Tx  Unknown Rx  Illegal Rx
   xe-0/0/3           4         3          0          0


jncie@vMX4> show lacp interfaces

Aggregated interface: ae0
  LACP state:    Role  Exp  Def Dist Col Syn Aggr Timeout Activity
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
   xe-0/0/3    Actor  No  No  Yes Yes Yes  Yes   Fast  Active
   xe-0/0/3  Partner  No  No  Yes Yes Yes  Yes   Fast  Active
  LACP protocol:    Receive State Transmit State     Mux State
   xe-0/0/3         Current  Fast periodic Collecting distributing

jncie@vMX4> show lacp statistics interfaces

Aggregated interface: ae0
  LACP Statistics:    LACP Rx   LACP Tx  Unknown Rx  Illegal Rx
   xe-0/0/3           4         4        0           0

jncie@vMX3> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up
 Interface index: 151, SNMP ifIndex: 547, Generation: 155
 Link-level type: Flexible-Ethernet, MTU: 1522, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
 Flow control: Disabled
 Pad to minimum frame size: Disabled
 Minimum links needed: 1, Minimum bandwidth needed: 1bps
 Device flags   : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Current address: 00:05:86:4d:34:c2, Hardware address: 00:05:86:4d:34:c2
 Ethernet segment value: 00:33:33:33:33:33:33:33:33:33, Mode: all-active
 Last flapped   : 2017-05-10 03:42:56 UTC (01:16:58 ago)
 Statistics last cleared: Never
 Traffic statistics:
  Input bytes :       4179468         1184 bps
  Output bytes :      4715282         1040 bps
  Input packets:        33853         1 pps
  Output packets:       36680         1 pps
  IPv6 transit statistics:
  Input bytes :             0
  Output bytes :            0
  Input packets:            0
  Output packets:           0
  Dropped traffic statistics due to STP State:
   Input bytes :            0
   Output bytes :           0
   Input packets:           0
   Output packets:          0
  Input errors:
   Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
   Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:    Queued packets Transmitted packets   Dropped packets
   0                 0              0                    0
   1                 0              0                    0
   2                 0              0                    0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 3                0          0          0
Egress queues: 8 supported, 4 in use
Queue counters:    Queued packets Transmitted packets   Dropped packets
 0                0          0          0
 1                0          0          0
 2                0          0          0
 3                0          0          0
Queue number:     Mapped forwarding classes
 0          best-effort
 1          expedited-forwarding
 2          assured-forwarding
 3          network-control
```

```
  Logical interface ae0.0 (Index 338) (SNMP ifIndex 559) (Generation 216)
   Flags: Up SNMP-Traps 0x20024000 Encapsulation: Ethernet-Bridge
   Statistics    Packets    pps     Bytes     bps
   Bundle:
     Input :      286     0     17160      0
     Output:      719     0     33074      0
   Adaptive Statistics:
     Adaptive Adjusts:     0
     Adaptive Scans :     0
     Adaptive Updates:     0
   Link:
    xe-0/0/3.0
     Input :      286     0     17160      0
     Output:      719     0     33074      0

   Aggregate member links: 1

   Marker Statistics:  Marker Rx   Resp Tx  Unknown Rx  Illegal Rx
    xe-0/0/3.0         0      0      0      0
   Protocol bridge, MTU: 1522, Generation: 301, Route table: 15, Mesh Group:
__esi_flood__ae0__, EVPN multi-homed status: Blocking BUM Traffic to ESI,
   EVPN multi-homed ESI Split Horizon Label: 321344
    Flags: Trunk-Mode

  Logical interface ae0.32767 (Index 339) (SNMP ifIndex 562) (Generation 217)
   Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
   Statistics    Packets    pps     Bytes     bps
   Bundle:
     Input :      0     0      0      0
     Output:      0     0      0      0
   Adaptive Statistics:
     Adaptive Adjusts:     0
     Adaptive Scans :     0
     Adaptive Updates:     0
   Link:
    xe-0/0/3.32767
     Input :      0     0      0      0
     Output:      0     0      0      0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  LACP info:     Role   System       System    Port   Port  Port
                priority     identifier  priority  number   key
   xe-0/0/3.32767  Actor    127 33:33:33:33:33:33    127    1    1
   xe-0/0/3.32767 Partner    127 02:05:86:71:7a:00    127    1    1
  LACP Statistics:    LACP Rx   LACP Tx  Unknown Rx  Illegal Rx
   xe-0/0/3.32767         0         0         0         0
  Marker Statistics:  Marker Rx   Resp Tx  Unknown Rx  Illegal Rx
   xe-0/0/3.32767         0         0         0         0
  Protocol multiservice, MTU: Unlimited, Generation: 302, Route table: 0
   Flags: None
   Policer: Input: __default_arp_policer__

jncie@vMX4> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up
 Interface index: 151, SNMP ifIndex: 546, Generation: 155
 Link-level type: Flexible-Ethernet, MTU: 1522, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
 Flow control: Disabled
 Pad to minimum frame size: Disabled
 Minimum links needed: 1, Minimum bandwidth needed: 1bps
 Device flags   : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Current address: 00:05:86:82:68:c2, Hardware address: 00:05:86:82:68:c2
 Ethernet segment value: 00:33:33:33:33:33:33:33:33:33, Mode: all-active
 Last flapped   : 2017-05-10 04:04:05 UTC (00:55:48 ago)
 Statistics last cleared: Never
 Traffic statistics:
  Input bytes :       4183804         1264 bps
  Output bytes :      4775652         1072 bps
  Input packets:        33953         1 pps
  Output packets:       37983         1 pps
  IPv6 transit statistics:
  Input bytes :             0
  Output bytes :            0
  Input packets:            0
  Output packets:           0
 Dropped traffic statistics due to STP State:
  Input bytes :             0
  Output bytes :            0
  Input packets:            0
  Output packets:           0
 Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
 Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
 Ingress queues: 8 supported, 4 in use
 Queue counters:    Queued packets Transmitted packets   Dropped packets
  0                 0              0              0
  1                 0              0              0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 2                 0          0          0
 3                 0          0          0
Egress queues: 8 supported, 4 in use
Queue counters:    Queued packets Transmitted packets   Dropped packets
 0                 0          0          0
 1                 0          0          0
 2                 0          0          0
 3                 0          0          0
Queue number:      Mapped forwarding classes
 0           best-effort
 1           expedited-forwarding
 2           assured-forwarding
 3           network-control

 Logical interface ae0.0 (Index 338) (SNMP ifIndex 553) (Generation 222)
  Flags: Up SNMP-Traps 0x20024000 Encapsulation: Ethernet-Bridge
  Statistics    Packets    pps     Bytes     bps
  Bundle:
    Input :      412     0     24720      0
    Output:     2054     0     97970      0
  Adaptive Statistics:
    Adaptive Adjusts:     0
    Adaptive Scans :      0
    Adaptive Updates:     0
  Link:
   xe-0/0/3.0
    Input :      412     0     24720      0
    Output:     2054     0     97970      0

  Aggregate member links: 1

  Marker Statistics:  Marker Rx   Resp Tx  Unknown Rx  Illegal Rx
   xe-0/0/3.0         0      0      0      0
  Protocol bridge, MTU: 1522, Generation: 309, Route table: 15, Mesh Group: __all_ces__, EVPN
multi-homed status: Forwarding,
  EVPN multi-homed ESI Split Horizon Label: 318976
   Flags: Trunk-Mode

Logical interface ae0.32767 (Index 339) (SNMP ifIndex 560) (Generation 223)
 Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
 Statistics    Packets    pps     Bytes     bps
 Bundle:
   Input :       0     0       0      0
   Output:       0     0       0      0
 Adaptive Statistics:
   Adaptive Adjusts:     0
   Adaptive Scans :      0
   Adaptive Updates:     0
 Link:
  xe-0/0/3.32767
   Input :       0     0       0      0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Output:        0      0       0      0
LACP info:     Role   System       System    Port   Port  Port
             priority     identifier  priority  number   key
 xe-0/0/3.32767  Actor   127 33:33:33:33:33:33    127    1    1
 xe-0/0/3.32767 Partner   127 02:05:86:71:7a:00    127    2    1
LACP Statistics:    LACP Rx   LACP Tx  Unknown Rx  Illegal Rx
 xe-0/0/3.32767         0      0       0      0
Marker Statistics:  Marker Rx   Resp Tx  Unknown Rx  Illegal Rx
 xe-0/0/3.32767         0      0       0      0
Protocol multiservice, MTU: Unlimited, Generation: 310, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__
```

LACP PDUs are exchanged with the vQFX and the AE interface is operational on both vMX routers.

Examining the signaling of EVPN ESI:

```
jncie@vMX3> show route advertising-protocol bgp 172.16.1.4

ie-142.inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
* 192.168.2.0/24     Self              100    I
* 192.168.2.1/32     Self              100    I
* 192.168.2.2/32     Self              100    I

ie-143.inet.0: 9 destinations, 12 routes (9 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
* 192.168.2.0/24     Self              100    I
* 192.168.2.1/32     Self              100    I
* 192.168.2.2/32     Self              100    I

jncie-south.evpn.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
 1:172.16.1.3:15::33333333333333333::0/304 AD/EVI
*              Self              100    I
 2:172.16.1.3:15::141::00:05:86:4d:34:f0/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::141::00:0c:29:b9:e9:5f/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::141::02:05:86:71:7a:00/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::142::00:05:86:4d:34:f0/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::142::00:0c:29:b9:e9:5f/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::142::02:05:86:71:7a:00/304 MAC/IP
*              Self              100    I
 2:172.16.1.3:15::143::00:0c:29:b9:e9:5f/304 MAC/IP
*              Self              100    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 2:172.16.1.3:15::143::02:05:86:71:7a:00/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::141::00:05:86:4d:34:f0::192.168.2.251/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::141::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::141::02:05:86:71:7a:00::192.168.2.1/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::142::00:05:86:4d:34:f0::192.168.2.251/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::142::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::142::02:05:86:71:7a:00::192.168.2.1/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::143::00:0c:29:b9:e9:5f::192.168.2.2/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::143::02:05:86:71:7a:00::192.168.2.1/304 MAC/IP
*              Self                  100    I
 2:172.16.1.3:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*              Self                  100    I
 3:172.16.1.3:15::141::172.16.1.3/304 IM
*              Self                  100    I
 3:172.16.1.3:15::142::172.16.1.3/304 IM
*              Self                  100    I
 3:172.16.1.3:15::143::172.16.1.3/304 IM
*              Self                  100    I

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
 Prefix         Nexthop       MED   Lclpref  AS path
 1:172.16.1.3:0::333333333333333::FFFF:FFFF/304 AD/ESI
*              Self                  100    I
 4:172.16.1.3:0::333333333333333:172.16.1.3/304 ES
*              Self                  100    I
```

```
jncie@vMX3> show route receive-protocol bgp 172.16.1.4

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

ie-142.inet.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
 Prefix         Nexthop       MED   Lclpref  AS path
 192.168.2.0/24     172.16.1.4          100    I
 192.168.2.3/32     172.16.1.4          100    I

ie-143.inet.0: 9 destinations, 12 routes (9 active, 0 holddown, 0 hidden)
 Prefix         Nexthop       MED   Lclpref  AS path
 192.168.2.0/24     172.16.1.4          100    I
 192.168.2.3/32     172.16.1.4          100    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
mpls.0: 27 destinations, 34 routes (27 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
 172.16.1.4:13:192.168.2.0/24
*              172.16.1.4           100    I
 172.16.1.4:13:192.168.2.3/32
*              172.16.1.4           100    I
 172.16.1.4:14:192.168.2.0/24
*              172.16.1.4           100    I
 172.16.1.4:14:192.168.2.3/32
*              172.16.1.4           100    I

bgp.evpn.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
 1:172.16.1.4:0::33333333333333333::FFFF:FFFF/304 AD/ESI
*              172.16.1.4           100    I
 1:172.16.1.4:15::33333333333333333::0/304 AD/EVI
*              172.16.1.4           100    I
 2:172.16.1.4:15::141::00:05:86:82:68:f0/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::141::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::142::00:05:86:82:68:f0/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::142::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::143::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::141::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::141::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::142::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::142::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::143::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4           100    I
 2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*              172.16.1.4           100    I
 3:172.16.1.4:15::141::172.16.1.4/304 IM
*              172.16.1.4           100    I
 3:172.16.1.4:15::142::172.16.1.4/304 IM
*              172.16.1.4           100    I
 3:172.16.1.4:15::143::172.16.1.4/304 IM
*              172.16.1.4           100    I
 4:172.16.1.4:0::33333333333333333:172.16.1.4/304 ES
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
*              172.16.1.4               100    I

jncie-south.evpn.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
 1:172.16.1.4:0::33333333333333333::FFFF:FFFF/304 AD/ESI
*              172.16.1.4               100    I
 1:172.16.1.4:15::33333333333333333::0/304 AD/EVI
*              172.16.1.4               100    I
 2:172.16.1.4:15::141::00:05:86:82:68:f0/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::141::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::142::00:05:86:82:68:f0/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::142::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::143::00:0c:29:b9:e9:69/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::141::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::141::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::142::00:05:86:82:68:f0::192.168.2.251/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::142::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::143::00:0c:29:b9:e9:69::192.168.2.3/304 MAC/IP
*              172.16.1.4               100    I
 2:172.16.1.4:15::143::aa:aa:aa:aa:aa:aa::192.168.2.254/304 MAC/IP
*              172.16.1.4               100    I
 3:172.16.1.4:15::141::172.16.1.4/304 IM
*              172.16.1.4               100    I
 3:172.16.1.4:15::142::172.16.1.4/304 IM
*              172.16.1.4               100    I
 3:172.16.1.4:15::143::172.16.1.4/304 IM
*              172.16.1.4               100    I

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
 Prefix          Nexthop       MED   Lclpref  AS path
 4:172.16.1.4:0::33333333333333333:172.16.1.4/304 ES
*              172.16.1.4               100    I
```

Displaying the routes related to the EVPN multihoming feature:

```
jncie@vMX3> show route receive-protocol bgp 172.16.1.4 table __default_evpn__.evpn.0 extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 4:172.16.1.4:0::33333333333333333:172.16.1.4/304 ES (1 entry, 1 announced)
   Import Accepted
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Route Distinguisher: 172.16.1.4:0
    Nexthop: 172.16.1.4
    Localpref: 100
    AS path: I
    Communities: es-import-target:33-33-33-33-33-33

jncie@vMX3> show route advertising-protocol bgp 172.16.1.4 table __default_evpn__.evpn.0
extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 1:172.16.1.3:0::333333333333333::FFFF:FFFF/304 AD/ESI (1 entry, 1 announced)
 BGP group ibgp type Internal
    Route Distinguisher: 172.16.1.3:0
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:65000:2 esi-label:all-active (label 321568)

* 4:172.16.1.3:0::333333333333333:172.16.1.3/304 ES (1 entry, 1 announced)
 BGP group ibgp type Internal
    Route Distinguisher: 172.16.1.3:0
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: es-import-target:33-33-33-33-33-33

jncie@vMX4> show route receive-protocol bgp 172.16.1.3 table __default_evpn__.evpn.0 extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 4:172.16.1.3:0::333333333333333:172.16.1.3/304 ES (1 entry, 1 announced)
    Import Accepted
    Route Distinguisher: 172.16.1.3:0
    Nexthop: 172.16.1.3
    Localpref: 100
    AS path: I
    Communities: es-import-target:33-33-33-33-33-33

jncie@vMX4> show route advertising-protocol bgp 172.16.1.3 table __default_evpn__.evpn.0
extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 1:172.16.1.4:0::333333333333333::FFFF:FFFF/304 AD/ESI (1 entry, 1 announced)
 BGP group ibgp type Internal
    Route Distinguisher: 172.16.1.4:0
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [65000] I
    Communities: target:65000:2 esi-label:all-active (label 318976)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
* 4:172.16.1.4:0::33333333333333333:172.16.1.4/304 ES (1 entry, 1 announced)
 BGP group ibgp type Internal
   Route Distinguisher: 172.16.1.4:0
   Nexthop: Self
   Flags: Nexthop Change
   Localpref: 100
   AS path: [65000] I
   Communities: es-import-target:33-33-33-33-33-33
```

Both vMX3 as well as vMX4 advertise the routes for the ESI to vMX1 and vMX2 as well.

```
jncie@vMX3> show route advertising-protocol bgp 172.16.1.1 table __default_evpn__.evpn.0
extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 1:172.16.1.3:0::33333333333333333::FFFF:FFFF/304 AD/ESI (1 entry, 1 announced)
 BGP group ibgp type Internal
   Route Distinguisher: 172.16.1.3:0
   Nexthop: Self
   Flags: Nexthop Change
   Localpref: 100
   AS path: [65000] I
   Communities: target:65000:2 esi-label:all-active (label 321568)

* 4:172.16.1.3:0::33333333333333333:172.16.1.3/304 ES (1 entry, 1 announced)
 BGP group ibgp type Internal
   Route Distinguisher: 172.16.1.3:0
   Nexthop: Self
   Flags: Nexthop Change
   Localpref: 100
   AS path: [65000] I
   Communities: es-import-target:33-33-33-33-33-33

jncie@vMX3> show route advertising-protocol bgp 172.16.1.2 table __default_evpn__.evpn.0
extensive

__default_evpn__.evpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
* 1:172.16.1.3:0::33333333333333333::FFFF:FFFF/304 AD/ESI (1 entry, 1 announced)
 BGP group ibgp type Internal
   Route Distinguisher: 172.16.1.3:0
   Nexthop: Self
   Flags: Nexthop Change
   Localpref: 100
   AS path: [65000] I
   Communities: target:65000:2 esi-label:all-active (label 321568)

* 4:172.16.1.3:0::33333333333333333:172.16.1.3/304 ES (1 entry, 1 announced)
 BGP group ibgp type Internal
   Route Distinguisher: 172.16.1.3:0
   Nexthop: Self
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Flags: Nexthop Change
Localpref: 100
AS path: [65000] I
Communities: es-import-target:33-33-33-33-33-33
```

However, vMX1 and vMX2 are configured with two other ESIs. Therefore, they will not import the community associated with the multihomed connection that is active on vMX3 and vXM4:

```
jncie@vMX1> show policy __vrf-import-__default_evpn__-internal__

Policy __vrf-import-__default_evpn__-internal__:
  Term unnamed:
    from community __vrf-community-__default_evpn__-import-internal__ [es-import-target:11-11-11-11-11-11 es-import-target:22-22-22-22-22-22 ]
    then accept
  Term unnamed:
    then reject

jncie@vMX2> show policy __vrf-import-__default_evpn__-internal__

Policy __vrf-import-__default_evpn__-internal__:
  Term unnamed:
    from community __vrf-community-__default_evpn__-import-internal__ [es-import-target:11-11-11-11-11-11 es-import-target:22-22-22-22-22-22 ]
    then accept
  Term unnamed:
    then reject
```

We can see that vMX1 and vMX2 only import the ESIs they have been configured with and will not import routes for the Ethernet Segment that is active on vMX3 and vMX4.

Examining the policies:

```
jncie@vMX3> show policy ?

Possible completions:
 <[Enter]>      Execute this command
 <policy>       Name of policy
 __vrf-export-ie-142-internal__
 __vrf-export-ie-143-internal__
 __vrf-export-jncie-south-internal__
 __vrf-import-__default_evpn__-internal__
 __vrf-import-ie-142-internal__
 __vrf-import-ie-143-internal__
 __vrf-import-jncie-south-internal__
 conditions      Show conditions used by policy
 damping        Show state of route flap damping
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
fabric          Internal fabric state
logical-system    Name of logical system, or 'all'
statistics      Statistics for policy
|           Pipe through a command
```

ES is not accepted in policies related to EVPN 'jncie-north':

```
jncie@vMX3> show policy __vrf-export-jncie-south-internal__

Policy __vrf-export-jncie-south-internal__:
  Term unnamed:
    then community + __vrf-community-jncie-south-common-internal__ [target:65000:2 ] accept
```

```
jncie@vMX3> show policy __vrf-import-jncie-south-internal__

Policy __vrf-import-jncie-south-internal__:
  Term unnamed:
    from community __vrf-community-jncie-south-common-internal__ [target:65000:2 ]
    then accept
  Term unnamed:
    then reject
```

Instead, we find that the ES is accepted in the following policy on both vMX3 as well as vMX4:

```
jncie@vMX3> show policy __vrf-import-__default_evpn__-internal__

Policy __vrf-import-__default_evpn__-internal__:
  Term unnamed:
    from community __vrf-community-__default_evpn__-import-internal__ [es-import-target:33-33-
33-33-33-33 ]
    then accept
  Term unnamed:
    then reject
```

```
jncie@vMX4> show policy __vrf-import-__default_evpn__-internal__

Policy __vrf-import-__default_evpn__-internal__:
  Term unnamed:
    from community __vrf-community-__default_evpn__-import-internal__ [es-import-target:33-33-
33-33-33-33 ]
    then accept
  Term unnamed:
    then reject
```

We can check the bridging table to see if we are learning any MAC addresses:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show bridge mac-table vlan-id 141

MAC flags    (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
  O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : jncie-south
 Bridging domain : bd-141, VLAN : 141
  MAC          MAC   Logical      NH   RTR
  addresssss     flags  interface    Index ID
  00:0c:29:b9:e9:5f  D    xe-0/0/5.0
  00:0c:29:b9:e9:69  DC             1048581 1048581
  02:05:86:71:7a:00  D    ae0.0

jncie@vMX3> show evpn database l2-domain-id 141
Instance: jncie-south
VLAN DomainId MAC address    Active source          Timestamp    IP address
141       00:05:86:4d:34:f0 irb.141            May 09 18:27:54 192.168.2.251
141       00:05:86:82:68:f0 172.16.1.4           May 10 02:39:25 192.168.2.251
141       00:0c:29:b9:e9:5f xe-0/0/5.0           May 10 05:28:27 192.168.2.2
141       00:0c:29:b9:e9:69 172.16.1.4           May 10 04:46:45 192.168.2.3
141       02:05:86:71:7a:00 00:33:33:33:33:33:33:33:33:33 May 10 05:24:56 192.168.2.1

jncie@vMX4> show bridge mac-table vlan-id 141

MAC flags    (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
  O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : jncie-south
 Bridging domain : bd-141, VLAN : 141
  MAC          MAC   Logical      NH   RTR
  addresssss     flags  interface    Index ID
  00:0c:29:b9:e9:5f  DC             1048574 1048574
  00:0c:29:b9:e9:69  D    xe-0/0/5.0
  02:05:86:71:7a:00  D    ae0.0

jncie@vMX4> show evpn database l2-domain-id 141

Instance: jncie-south
VLAN DomainId MAC address    Active source          Timestamp    IP address
141       00:05:86:4d:34:f0 172.16.1.3           May 10 02:39:25 192.168.2.251
141       00:05:86:82:68:f0 irb.141            May 09 18:28:11 192.168.2.251
141       00:0c:29:b9:e9:5f 172.16.1.3           May 10 04:55:27 192.168.2.2
141       00:0c:29:b9:e9:69 xe-0/0/5.0           May 10 05:27:36 192.168.2.3
141       02:05:86:71:7a:00 00:33:33:33:33:33:33:33:33:33 May 10 05:26:46 192.168.2.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4> show interfaces irb | match Har

 Current address: 02:05:86:71:7a:00, Hardware address: 02:05:86:71:7a:00

{master:0}
jncie@vQFX4>

jncie@vMX3> show evpn database extensive l2-domain-id 141 mac-address 02:05:86:71:7a:00
Instance: jncie-south

VLAN ID: 141, MAC address: 02:05:86:71:7a:00
 Source: 00:33:33:33:33:33:33:33:33:33, Rank: 1, Status: Active
  Remote origin: 172.16.1.4
  Timestamp: May 10 05:24:56 (0x5912a428)
  State: <Remote-To-Local-Adv-Done>
  IP address: 192.168.2.1
   Remote origin: 172.16.1.4
   L3 route: 192.168.2.1/32, L3 context: master (irb.141)

jncie@vMX4> show evpn database extensive l2-domain-id 141 mac-address 02:05:86:71:7a:00
Instance: jncie-south

VLAN ID: 141, MAC address: 02:05:86:71:7a:00
 Source: 00:33:33:33:33:33:33:33:33:33, Rank: 1, Status: Active
  Local origin: ae0.0
  Timestamp: May 10 05:26:46 (0x5912a496)
  State: <Local-MAC-Only Local-To-Remote-Adv-Allowed>
  IP address: 192.168.2.1
   Local origin: ae0.0
   L3 route: 192.168.2.1/32, L3 context: master (irb.141)
```

Possibly interesting output from vMX1 and vMX2:

```
jncie@vMX1> show evpn instance jncie-north extensive

Instance: jncie-north
 Route Distinguisher: 172.16.1.1:15
 Per-instance MAC route label: 325424
 Per-instance multicast route label: 325408
 MAC database status              Local Remote
  MAC advertisements:              11    6
  MAC+IP advertisements:             11    6
  Default gateway MAC advertisements:    3    3
 Number of local interfaces: 3 (3 up)
  Interface name ESI                Mode       Status
  ae0.0      00:11:11:11:11:11:11:11:11:11 all-active    Up
  ae1.0      00:22:22:22:22:22:22:22:22:22 all-active    Up
  xe-0/0/5.0   00:00:00:00:00:00:00:00:00:00 single-homed   Up
 Number of IRB interfaces: 3 (3 up)
  Interface name VLAN  VNI  Status L3 context
  irb.141     141      Up   master
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 irb.142     142      Up    ie-142
 irb.143     143      Up    ie-143
Number of bridge domains: 3
 VLAN Domain ID  Intfs / up  IRB intf  Mode       MAC sync IM route label
 141         3   3   irb.141  Extended     Enabled  325408
 142         3   3   irb.142  Extended     Enabled  325408
 143         3   3   irb.143  Extended     Enabled  325408
Number of neighbors: 1
 172.16.1.2
  Received routes
   MAC address advertisement:       6
   MAC+IP address advertisement:      6
   Inclusive multicast:         3
   Ethernet auto-discovery:        4
Number of ethernet segments: 2
 ESI: 00:11:11:11:11:11:11:11:11:11
  Status: Resolved by IFL ae0.0
  Local interface: ae0.0, Status: Up/Forwarding
  Number of remote PEs connected: 1
   Remote PE    MAC label Aliasing label Mode
   172.16.1.2    0     321808     all-active
  Designated forwarder: 172.16.1.2
  Backup forwarder: 172.16.1.1
  Last designated forwarder update: May 10 05:09:55
  Advertised MAC label: 325776
  Advertised aliasing label: 325776
  Advertised split horizon label: 334288
 ESI: 00:22:22:22:22:22:22:22:22:22
  Status: Resolved by IFL ae1.0
  Local interface: ae1.0, Status: Up/Forwarding
  Number of remote PEs connected: 1
   Remote PE    MAC label Aliasing label Mode
   172.16.1.2    0     321824     all-active
  Designated forwarder: 172.16.1.2
  Backup forwarder: 172.16.1.1
  Last designated forwarder update: May 10 05:20:52
  Advertised MAC label: 325792
  Advertised aliasing label: 325792
  Advertised split horizon label: 334576

jncie@vMX2> show evpn instance jncie-north extensive

Instance: jncie-north
 Route Distinguisher: 172.16.1.2:15
 Per-instance MAC route label: 321472
 Per-instance multicast route label: 321456
 MAC database status           Local Remote
  MAC advertisements:            6   11
  MAC+IP advertisements:          6   11
  Default gateway MAC advertisements:    3    3
 Number of local interfaces: 3 (3 up)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 Interface name ESI                Mode       Status
 ae0.0      00:11:11:11:11:11:11:11:11:11 all-active    Up
 ae1.0      00:22:22:22:22:22:22:22:22:22 all-active    Up
 xe-0/0/5.0   00:00:00:00:00:00:00:00:00:00 single-homed   Up
Number of IRB interfaces: 3 (3 up)
 Interface name VLAN  VNI  Status L3 context
 irb.141     141      Up   master
 irb.142     142      Up   ie-142
 irb.143     143      Up   ie-143
Number of bridge domains: 3
 VLAN Domain ID  Intfs / up  IRB intf  Mode       MAC sync IM route label
 141         3  3   irb.141  Extended     Enabled  321456
 142         3  3   irb.142  Extended     Enabled  321456
 143         3  3   irb.143  Extended     Enabled  321456
Number of neighbors: 1
 172.16.1.1
  Received routes
   MAC address advertisement:       11
   MAC+IP address advertisement:     11
   Inclusive multicast:          3
   Ethernet auto-discovery:        4
Number of ethernet segments: 2
 ESI: 00:11:11:11:11:11:11:11:11:11
  Status: Resolved by IFL ae0.0
  Local interface: ae0.0, Status: Up/Forwarding
  Number of remote PEs connected: 1
   Remote PE    MAC label Aliasing label Mode
   172.16.1.1    325776   325776     all-active
  Designated forwarder: 172.16.1.2
  Backup forwarder: 172.16.1.1
  Last designated forwarder update: May 10 05:09:57
  Advertised MAC label: 321808
  Advertised aliasing label: 321808
  Advertised split horizon label: 324784
 ESI: 00:22:22:22:22:22:22:22:22:22
Status: Resolved by IFL ae1.0
Local interface: ae1.0, Status: Up/Forwarding
Number of remote PEs connected: 1
 Remote PE    MAC label Aliasing label Mode
 172.16.1.1    325792   325792     all-active
Designated forwarder: 172.16.1.2
Backup forwarder: 172.16.1.1
Last designated forwarder update: May 10 04:57:27
Advertised MAC label: 321824
Advertised aliasing label: 321824
Advertised split horizon label: 324720
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.19: VLAN-Based EVPN Configuration

- Configure a VLAN-based EVPN called 'ie-north' on vMX1 and vMX2.
- Configure the following IP addresses on vMX1 and vMX2:

PE device	IP address	Virtual Gateway address	VLAN
vMX1	192.168.100.251/24	192.168.100.254/24	100
vMX2	192.168.100.252/24	192.168.100.254/24	100

- Configure a VLAN-based EVPN called 'ie-south' on vMX3 and vMX4.
- Configure the following IP addresses on vMX3 and vMX4:

PE device	IP address	Virtual Gateway address	VLAN
vMX3	192.168.200.251/24	192.168.200.254/24	2
vMX4	192.168.200.252/24	192.168.200.254/24	200

- Configure an MPLS L3VPN named 'jncie-vpn' for all newly configured layer 3 interfaces.

Solution

vMX1:

Since the EVPN will be VLAN-based, we configure the logical interface for 'encapsulation vlan-bridge' as opposed to the 'family bridge' we used when configuring the VLAN-aware EVPN:

```
set interfaces ae0 unit 100 encapsulation vlan-bridge
set interfaces ae0 unit 100 vlan-id 100

set interfaces irb unit 100 family inet address 192.168.100.251/24 virtual-gateway-address
192.168.100.254
```

Apart from setting the instance-type to EVPN, we also need to assign the newly created AE interface to the EVPN. In addition to this, we assign the EVPN a unique route-target:

```
set routing-instances ie-north instance-type evpn
set routing-instances ie-north vlan-id 100
set routing-instances ie-north interface ae0.100
set routing-instances ie-north routing-interface irb.100
set routing-instances ie-north vrf-target target:65000:3
set routing-instances ie-north protocols evpn

set routing-instances jncie-vpn instance-type vrf
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances jncie-vpn interface irb.100
set routing-instances jncie-vpn vrf-target target:65000:100200
set routing-instances jncie-vpn vrf-table-label
```

When we configure the other PE routers, we make sure that the 'ie-north' and 'ie-south' EVPN are configured with a different route-target. The 'jncie' L3VPN is configured with the same route-target on all PE routers.

vMX2:

```
set interfaces ae0 unit 100 encapsulation vlan-bridge
set interfaces ae0 unit 100 vlan-id 100

set interfaces irb unit 100 family inet address 192.168.100.252/24 virtual-gateway-address
192.168.100.254

set routing-instances ie-north instance-type evpn
set routing-instances ie-north vlan-id 100
set routing-instances ie-north interface ae0.100
set routing-instances ie-north routing-interface irb.100
set routing-instances ie-north vrf-target target:65000:3
set routing-instances ie-north protocols evpn

set routing-instances jncie-vpn instance-type vrf
set routing-instances jncie-vpn interface irb.100
set routing-instances jncie-vpn vrf-target target:65000:100200
set routing-instances jncie-vpn vrf-table-label
```

vMX3:

```
set interfaces ae0 unit 200 encapsulation vlan-bridge
set interfaces ae0 unit 200 vlan-id 200

set interfaces irb unit 200 family inet address 192.168.200.251/24 virtual-gateway-address
192.168.200.254

set routing-instances ie-south instance-type evpn
set routing-instances ie-south vlan-id 200
set routing-instances ie-south interface ae0.200
set routing-instances ie-south routing-interface irb.200
set routing-instances ie-south vrf-target target:65000:4
set routing-instances ie-south protocols evpn

set routing-instances jncie-vpn instance-type vrf
set routing-instances jncie-vpn interface irb.200
set routing-instances jncie-vpn vrf-target target:65000:100200
set routing-instances jncie-vpn vrf-table-label
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set interfaces ae0 unit 200 encapsulation vlan-bridge
set interfaces ae0 unit 200 vlan-id 200

set interfaces irb unit 200 family inet address 192.168.200.252/24 virtual-gateway-address
192.168.200.254

set routing-instances ie-south instance-type evpn
set routing-instances ie-south vlan-id 200
set routing-instances ie-south interface ae0.200
set routing-instances ie-south routing-interface irb.200
set routing-instances ie-south vrf-target target:65000:4
set routing-instances ie-south protocols evpn

set routing-instances jncie-vpn instance-type vrf
set routing-instances jncie-vpn interface irb.200
set routing-instances jncie-vpn vrf-target target:65000:100200
set routing-instances jncie-vpn vrf-table-label
```

Verification

We verify IP connectivity from the QFX to both PE routers in the 'ie-north' EVPN:

```
jncie@vQFX1> ping source 192.168.100.1 192.168.100.251 rapid

PING 192.168.100.251 (192.168.100.251): 56 data bytes
!!!!!
--- 192.168.100.251 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.857/11.206/12.014/0.411 ms

{master:0}
jncie@vQFX1> ping source 192.168.100.1 192.168.100.252 rapid

PING 192.168.100.252 (192.168.100.252): 56 data bytes
!!!!!
--- 192.168.100.252 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 8.356/9.022/9.728/0.483 ms
```

Verifying IP connectivity from vQFX1 to vQFX4 using the newly configured 'jncie-vpn' MPLS L3VPN:

```
jncie@vQFX1> ping source 192.168.100.1 192.168.200.1 rapid

PING 192.168.200.1 (192.168.200.1): 56 data bytes
!!!!!
--- 192.168.200.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
round-trip min/avg/max/stddev = 32.304/37.010/47.681/5.500 ms
```

Verifying the EVPN instance on vMX1 and vMX2:

```
jncie@vMX1> show evpn instance ie-north extensive

Instance: ie-north
  Route Distinguisher: 172.16.1.1:17
  VLAN ID: 100
  Per-instance MAC route label: 334720
  Per-instance multicast route label: 334704
  MAC database status                         Local  Remote
  MAC advertisements:                         1      2
  MAC+IP advertisements:                      1      2
  Default gateway MAC advertisements:         1      1
Number of local interfaces: 1 (1 up)
  Interface name  ESI                                Mode             Status
  ae0.100         00:11:11:11:11:11:11:11:11:11      all-active       Up
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN   VNI    Status  L3 context
  irb.100         100             Up     jncie-vpn
Number of bridge domains: 1
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  100                  1    1     irb.100    Extended        Enabled   334704
Number of neighbors: 1
  172.16.1.2
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           2
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved by IFL ae0.100
    Local interface: ae0.100, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.16.1.2       324896     324896          all-active
    Designated forwarder: 172.16.1.1
    Backup forwarder: 172.16.1.2
    Last designated forwarder update: May 10 05:53:17
    Advertised MAC label: 334736
    Advertised aliasing label: 334736
    Advertised split horizon label: 334896

jncie@vMX2> show evpn instance ie-north extensive

Instance: ie-north
  Route Distinguisher: 172.16.1.2:17
  VLAN ID: 100
  Per-instance MAC route label: 324848
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Per-instance multicast route label: 324832
MAC database status                         Local  Remote
  MAC advertisements:                           2       1
  MAC+IP advertisements:                        2       1
  Default gateway MAC advertisements:           1       1
Number of local interfaces: 1 (1 up)
  Interface name  ESI                                Mode             Status
  ae0.100         00:11:11:11:11:11:11:11:11:11      all-active       Up
Number of IRB interfaces: 1 (1 up)
  Interface name  VLAN   VNI    Status  L3 context
  irb.100         100             Up     jncie-vpn
Number of bridge domains: 1
  VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
  100               1    1      irb.100    Extended        Enabled   324832
Number of neighbors: 1
  172.16.1.1
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           1
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
Number of ethernet segments: 1
  ESI: 00:11:11:11:11:11:11:11:11:11
    Status: Resolved by IFL ae0.100
    Local interface: ae0.100, Status: Up/Forwarding
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.16.1.1       334736     334736          all-active
    Designated forwarder: 172.16.1.1
    Backup forwarder: 172.16.1.2
    Last designated forwarder update: May 10 05:53:20
    Advertised MAC label: 324896
    Advertised aliasing label: 324896
    Advertised split horizon label: 325136
```

We issue the same verification command on vMX3 and vMX4.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Part 4: Controllerless Overlay and EVPN DCI

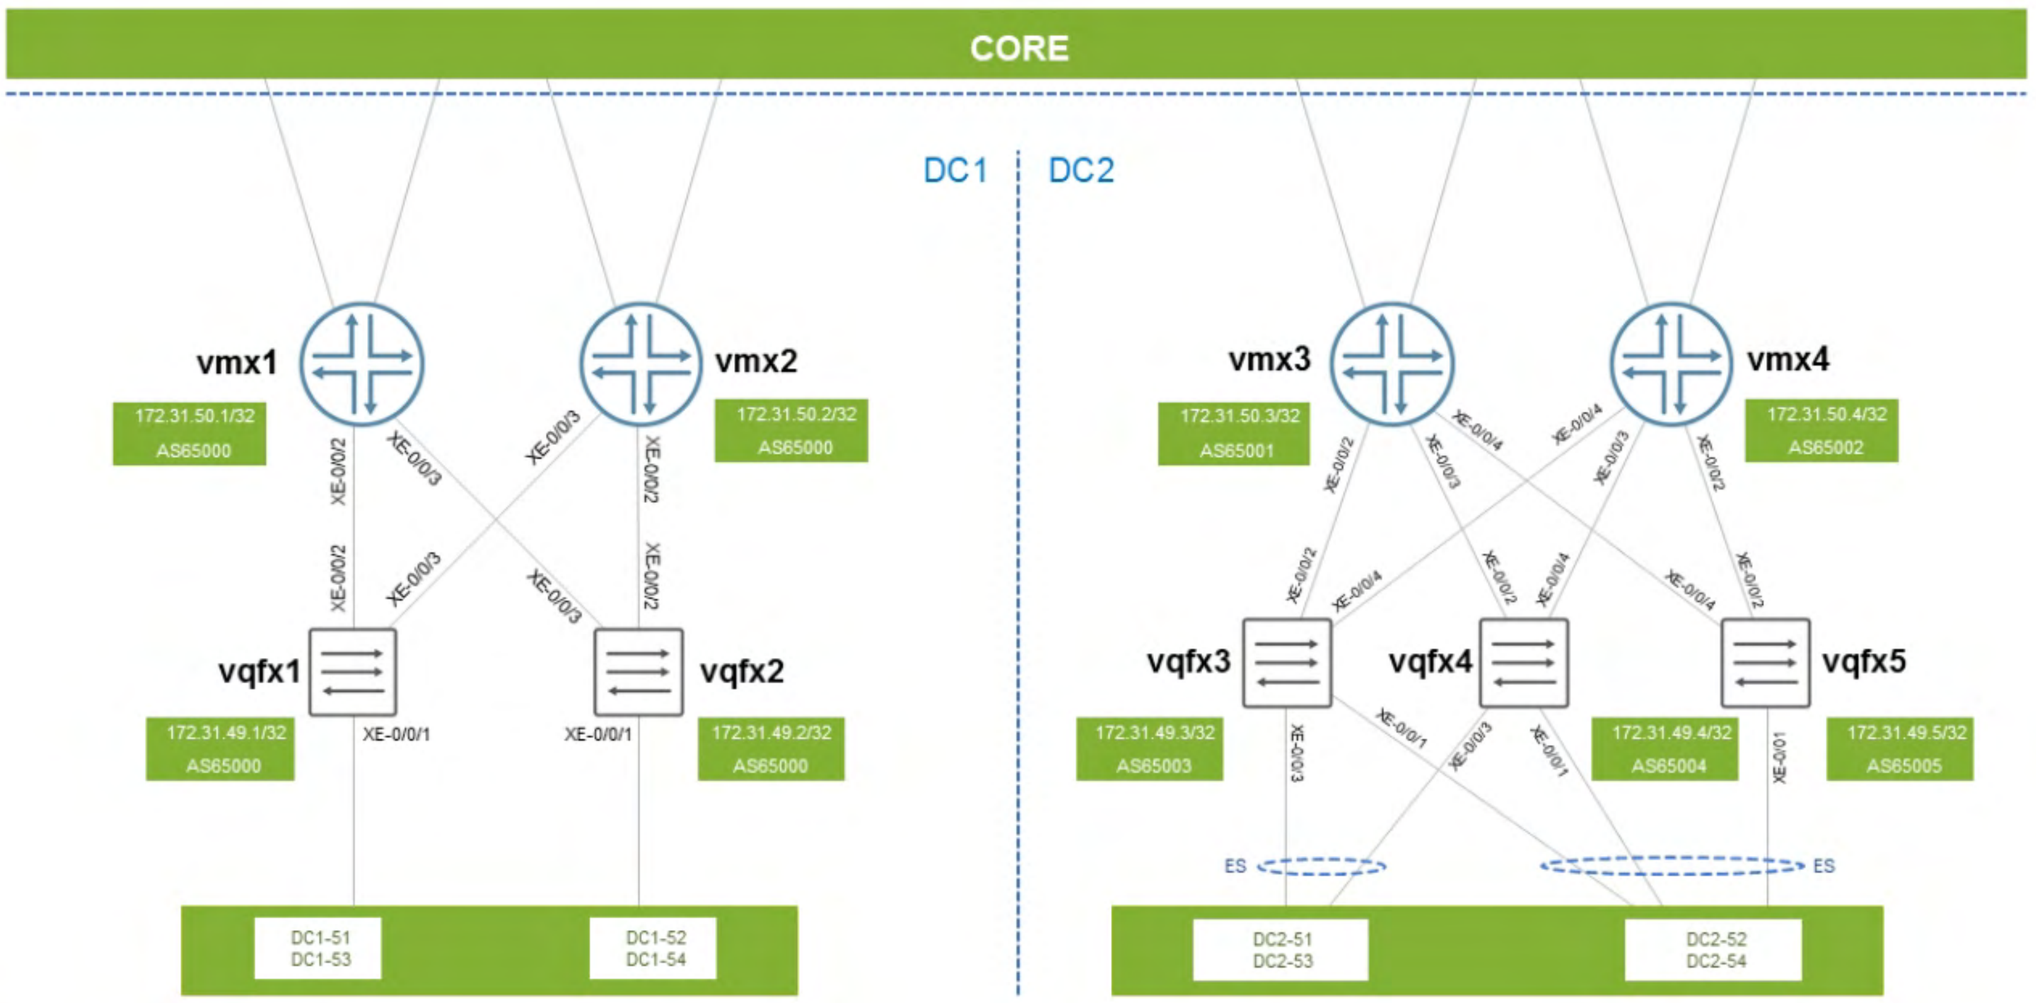

Figure1 –Network topology

Notes:

- Starting topology is configured with IP-addressing.
- Some links also have OSPF and LDP enabled.

Task 4.20: DC1 Controllerless Overlay Part 1: routing and signaling

- Ensure IP-connectivity between the loopback IPs of the devices in DC1 by configuring OSPF on the links between the vQFX and the vMX in DC1.
- Source BGP sessions from the loopback interface and enable BGP for EVPN signaling.
- The vQFX devices in DC1 are only allowed to peer with the vMX routers and should receive each other's EVPN routes.
- Use the AS numbers shown in the topology.
- Enable the devices to load-share traffic across links.

Solution

vMX1:

```
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.50.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 0.0.0.1
set protocols bgp group overlay neighbor 172.31.49.1
set protocols bgp group overlay neighbor 172.31.49.2
set protocols bgp group overlay neighbor 172.31.50.2

set routing-options autonomous-system 65000
```

vMX2:

```
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.50.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 0.0.0.2
set protocols bgp group overlay neighbor 172.31.49.1
set protocols bgp group overlay neighbor 172.31.49.2
set protocols bgp group overlay neighbor 172.31.50.1

set routing-options autonomous-system 65000
```

vQFX1:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.49.1
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay neighbor 172.31.50.1
set protocols bgp group overlay neighbor 172.31.50.2

set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
```

vQFX2:

```
set routing-options autonomous-system 65000

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.49.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay neighbor 172.31.50.1
set protocols bgp group overlay neighbor 172.31.50.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
```

Verification

Check the OSPF neighbor relationships between vQFX and vMX:

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri   Dead
172.16.48.41     xe-0/0/2.0             Full      172.31.49.1      128     35
172.16.48.45     xe-0/0/3.0             Full      172.31.49.2      128     39
172.16.24.34     xe-0/0/5.12            Full      172.31.50.5      128     38
172.16.48.34     xe-0/0/6.12            Full      172.31.50.5      128     35
```

Verify that there is a full mesh of IBGP sessions for the overlay network:

```
jncie@vMX1> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.evpn.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.31.49.1           65000          4          2       0       0          13 Establ
  bgp.evpn.0: 0/0/0/0
172.31.49.2           65000          4          2       0       0           4 Establ
  bgp.evpn.0: 0/0/0/0
172.31.50.2           65000          5          4       0       0        1:23 Establ
  bgp.evpn.0: 0/0/0/0
```

The vMX1 and vMX2 routers should be functioning as route-reflectors:

```
jncie@vMX1> show bgp neighbor 172.31.49.2

Peer: 172.31.49.2+52678 AS 65000 Local: 172.31.50.1+179 AS 65000
  Group: overlay                Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established  (route reflector client)Flags: <Initializing>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: evpn
  Local Address: 172.31.50.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.31.49.2     Local ID: 172.31.50.1       Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
I/O Session Thread: bgpio-0 State: Enabled
BFD: disabled, down
NLRI for restart configured on peer: evpn
NLRI advertised by peer: evpn
NLRI for this session: evpn
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: evpn
NLRI of received end-of-rib markers: evpn
NLRI of all end-of-rib markers sent: evpn
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65000)
Peer does not support Addpath
Table bgp.evpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not in sync
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
  Advertised prefixes:          0
Last traffic (seconds): Received 205589 Sent 100 Checked 205589
Input messages:  Total 7      Updates 1       Refreshes 0     Octets 188
Output messages: Total 4      Updates 0       Refreshes 0     Octets 87
Output Queue[1]: 0            (bgp.evpn.0, evpn)
```

Task 4.21: DC1 Controllerless Overlay Part 2: the overlay

- Configure a VXLAN layer 2 gateway on the vQFX switches for VLAN 200.
- Server-facing interfaces should support tagged traffic. Ensure that the servers can reach the core network.
- The devices should automatically assign RTs for the VXLAN segments.
- Configure the following layer 3 gateway on vMX1 and vMX2:

VLAN	IPv4 gateway	MAC address
200	192.168.200.254	00:00:00:00:02:00

Solution

vMX1:

```
set interfaces irb unit 200 family inet address 192.168.200.254/24
set interfaces irb unit 200 mac 00:00:00:00:02:00

set routing-options route-distinguisher-id 172.31.50.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols ospf area 0.0.0.0 interface irb.200 passive

set routing-instances overlay vtep-source-interface lo0.0
set routing-instances overlay instance-type virtual-switch
set routing-instances overlay route-distinguisher 172.31.50.1:65000
set routing-instances overlay vrf-target target:65000:100
set routing-instances overlay vrf-target auto
set routing-instances overlay protocols evpn encapsulation vxlan
set routing-instances overlay protocols evpn extended-vni-list all
set routing-instances overlay protocols evpn multicast-mode ingress-replication
set routing-instances overlay bridge-domains bd-200 vlan-id 200
set routing-instances overlay bridge-domains bd-200 routing-interface irb.200
set routing-instances overlay bridge-domains bd-200 vxlan vni 200
```

vMX2:

```
set interfaces irb unit 200 family inet address 192.168.200.254/24
set interfaces irb unit 200 mac 00:00:00:00:02:00

set routing-options route-distinguisher-id 172.31.50.2

set protocols ospf area 0.0.0.0 interface irb.200 passive

set routing-instances overlay vtep-source-interface lo0.0
set routing-instances overlay instance-type virtual-switch
set routing-instances overlay route-distinguisher 172.31.50.2:65000
set routing-instances overlay vrf-target target:65000:100
set routing-instances overlay vrf-target auto
set routing-instances overlay protocols evpn encapsulation vxlan
set routing-instances overlay protocols evpn extended-vni-list all
set routing-instances overlay protocols evpn multicast-mode ingress-replication
set routing-instances overlay bridge-domains bd-200 vlan-id 200
set routing-instances overlay bridge-domains bd-200 routing-interface irb.200
set routing-instances overlay bridge-domains bd-200 vxlan vni 200
```

vQFX1:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 200

set routing-options router-id 172.31.49.1

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list all
set protocols evpn multicast-mode ingress-replication

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.31.49.1:65000
set switch-options vrf-target target:65000:100
set switch-options vrf-target auto
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set vlans vlan-200 vlan-id 200
set vlans vlan-200 vxlan vni 200
```

vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 200

set routing-options router-id 172.31.49.2

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list all
set protocols evpn multicast-mode ingress-replication

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.31.49.2:65000
set switch-options vrf-target target:65000:100
set switch-options vrf-target auto

set vlans vlan-200 vlan-id 200
set vlans vlan-200 vxlan vni 200
```

Verification

Verify that both vMX routers are advertising the subnet that the servers are using:

```
jncie@vMX1> show ospf database advertising-router self router detail | match 192

 id 192.168.200.0, data 255.255.255.0, Type Stub (3)

jncie@vMX2> show ospf database advertising-router self router detail | match 192

 id 192.168.200.0, data 255.255.255.0, Type Stub (3)
```

Check connectivity from the servers to the core:

```
jncie@vMX6> ping routing-instance DC1-S1 rapid 172.31.50.5

PING 172.31.50.5 (172.31.50.5): 56 data bytes
!!!!!
--- 172.31.50.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 30.533/33.526/36.373/2.236 ms

jncie@vMX6> ping routing-instance DC1-S2 rapid 172.31.50.5

PING 172.31.50.5 (172.31.50.5): 56 data bytes
!!!!!
--- 172.31.50.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
round-trip min/avg/max/stddev = 42.313/297.653/418.316/131.888 ms
```

Check the EVPN neighbor relationship on the members of the overlay:

```
jncie@vQFX1> show evpn instance extensive

Instance: __default_evpn__
 Route Distinguisher: 172.31.49.1:0
 Number of bridge domains: 0
 Number of neighbors: 0

Instance: default-switch
 Route Distinguisher: 172.31.49.1:65000
 Encapsulation type: VXLAN
 MAC database status                Local  Remote
  MAC advertisements:                 1    2
  MAC+IP advertisements:              0    1
  Default gateway MAC advertisements:    0    0
 Number of local interfaces: 1 (1 up)
  Interface name ESI                  Mode       Status
  xe-0/0/1.0   00:00:00:00:00:00:00:00:00:00 single-homed   Up
 Number of IRB interfaces: 0 (0 up)
 Number of bridge domains: 1
  VLAN VNI  Intfs / up  IRB intf  Mode       MAC sync IM route label
  200  200    1  1         Extended     Enabled  200
 Number of neighbors: 3
  172.31.49.2
   Received routes
    MAC address advertisement:       1
    MAC+IP address advertisement:      0
    Inclusive multicast:          1
    Ethernet auto-discovery:        0
  172.31.50.1
   Received routes
    MAC address advertisement:       1
    MAC+IP address advertisement:      1
    Inclusive multicast:          1
    Ethernet auto-discovery:        1
  172.31.50.2
   Received routes
    MAC address advertisement:       0
    MAC+IP address advertisement:      0
    Inclusive multicast:          1
    Ethernet auto-discovery:        1
 Number of ethernet segments: 1
  ESI: 05:00:00:fd:e8:00:00:00:c8:00
   Status: Resolved
   Number of remote PEs connected: 2
    Remote PE    MAC label Aliasing label Mode
    172.31.50.1   0     0        all-active
    172.31.50.2   0     0        all-active
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
 Router-ID: 172.31.49.1
 Source VTEP interface IP: 172.31.49.1

jncie@vMX1> show evpn instance extensive

Instance: __default_evpn__
 Route Distinguisher: 172.31.50.1:0
 Number of bridge domains: 0
 Number of neighbors: 0

Instance: overlay
 Route Distinguisher: 172.31.50.1:65000
 Encapsulation type: VXLAN
MAC database status               Local Remote
 MAC advertisements:                1    2
 MAC+IP advertisements:             1    0
 Default gateway MAC advertisements:    1    0
Number of local interfaces: 0 (0 up)
Number of IRB interfaces: 1 (1 up)
 Interface name VLAN  VNI  Status L3 context
 irb.200        200   Up   master
Number of bridge domains: 1
 VLAN Domain ID  Intfs / up  IRB intf  Mode       MAC sync IM route label
 200  200      0  0   irb.200  Extended     Enabled  200
Number of neighbors: 3
 172.31.49.1
  Received routes
   MAC address advertisement:       1
   MAC+IP address advertisement:      0
   Inclusive multicast:           1
   Ethernet auto-discovery:         0
 172.31.49.2
  Received routes
   MAC address advertisement:       1
   MAC+IP address advertisement:      0
   Inclusive multicast:           1
   Ethernet auto-discovery:         0
 172.31.50.2
  Received routes
   MAC address advertisement:       0
   MAC+IP address advertisement:      0
   Inclusive multicast:           1
   Ethernet auto-discovery:         1
Number of ethernet segments: 1
 ESI: 05:00:00:fd:e8:00:00:00:c8:00
  Status: Resolved by IFL irb.200
  Local interface: irb.200, Status: Up/Forwarding
  Number of remote PEs connected: 1
   Remote PE    MAC label Aliasing label Mode
   172.31.50.2   0     0        all-active
Router-ID: 172.31.50.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Source VTEP interface IP: 172.31.50.1
```

Task 4.22: DC2 Controllerless Overlay Part 1: the underlay

- Ensure IP-connectivity between the loopback IPs. Use EBGP to accomplish your goal.
- Enable EVPN signaling between the devices. Source the BGP sessions used for EVPN signaling from the loopback interface.
- Use the AS numbers shown in the topology.
- Enable the devices to load-share traffic across links.

Solution

vMX3:

```
set routing-options autonomous-system 65001
set routing-options forwarding-table export lbpp

set protocols bgp group underlay export export-lo0
set protocols bgp group underlay local-as 65001
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 172.16.48.2 peer-as 65003
set protocols bgp group underlay neighbor 172.16.48.6 peer-as 65004
set protocols bgp group underlay neighbor 172.16.48.10 peer-as 65005

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.50.3
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay neighbor 172.31.50.4
set protocols bgp group overlay neighbor 172.31.49.3
set protocols bgp group overlay neighbor 172.31.49.4
set protocols bgp group overlay neighbor 172.31.49.5

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vMX4:

```
set routing-options autonomous-system 65002
set routing-options forwarding-table export lbpp

set protocols bgp group underlay export export-lo0
set protocols bgp group underlay local-as 65002
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 172.16.48.14 peer-as 65003
set protocols bgp group underlay neighbor 172.16.48.18 peer-as 65004
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay neighbor 172.16.48.22 peer-as 65005

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.50.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay neighbor 172.31.50.3
set protocols bgp group overlay neighbor 172.31.49.3
set protocols bgp group overlay neighbor 172.31.49.4
set protocols bgp group overlay neighbor 172.31.49.5

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vQFX3:

```
set routing-options autonomous-system 65003
set routing-options forwarding-table export lbpp

set protocols bgp group underlay export export-lo0
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 172.16.48.1 peer-as 65001
set protocols bgp group underlay neighbor 172.16.48.13 peer-as 65002

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.49.3
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay neighbor 172.31.50.3
set protocols bgp group overlay neighbor 172.31.50.4
set protocols bgp group overlay neighbor 172.31.49.4
set protocols bgp group overlay neighbor 172.31.49.5

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vQFX4:

```
set routing-options autonomous-system 65004
set routing-options forwarding-table export lbpp

set protocols bgp group underlay export export-lo0
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 172.16.48.5 peer-as 65001
set protocols bgp group underlay neighbor 172.16.48.17 peer-as 65002

set protocols bgp group overlay type internal
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay local-address 172.31.49.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay neighbor 172.31.50.3
set protocols bgp group overlay neighbor 172.31.50.4
set protocols bgp group overlay neighbor 172.31.49.3
set protocols bgp group overlay neighbor 172.31.49.5

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

vQFX5:

```
set routing-options autonomous-system 65005
set routing-options forwarding-table export lbpp

set protocols bgp group underlay export export-lo0
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 172.16.48.9 peer-as 65001
set protocols bgp group underlay neighbor 172.16.48.21 peer-as 65002

set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 172.31.49.5
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65500
set protocols bgp group overlay neighbor 172.31.50.3
set protocols bgp group overlay neighbor 172.31.50.4
set protocols bgp group overlay neighbor 172.31.49.3
set protocols bgp group overlay neighbor 172.31.49.4

set policy-options policy-statement export-lo0 term lo0 from interface lo0.0
set policy-options policy-statement export-lo0 term lo0 then accept

set policy-options policy-statement lbpp term lbpp then load-balance per-packet
```

Verification

Verifying the BGP sessions after configuring the underlay:

```
jncie@vMX3> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       6          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
172.16.48.2           65003        8          6       0       0        1:06 1/2/2/0
0/0/0/0
172.16.48.6           65004        7          7       0       0        1:06 1/2/2/0
0/0/0/0
172.16.48.10          65005        8          6       0       0        1:07 1/2/2/0
0/0/0/0

jncie@vMX4> show bgp summary

Groups: 1 Peers: 3 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        12          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.48.14          65003          9          8       0       0        1:10 1/4/4/0
0/0/0/0
172.16.48.18          65004          9          6       0       0        1:14 1/4/4/0
0/0/0/0
172.16.48.22          65005          9         10       0       0        1:11 1/4/4/0
0/0/0/0
```

Verifying IP connectivity:

```
jncie@vQFX3> show route protocol bgp | match /32

172.31.49.4/32     *[BGP/170] 00:00:15, localpref 100
172.31.49.5/32     *[BGP/170] 00:00:15, localpref 100
172.31.50.3/32     *[BGP/170] 00:00:15, localpref 100
172.31.50.4/32     *[BGP/170] 00:00:10, localpref 100

jncie@vQFX3> ping source 172.31.49.3 rapid 172.31.49.4

PING 172.31.49.4 (172.31.49.4): 56 data bytes
!!!!!
--- 172.31.49.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 36.036/42.597/47.264/4.969 ms

{master:0}
jncie@vQFX3> ping source 172.31.49.3 rapid 172.31.49.5

PING 172.31.49.5 (172.31.49.5): 56 data bytes
!!!!!
--- 172.31.49.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 38.292/44.925/66.235/10.774 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
{master:0}
jncie@vQFX3> ping source 172.31.49.3 rapid 172.31.50.3

PING 172.31.50.3 (172.31.50.3): 56 data bytes
!!!!!
--- 172.31.50.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 18.001/18.951/19.869/0.780 ms

{master:0}
jncie@vQFX3> ping source 172.31.49.3 rapid 172.31.50.4

PING 172.31.50.4 (172.31.50.4): 56 data bytes
!!!!!
--- 172.31.50.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 18.842/19.490/20.825/0.723 ms
```

Verify load balancing:

```
jncie@vQFX3> show route forwarding-table destination 172.31.49.5

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
172.31.49.5/32     user     0                    ulst   131070     3
                            172.16.48.1          ucst     1707     5 xe-0/0/2.0
                            172.16.48.13         ucst     1706     5 xe-0/0/4.0
```

Task 4.23: DC2 Controllerless Overlay Part 2: the overlay

- Configure a VXLAN layer 2 gateway on the vQFX switches for VLAN 100.
- Server-facing interfaces should support tagged traffic. Ensure that the servers can reach the core network.
- Configure the following layer 3 gateway on vMX3 and vMX4:

VLAN	IPv4 gateway	MAC address
100	192.168.100.254	0:00:00:00:01:00

- You should be able to hide VXLAN segments on a per-device basis.
- Configure the Ether Segments shown in the topology.

Solution

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX3:

```
set interfaces irb unit 100 family inet address 192.168.100.254/24
set interfaces irb unit 100 mac 00:00:00:00:01:00

set routing-options route-distinguisher-id 172.31.50.3

set protocols ospf area 0.0.0.0 interface irb.100 passive

set policy-options policy-statement fabric-import term leaf-nodes from community leaf-nodes
set policy-options policy-statement fabric-import term leaf-nodes then accept
set policy-options policy-statement fabric-import term spine-nodes from community spine-nodes
set policy-options policy-statement fabric-import term spine-nodes then accept
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept

set policy-options community leaf-nodes members target:65000:100
set policy-options community spine-nodes members target:65000:101
set policy-options community vni-100 members target:65500:100

set routing-instances overlay vtep-source-interface lo0.0
set routing-instances overlay instance-type virtual-switch
set routing-instances overlay vrf-import fabric-import
set routing-instances overlay vrf-target target:65000:101
set routing-instances overlay protocols evpn encapsulation vxlan
set routing-instances overlay protocols evpn extended-vni-list 100
set routing-instances overlay protocols evpn vni-options vni 100 vrf-target target:65500:100
set routing-instances overlay protocols evpn multicast-mode ingress-replication
set routing-instances overlay bridge-domains bd-100 vlan-id 100
set routing-instances overlay bridge-domains bd-100 routing-interface irb.100
set routing-instances overlay bridge-domains bd-100 vxlan vni 100
```

vMX4:

```
set interfaces irb unit 100 family inet address 192.168.100.254/24
set interfaces irb unit 100 mac 00:00:00:00:01:00

set routing-options route-distinguisher-id 172.31.50.4

set protocols ospf area 0.0.0.0 interface irb.100 passive

set policy-options policy-statement fabric-import term leaf-nodes from community leaf-nodes
set policy-options policy-statement fabric-import term leaf-nodes then accept
set policy-options policy-statement fabric-import term spine-nodes from community spine-nodes
set policy-options policy-statement fabric-import term spine-nodes then accept
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept

set policy-options community leaf-nodes members target:65000:100
set policy-options community spine-nodes members target:65000:101
set policy-options community vni-100 members target:65500:100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances overlay vtep-source-interface lo0.0
set routing-instances overlay instance-type virtual-switch
set routing-instances overlay vrf-import fabric-import
set routing-instances overlay vrf-target target:65000:101
set routing-instances overlay protocols evpn encapsulation vxlan
set routing-instances overlay protocols evpn extended-vni-list 100
set routing-instances overlay protocols evpn vni-options vni 100 vrf-target target:65500:100
set routing-instances overlay protocols evpn multicast-mode ingress-replication
set routing-instances overlay bridge-domains bd-100 vlan-id 100
set routing-instances overlay bridge-domains bd-100 routing-interface irb.100
set routing-instances overlay bridge-domains bd-100 vxlan vni 100
```

vQFX3:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae0 esi 00:03:04:05:03:04:05:03:04:05
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 03:04:05:03:04:05
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set interfaces ae1 esi 00:03:04:03:04:03:04:03:04:03
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 03:04:03:04:03:04
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 100

set routing-options router-id 172.31.49.3

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list 100
set protocols evpn multicast-mode ingress-replication
set protocols evpn vni-options vni 100 vrf-target export target:65500:100

set policy-options policy-statement fabric-import term leaf-nodes from community leaf-nodes
set policy-options policy-statement fabric-import term leaf-nodes then accept
set policy-options policy-statement fabric-import term spine-nodes from community spine-nodes
set policy-options policy-statement fabric-import term spine-nodes then accept
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept

set policy-options community leaf-nodes members target:65000:100
set policy-options community spine-nodes members target:65000:101
set policy-options community vni-100 members target:65500:100
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.31.49.3:65003
set switch-options vrf-import fabric-import
set switch-options vrf-target target:65000:100

set vlans vlan-100 vlan-id 100
set vlans vlan-100 vxlan vni 100
```

vQFX4:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces xe-0/0/1 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae1

set interfaces ae0 esi 00:03:04:05:03:04:05:03:04:05
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 03:04:05:03:04:05
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set interfaces ae1 esi 00:03:04:03:04:03:04:03:04:03
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 03:04:03:04:03:04
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members 100

set routing-options router-id 172.31.49.4

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list 100
set protocols evpn multicast-mode ingress-replication
set protocols evpn vni-options vni 100 vrf-target export target:65500:100

set policy-options policy-statement fabric-import term leaf-nodes from community leaf-nodes
set policy-options policy-statement fabric-import term leaf-nodes then accept
set policy-options policy-statement fabric-import term spine-nodes from community spine-nodes
set policy-options policy-statement fabric-import term spine-nodes then accept
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept


set policy-options community leaf-nodes members target:65000:100
set policy-options community spine-nodes members target:65000:101
set policy-options community vni-100 members target:65500:100

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.31.49.4:65003
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set switch-options vrf-import fabric-import
set switch-options vrf-target target:65000:100

set vlans vlan-100 vlan-id 100
set vlans vlan-100 vxlan vni 100
```

vQFX5:

```
set chassis aggregated-devices ethernet device-count 1

set interfaces xe-0/0/1 ether-options 802.3ad ae0

set interfaces ae0 esi 00:03:04:05:03:04:05:03:04:05
set interfaces ae0 esi all-active
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 03:04:05:03:04:05
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members 100

set routing-options router-id 172.31.49.5

set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list 100
set protocols evpn multicast-mode ingress-replication
set protocols evpn vni-options vni 100 vrf-target export target:65500:100

set policy-options policy-statement fabric-import term leaf-nodes from community leaf-nodes
set policy-options policy-statement fabric-import term leaf-nodes then accept
set policy-options policy-statement fabric-import term spine-nodes from community spine-nodes
set policy-options policy-statement fabric-import term spine-nodes then accept
set policy-options policy-statement fabric-import term vni-100 from community vni-100
set policy-options policy-statement fabric-import term vni-100 then accept

set policy-options community leaf-nodes members target:65000:100
set policy-options community spine-nodes members target:65000:101
set policy-options community vni-100 members target:65500:100

set switch-options vtep-source-interface lo0.0
set switch-options route-distinguisher 172.31.49.5:65003
set switch-options vrf-import fabric-import
set switch-options vrf-target target:65000:100

set vlans vlan-100 vlan-id 100
set vlans vlan-100 vxlan vni 100
```

Verification

Verifying IP connectivity from the servers to the core:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping routing-instance DC2-S1 rapid 172.31.50.5

PING 172.31.50.5 (172.31.50.5): 56 data bytes
!!!!!
--- 172.31.50.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 21.760/37.698/52.050/10.783 ms

jncie@vMX6> ping routing-instance DC2-S2 rapid 172.31.50.5

PING 172.31.50.5 (172.31.50.5): 56 data bytes
!!!!!
--- 172.31.50.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 42.309/138.212/271.956/86.909 ms
```

Verifying the EVPN instance from both spine routers:

```
jncie@vMX3> show evpn instance overlay extensive

Instance: overlay
  Route Distinguisher: 172.31.50.3:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           1       3
    MAC+IP advertisements:                        1       1
    Default gateway MAC advertisements:           1       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0     irb.100    Extended        Enabled   100
  Number of neighbors: 4
172.31.49.3
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                4
172.31.49.4
    Received routes
      MAC address advertisement:              0
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                4
172.31.49.5
    Received routes
      MAC address advertisement:              0
      MAC+IP address advertisement:           0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Inclusive multicast:                    1
      Ethernet auto-discovery:                2
  172.31.50.4
    Received routes
      MAC address advertisement:              1
      MAC+IP address advertisement:           1
      Inclusive multicast:                    1
      Ethernet auto-discovery:                1
Number of ethernet segments: 4
  ESI: 00:03:04:03:04:03:04:03:04:03
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      172.31.49.3      100        0               all-active
      172.31.49.4      0          0               all-active
    ESI: 00:03:04:05:03:04:05:03:04:05
      Status: Resolved
      Number of remote PEs connected: 3
        Remote PE        MAC label  Aliasing label  Mode
        172.31.49.3      100        0               all-active
        172.31.49.5      0          0               all-active
        172.31.49.4      0          0               all-active
    ESI: 05:00:00:fd:e9:00:00:00:64:00
      Status: Resolved by IFL irb.100
      Local interface: irb.100, Status: Up/Forwarding
    ESI: 05:00:00:fd:ea:00:00:00:64:00
      Status: Resolved
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        172.31.50.4      0          0               all-active
  Router-ID: 172.31.50.3
  Source VTEP interface IP: 172.31.50.3

jncie@vMX4> show evpn instance overlay extensive

Instance: overlay
  Route Distinguisher: 172.31.50.4:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           1       3
    MAC+IP advertisements:                        1       1
    Default gateway MAC advertisements:           1       0
  Number of local interfaces: 0 (0 up)
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0     irb.100    Extended        Enabled   100
  Number of neighbors: 4
    172.31.49.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Received routes
    MAC address advertisement:              2
    MAC+IP address advertisement:           0
    Inclusive multicast:                    1
    Ethernet auto-discovery:                4
172.31.49.4
  Received routes
    MAC address advertisement:              0
    MAC+IP address advertisement:           0
    Inclusive multicast:                    1
    Ethernet auto-discovery:                4
172.31.49.5
  Received routes
    MAC address advertisement:              0
    MAC+IP address advertisement:           0
    Inclusive multicast:                    1
    Ethernet auto-discovery:                2
  172.31.50.3
    Received routes
      MAC address advertisement:            1
      MAC+IP address advertisement:         1
      Inclusive multicast:                  1
      Ethernet auto-discovery:              1
Number of ethernet segments: 4
  ESI: 00:03:04:03:04:03:04:03:04:03
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      172.31.49.3      100        0               all-active
      172.31.49.4      0          0               all-active
  ESI: 00:03:04:05:03:04:05:03:04:05
    Status: Resolved
    Number of remote PEs connected: 3
      Remote PE        MAC label  Aliasing label  Mode
      172.31.49.3      100        0               all-active
      172.31.49.4      0          0               all-active
      172.31.49.5      0          0               all-active
  ESI: 05:00:00:fd:e9:00:00:00:64:00
    Status: Resolved
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.31.50.3      0          0               all-active
  ESI: 05:00:00:fd:ea:00:00:00:64:00
    Status: Resolved by IFL irb.100
    Local interface: irb.100, Status: Up/Forwarding
Router-ID: 172.31.50.4
Source VTEP interface IP: 172.31.50.4
```

Verifying LACP on the leaf switches:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX3> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3                  Current   Fast periodic Collecting distributing

jncie@vQFX4> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/3                  Current   Fast periodic Collecting distributing

jncie@vQFX5> show lacp interfaces

Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/1                  Current   Fast periodic Collecting distributing
```

Verifying the Ethernet switching table on vQFX:

```
jncie@vQFX3> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
   Vlan                 MAC                  MAC      Logical                  Active
   name                 address              flags    interface                source
   vlan-100             00:00:00:00:01:00    D        vtep.32770               172.31.50.4
   vlan-100             00:03:03:03:03:03    DL       ae0.0
   vlan-100             00:04:04:04:04:04    DL       ae1.0

jncie@vQFX4> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
   Vlan                 MAC                  MAC      Logical                  Active
   name                 address              flags    interface                source
   vlan-100             00:00:00:00:01:00    D        vtep.32771               172.31.50.3
   vlan-100             00:03:03:03:03:03    DR       ae0.0
   vlan-100             00:04:04:04:04:04    DR       ae1.0

jncie@vQFX5> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 3 entries, 3 learned
Routing instance : default-switch
   Vlan                 MAC                  MAC      Logical                  Active
   name                 address              flags    interface                source
   vlan-100             00:00:00:00:01:00    D        vtep.32772               172.31.50.3
   vlan-100             00:03:03:03:03:03    DR       ae0.0
   vlan-100             00:04:04:04:04:04    DR       esi.1721
00:03:04:03:04:03:04:03:04:03
```

Verifying the EVPN database on vMX:

```
jncie@vMX3> show evpn database

Instance: overlay
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     100        00:00:00:00:01:00  irb.100                        Jun 10 11:22:23
192.168.100.254
     100        00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:22:48
     100        00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:22:48
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX4> show evpn database

Instance: overlay
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     100        00:00:00:00:01:00  irb.100                        Jun 10 11:22:23
192.168.100.254
     100        00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:22:38
     100        00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:22:38
```

Task 4.24: EVPN Stitching

- Configure a layer 2 VXLAN gateway on all server facing interfaces for VLAN 250. This VXLAN gateway should be active on every vQFX.
- Configure a layer 2 gateway for VLAN 250 on both vMX2 as well as vMX3.
- Any inter-datacenter traffic inside VLAN 250 should be transported across the MPLS LSPs.

Solution

vMX2:

To be able to complete this assignment, we will need to connect an EVPN over MPLS domain to an EVPN over VXLAN domain. On an MX, this can be done using a logical tunnel interface. First, we need to enable the MX to support logical tunnel interfaces:

```
set chassis fpc 0 pic 0 tunnel-services
```

After this, we create two logical tunnel interfaces that we connect together. We enable the interfaces for the proper encapsulation and VLAN-ID:

```
set interfaces lt-0/0/0 unit 0 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 0 peer-unit 1
set interfaces lt-0/0/0 unit 0 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 0 family bridge vlan-id-list 250

set interfaces lt-0/0/0 unit 1 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 1 family bridge vlan-id-list 250
```

Next, we create an EVPN instance and place one of the logical tunnel interfaces in this routing-instance:

```
set routing-instances evpn instance-type virtual-switch
set routing-instances evpn interface lt-0/0/0.0
set routing-instances evpn vrf-target target:6500:250
set routing-instances evpn protocols evpn extended-vlan-list 250
set routing-instances evpn bridge-domains bd-250 vlan-id 250
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

The previous instance will be the EVPN over MPLS instance. To connect this EVPN over MPLS domain with the EVPN over VXLAN domain, we place the other logical tunnel interface into the already existing overlay instance:

```
set routing-instances overlay interface lt-0/0/0.1
```

And we enable that overlay instance for the additional VLAN/VXLAN:

```
set routing-instances overlay bridge-domains bd-250 vlan-id 250
set routing-instances overlay bridge-domains bd-250 vxlan vni 250
```

The final task on vMX2 is to configure a BGP peering session with vMX3 so that they can exchange EVPN routes:

```
set protocols bgp group dci type internal
set protocols bgp group dci local-address 172.31.50.2
set protocols bgp group dci family evpn signaling
set protocols bgp group dci neighbor 172.31.50.3
```

vMX3:

The configuration tasks on vMX3 are mostly similar to those on vMX2. We enable the vMX for logical tunnel interfaces:

```
set chassis fpc 0 pic 0 tunnel-services
```

We configure two logical tunnel interfaces and tie them together:

```
set interfaces lt-0/0/0 unit 0 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 0 peer-unit 1
set interfaces lt-0/0/0 unit 0 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 0 family bridge vlan-id-list 250
set interfaces lt-0/0/0 unit 1 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 1 family bridge vlan-id-list 250
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We then configure the routing instance that will use MPLS as the encapsulation protocol for EVPN:

```
set routing-instances evpn instance-type virtual-switch
set routing-instances evpn interface lt-0/0/0.0
set routing-instances evpn vrf-target target:6500:250
set routing-instances evpn protocols evpn extended-vlan-list 250
set routing-instances evpn bridge-domains bd-250 vlan-id 250
```

After this, we enable the already existing overlay for the new VXLAN segment:

```
set routing-instances overlay interface lt-0/0/0.1
set routing-instances overlay protocols evpn extended-vni-list 250
set routing-instances overlay protocols evpn vni-options vni 250 vrf-target target:65500:250
set routing-instances overlay bridge-domains bd-250 vlan-id 250
set routing-instances overlay bridge-domains bd-250 routing-interface irb.250
set routing-instances overlay bridge-domains bd-250 vxlan vni 250

set policy-options policy-statement fabric-import term vni-250 from community vni-250
set policy-options policy-statement fabric-import term vni-250 then accept

set policy-options community vni-250 members target:65500:250
```

And finally, we configure a BGP peering session with vMX2:

```
set protocols bgp group dci type internal
set protocols bgp group dci local-address 172.31.50.3
set protocols bgp group dci family evpn signaling
set protocols bgp group dci local-as 65000
set protocols bgp group dci neighbor 172.31.50.2
```

On the vQFX devices, we add the configuration that is necessary to enable the overlay for the additional VXLAN.

vQFX1:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 250

set vlans vlan-250 vlan-id 250
set vlans vlan-250 vxlan vni 250
```

vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members 250

set vlans vlan-250 vlan-id 250
set vlans vlan-250 vxlan vni 250
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set interfaces ae0 unit 0 family ethernet-switching vlan members 250
set interfaces ae1 unit 0 family ethernet-switching vlan members 250

set protocols evpn extended-vni-list 250
set protocols evpn vni-options vni 250 vrf-target export target:65500:250

set policy-options policy-statement fabric-import term vni-250 from community vni-250
set policy-options policy-statement fabric-import term vni-250 then accept

set policy-options community vni-250 members target:65500:250

set vlans vlan-250 vlan-id 250
set vlans vlan-250 vxlan vni 250
```

vQFX4:

```
set interfaces ae0 unit 0 family ethernet-switching vlan members 250
set interfaces ae1 unit 0 family ethernet-switching vlan members 250

set protocols evpn extended-vni-list 250
set protocols evpn vni-options vni 250 vrf-target export target:65500:250

set policy-options policy-statement fabric-import term vni-250 from community vni-250
set policy-options policy-statement fabric-import term vni-250 then accept

set policy-options community vni-250 members target:65500:250

set vlans vlan-250 vlan-id 250
set vlans vlan-250 vxlan vni 250
```

vQFX5:

```
set interfaces ae0 unit 0 family ethernet-switching vlan members 250

set protocols evpn extended-vni-list 250
set protocols evpn vni-options vni 250 vrf-target export target:65500:250

set policy-options policy-statement fabric-import term vni-250 from community vni-250
set policy-options policy-statement fabric-import term vni-250 then accept

set policy-options community vni-250 members target:65500:250

set vlans vlan-250 vlan-id 250
set vlans vlan-250 vxlan vni 250
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

LSPs were already present between vMX2 and vMX3:

```
jncie@vMX2> show route table inet.3 172.31.50.3

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.50.3/32     *[LDP/9] 00:01:49, metric 1
                    > to 172.16.48.38 via xe-0/0/6.13, Push 302304
                      to 172.16.24.38 via xe-0/0/5.13, Push 302304
```

```
jncie@vMX3> show route table inet.3 172.31.50.2

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.50.2/32     *[LDP/9] 00:02:25, metric 1
                    > to 172.16.48.26 via xe-0/0/6.10, Push 302224
                      to 172.16.24.26 via xe-0/0/5.10, Push 302224
```

Verify the newly configured BGP session for EVPN signaling between vMX2 and vMX3:

```
jncie@vMX2> show bgp neighbor 172.31.50.3

Peer: 172.31.50.3+52245 AS 65000 Local: 172.31.50.2+179 AS 65000
  Group: dci                    Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
  Address families configured: evpn
  Local Address: 172.31.50.2 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Error: 'Open Message Error' Sent: 1 Recv: 0
  Peer ID: 172.31.50.3     Local ID: 172.31.50.2       Active Holdtime: 90
  Keepalive Interval: 30         Group index: 1    Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: evpn
  NLRI advertised by peer: evpn
  NLRI for this session: evpn
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: evpn
  NLRI of received end-of-rib markers: evpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
NLRI of all end-of-rib markers sent: evpn
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65000)
Peer does not support Addpath
Table bgp.evpn.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:              8
  Received prefixes:            8
  Accepted prefixes:            8
  Suppressed due to damping:    0
  Advertised prefixes:          20
Table overlay.evpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
Table __default_evpn__.evpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
Table evpn.evpn.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:              1
  Received prefixes:            1
  Accepted prefixes:            1
  Suppressed due to damping:    0
Last traffic (seconds): Received 282863 Sent 96 Checked 282863
Input messages:  Total 13     Updates 8       Refreshes 0     Octets 916
Output messages: Total 22     Updates 18      Refreshes 0     Octets 2135
Output Queue[1]: 0            (bgp.evpn.0, evpn)
```

After tunnel-services is configured, the ‘lt-‘ interfaces should be listed when we issue the following command:

```
jncie@vMX2> show interfaces terse | match lt-0

lt-0/0/0                up    up
lt-0/0/0.0              up    up   bridge
lt-0/0/0.1              up    up   bridge
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show interfaces terse | match lt-0

lt-0/0/0                up    up
lt-0/0/0.0              up    up   bridge
lt-0/0/0.1              up    up   bridge
```

Verify the newly created EVPN instance that will use MPLS as the encapsulation method:

```
jncie@vMX2> show evpn instance evpn extensive

Instance: evpn
  Route Distinguisher: 172.31.50.2:9
  Per-instance MAC route label: 299936
  Per-instance multicast route label: 299920
  MAC database status                       Local  Remote
    MAC advertisements:                         2       2
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                               Mode             Status
    lt-0/0/0.0      00:00:00:00:00:00:00:00:00:00     single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode             MAC sync  IM route label
    250                  1    1               Extended         Enabled   299920
  Number of neighbors: 1
    172.31.50.3
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    1
        Ethernet auto-discovery:                0
  Number of ethernet segments: 0

jncie@vMX3> show evpn instance evpn extensive

Instance: evpn
  Route Distinguisher: 172.31.50.3:9
  Per-instance MAC route label: 299856
  Per-instance multicast route label: 299840
  MAC database status                       Local  Remote
    MAC advertisements:                         2       2
    MAC+IP advertisements:                      0       0
    Default gateway MAC advertisements:         0       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                               Mode             Status
    lt-0/0/0.0      00:00:00:00:00:00:00:00:00:00     single-homed     Up
  Number of IRB interfaces: 0 (0 up)
  Number of bridge domains: 1
    VLAN  Domain ID   Intfs / up    IRB intf   Mode             MAC sync  IM route label
    250                  1    1               Extended         Enabled   299840
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Number of neighbors: 1
  172.31.50.2
    Received routes
      MAC address advertisement:              2
      MAC+IP address advertisement:           0
      Inclusive multicast:                    1
      Ethernet auto-discovery:                0
Number of ethernet segments: 0
```

Verify that the 'lt-' interface that is used for stitching is also present in the overlay instance:

```
jncie@vMX2> show evpn instance overlay extensive
Instance: overlay
  Route Distinguisher: 172.31.50.2:65000
  Encapsulation type: VXLAN
  MAC database status                       Local  Remote
    MAC advertisements:                         3       5
    MAC+IP advertisements:                      1       1
    Default gateway MAC advertisements:         1       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    lt-0/0/0.1      00:00:00:00:00:00:00:00:00:00      single-homed     Up
  Number of IRB interfaces: 1 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.200                200     Up     master
  Number of bridge domains: 2
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    200   200            0    0     irb.200    Extended        Enabled   200
    250   250            1    1                Extended        Enabled   250
  Number of neighbors: 3
    172.31.49.1
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    2
        Ethernet auto-discovery:                0
    172.31.49.2
      Received routes
        MAC address advertisement:              2
        MAC+IP address advertisement:           0
        Inclusive multicast:                    2
        Ethernet auto-discovery:                0
    172.31.50.1
      Received routes
        MAC address advertisement:              1
        MAC+IP address advertisement:           1
        Inclusive multicast:                    1
        Ethernet auto-discovery:                1
  Number of ethernet segments: 1
    ESI: 05:00:00:fd:e8:00:00:00:c8:00
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Status: Resolved by IFL irb.200
      Local interface: irb.200, Status: Up/Forwarding
      Number of remote PEs connected: 1
        Remote PE        MAC label  Aliasing label  Mode
        172.31.50.1      0          0               all-active
  Router-ID: 172.31.50.2
  Source VTEP interface IP: 172.31.50.2

jncie@vMX3> show evpn instance overlay extensive

Instance: overlay
  Route Distinguisher: 172.31.50.3:8
  Encapsulation type: VXLAN
  MAC database status                         Local  Remote
    MAC advertisements:                           3       6
    MAC+IP advertisements:                        1       1
    Default gateway MAC advertisements:           1       0
  Number of local interfaces: 1 (1 up)
    Interface name  ESI                                Mode             Status
    lt-0/0/0.1      00:00:00:00:00:00:00:00:00:00      single-homed     Up
  Number of IRB interfaces: 2 (1 up)
    Interface name  VLAN   VNI    Status  L3 context
    irb.100                100     Up     master
    irb.250                250     Down   Not assigned
  Number of bridge domains: 2
    VLAN  Domain ID   Intfs / up    IRB intf   Mode            MAC sync  IM route label
    100   100            0    0     irb.100    Extended        Enabled   100
    250   250            1    1     irb.250    Extended        Enabled   250
  Number of neighbors: 6
    172.31.49.3
      Received routes
        MAC address advertisement:             1
        MAC+IP address advertisement:          0
        Inclusive multicast:                   2
        Ethernet auto-discovery:               4
    172.31.49.4
      Received routes
        MAC address advertisement:             2
        MAC+IP address advertisement:          0
        Inclusive multicast:                   2
        Ethernet auto-discovery:               4
    172.31.49.5
      Received routes
        MAC address advertisement:             2
        MAC+IP address advertisement:          0
        Inclusive multicast:                   2
        Ethernet auto-discovery:               2
    172.31.50.1
      Received routes
        MAC address advertisement:             0
        MAC+IP address advertisement:          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Inclusive multicast:                    0
    Ethernet auto-discovery:                1
172.31.50.2
  Received routes
    MAC address advertisement:              0
    MAC+IP address advertisement:           0
    Inclusive multicast:                    0
    Ethernet auto-discovery:                1
172.31.50.4
  Received routes
    MAC address advertisement:              1
    MAC+IP address advertisement:           1
    Inclusive multicast:                    1
    Ethernet auto-discovery:                1
Number of ethernet segments: 6
  ESI: 00:03:04:03:04:03:04:03:04:03
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      172.31.49.4      250        0               all-active
      172.31.49.3      0          0               all-active
  ESI: 00:03:04:05:03:04:05:03:04:05
    Status: Resolved
    Number of remote PEs connected: 3
      Remote PE        MAC label  Aliasing label  Mode
      172.31.49.5      250        0               all-active
      172.31.49.3      250        0               all-active
      172.31.49.4      0          0               all-active
  ESI: 05:00:00:fd:e8:00:00:00:c8:00
    Status: Resolved
    Number of remote PEs connected: 2
      Remote PE        MAC label  Aliasing label  Mode
      172.31.50.1      0          0               all-active
      172.31.50.2      0          0               all-active
  ESI: 05:00:00:fd:e9:00:00:00:64:00
    Status: Resolved by IFL irb.100
    Local interface: irb.100, Status: Up/Forwarding
  ESI: 05:00:00:fd:e9:00:00:00:fa:00
    Status: Resolved
    Local interface: irb.250, Status: Down
  ESI: 05:00:00:fd:ea:00:00:00:64:00
    Status: Resolved
    Number of remote PEs connected: 1
      Remote PE        MAC label  Aliasing label  Mode
      172.31.50.4      0          0               all-active
Router-ID: 172.31.50.3
Source VTEP interface IP: 172.31.50.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verify the EVPN table on the devices in DC1:

```
jncie@vQFX1> show evpn database l2-domain-id 250

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      250   00:01:01:01:01:01  xe-0/0/1.0                     Jun 10 11:33:58
      250   00:02:02:02:02:02  172.31.49.2                    Jun 10 11:34:04
      250   00:03:03:03:03:03  172.31.50.2                    Jun 10 11:40:31
      250   00:04:04:04:04:04  172.31.50.2                    Jun 10 11:40:47
```

```
jncie@vQFX2> show evpn database l2-domain-id 250

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      250   00:01:01:01:01:01  172.31.49.1                    Jun 10 11:33:59
      250   00:02:02:02:02:02  xe-0/0/1.0                     Jun 10 11:34:03
      250   00:03:03:03:03:03  172.31.50.2                    Jun 10 11:40:31
      250   00:04:04:04:04:04  172.31.50.2                    Jun 10 11:40:47
```

```
jncie@vMX2> show evpn database l2-domain-id 250

Instance: evpn
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
250             00:01:01:01:01:01  lt-0/0/0.0                     Jun 10 11:33:56
250             00:02:02:02:02:02  lt-0/0/0.0                     Jun 10 11:34:02
250             00:03:03:03:03:03  172.31.50.3                    Jun 10 11:40:29
250             00:04:04:04:04:04  172.31.50.3                    Jun 10 11:40:45

Instance: overlay
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     250        00:01:01:01:01:01  172.31.49.1                    Jun 10 11:33:58
     250        00:02:02:02:02:02  172.31.49.2                    Jun 10 11:34:03
     250        00:03:03:03:03:03  lt-0/0/0.1                     Jun 10 11:40:28
     250        00:04:04:04:04:04  lt-0/0/0.1                     Jun 10 11:40:44
```

Verify the EVPN table on the devices in DC2:

```
jncie@vQFX3> show evpn database l2-domain-id 250

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      250   00:01:01:01:01:01  172.31.50.3                    Jun 10 11:34:15
      250   00:02:02:02:02:02  172.31.50.3                    Jun 10 11:34:19
      250   00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:40:31
      250   00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:40:46
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vQFX4> show evpn database l2-domain-id 250

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      250   00:01:01:01:01:01  172.31.50.3                    Jun 10 11:34:14
      250   00:02:02:02:02:02  172.31.50.3                    Jun 10 11:34:18
      250   00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:40:31
      250   00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:40:45

jncie@vQFX5> show evpn database l2-domain-id 250

Instance: default-switch
VLAN  VNI  MAC address        Active source                  Timestamp        IP address
      250   00:01:01:01:01:01  172.31.50.3                    Jun 10 11:34:24
      250   00:02:02:02:02:02  172.31.50.3                    Jun 10 11:34:24
      250   00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:40:30
      250   00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:40:45

jncie@vMX3> show evpn database l2-domain-id 250

Instance: evpn
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
250             00:01:01:01:01:01  172.31.50.2                    Jun 10 11:34:03
250             00:02:02:02:02:02  172.31.50.2                    Jun 10 11:34:03
250             00:03:03:03:03:03  lt-0/0/0.0                     Jun 10 11:40:28
250             00:04:04:04:04:04  lt-0/0/0.0                     Jun 10 11:40:43

Instance: overlay
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     250        00:01:01:01:01:01  lt-0/0/0.1                     Jun 10 11:34:11
     250        00:02:02:02:02:02  lt-0/0/0.1                     Jun 10 11:34:15
     250        00:03:03:03:03:03  00:03:04:05:03:04:05:03:04:05  Jun 10 11:40:29
     250        00:04:04:04:04:04  00:03:04:03:04:03:04:03:04:03  Jun 10 11:40:44
```

Verify IP connectivity within the VXLAN segment:

```
jncie@vMX6> ping routing-instance DC1-S3 rapid 192.168.250.2

PING 192.168.250.2 (192.168.250.2): 56 data bytes
!!!!!
--- 192.168.250.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 71.240/96.200/122.580/18.285 ms

jncie@vMX6> ping routing-instance DC1-S3 rapid 192.168.250.3

PING 192.168.250.3 (192.168.250.3): 56 data bytes
!!!!!
--- 192.168.250.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 76.250/107.790/151.056/30.791 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX6> ping routing-instance DC1-S3 rapid 192.168.250.4
PING 192.168.250.4 (192.168.250.4): 56 data bytes
!!!!!
--- 192.168.250.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 101.426/210.819/309.963/72.922 ms
```

Appendix – Chapter 5: Security

Part 1: Securing the Data Center

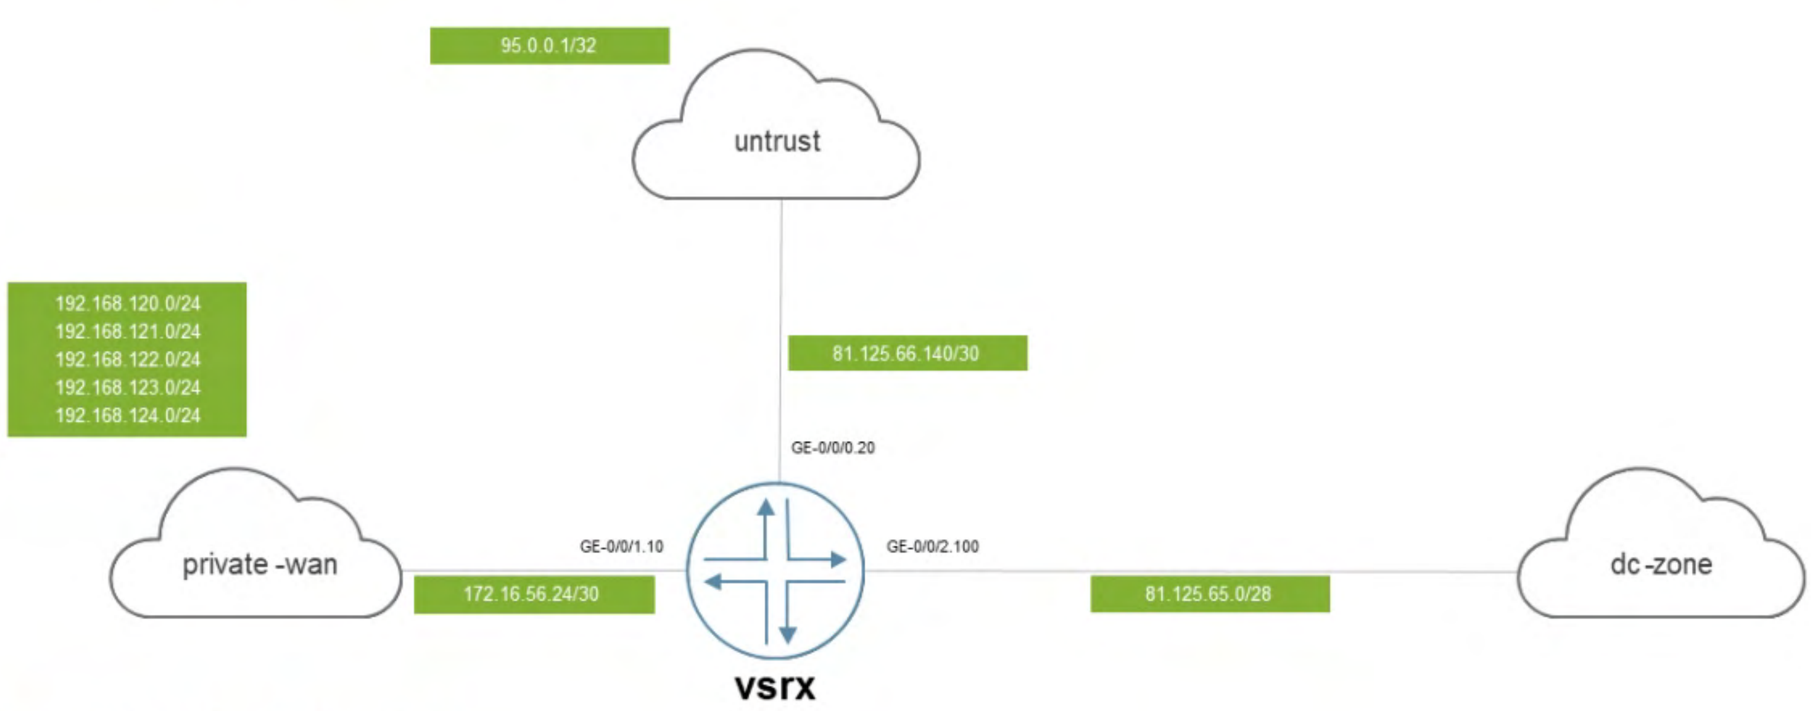

Figure 25 - Network topology

Task 5.1: vSRX Zone and Interface Configuration

- Configure the following zones:

Interface	IP address	VLAN-ID	Zone
GE-0/0/0	81.125.66.141/30	20	untrust
GE-0/0/1	172.16.56.25/30	10	private-wan
GE-0/0/2	81.125.65.1/28	100	dc-zone

- Users in the dc-zone and the private-wan zone should be able to ping the address configured on the vSRX.
- Configure a static default route towards the 81.125.66.142 address.
- Configure the vSRX with autonomous system 65000 and set up peering with 172.16.56.26 inside autonomous system 65010. Advertise a default route into the private-wan domain.

Solution

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vSRX:

Interface configuration:

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 20 vlan-id 20
set interfaces ge-0/0/0 unit 20 family inet address 81.125.66.141/30

set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 10 vlan-id 10
set interfaces ge-0/0/1 unit 10 family inet address 172.16.56.25/30

set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2 unit 100 vlan-id 100
set interfaces ge-0/0/2 unit 100 family inet address 81.125.65.1/28
```

Zone configuration:

```
set security zones security-zone private-wan host-inbound-traffic system-services ping
set security zones security-zone private-wan interfaces ge-0/0/1.10

set security zones security-zone untrust interfaces ge-0/0/0.20

set security zones security-zone dc-zone host-inbound-traffic system-services ping
set security zones security-zone dc-zone interfaces ge-0/0/2.100
```

Static default route towards the untrust domain:

```
set routing-options static route 0.0.0.0/0 next-hop 81.125.66.142
```

Peering configuration with the private-wan router, including the policy to advertise the default route into the private-wan domain:

```
set routing-options autonomous-system 65000

set protocols bgp group private-wan peer-as 65010
set protocols bgp group private-wan neighbor 172.16.56.26
set protocols bgp group private-wan export static

set policy-options policy-statement static term 1 from protocol static
set policy-options policy-statement static term 1 then accept
```

Verification

To verify which zones exist on the vSRX, we issue the following command:

```
jncie@vSRX> show security zones

Functional zone: management
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Policy configurable: No
  Interfaces bound: 1
  Interfaces:
    ge-0/0/3.0

Security zone: dc-zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/2.100

Security zone: private-wan
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.10

Security zone: untrust
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/0.20

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

We can see the same information when we examine the individual interfaces. In addition to verifying what zone an interface is in, we can also see the 'Allowed host-inbound traffic' (if any is permitted through configuration):

```
jncie@vSRX> show interfaces ge-0/0/2.100

  Logical interface ge-0/0/2.100 (Index 76) (SNMP ifIndex 527)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 1
    Security: Zone: dc-zone
    Allowed host-inbound traffic : ping
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 81.125.65.0/28, Local: 81.125.65.1, Broadcast: 81.125.65.15

jncie@vSRX> show interfaces ge-0/0/1.10
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Logical interface ge-0/0/1.10 (Index 74) (SNMP ifIndex 523)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
    Input packets : 12
    Output packets: 1
    Security: Zone: private-wan
    Allowed host-inbound traffic : ping
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 172.16.56.24/30, Local: 172.16.56.25, Broadcast: 172.16.56.27

jncie@vSRX> show interfaces ge-0/0/0.20

  Logical interface ge-0/0/0.20 (Index 69) (SNMP ifIndex 525)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.20 ]  Encapsulation: ENET2
    Input packets : 1
    Output packets: 2
    Security: Zone: untrust
    Allowed host-inbound traffic : bgp
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 81.125.66.140/30, Local: 81.125.66.141, Broadcast: 81.125.66.143
```

Verifying the BGP session:

```
jncie@vSRX> show bgp summary

Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.16.56.26          65010          8          9       0       0        2:47 5/5/5/0
0/0/0/0
```

After checking the state of the BGP session, we examine the routing table. Here we can see that we are learning BGP routes from the private WAN domain. In addition to this, we can also see the entry for the static route:

```
jncie@vSRX> show route

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

0.0.0.0/0          *[Static/5] 00:02:56
                    > to 81.125.66.142 via ge-0/0/0.20
10.10.20.0/24      *[Direct/0] 3d 23:25:53
                    > via ge-0/0/3.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
10.10.20.7/32      *[Local/0] 3d 23:25:57
                      Local via ge-0/0/3.0
81.125.65.0/28     *[Direct/0] 00:11:52
                    > via ge-0/0/2.100
81.125.65.1/32     *[Local/0] 00:11:52
                      Local via ge-0/0/2.100
81.125.66.140/30   *[Direct/0] 00:11:52
                    > via ge-0/0/0.20
81.125.66.141/32   *[Local/0] 00:11:52
                      Local via ge-0/0/0.20
172.16.56.24/30    *[Direct/0] 00:11:52
                    > via ge-0/0/1.10
172.16.56.25/32    *[Local/0] 00:11:52
                      Local via ge-0/0/1.10
192.168.120.0/24   *[BGP/170] 00:02:17, localpref 100
                      AS path: 65010 I
                    > to 172.16.56.26 via ge-0/0/1.10
192.168.121.0/24   *[BGP/170] 00:02:17, localpref 100
                      AS path: 65010 I
                    > to 172.16.56.26 via ge-0/0/1.10
192.168.122.0/24   *[BGP/170] 00:02:17, localpref 100
                      AS path: 65010 I
                    > to 172.16.56.26 via ge-0/0/1.10
192.168.123.0/24   *[BGP/170] 00:02:17, localpref 100
                      AS path: 65010 I
                    > to 172.16.56.26 via ge-0/0/1.10
192.168.124.0/24   *[BGP/170] 00:02:17, localpref 100
                      AS path: 65010 I
                    > to 172.16.56.26 via ge-0/0/1.10
```

The last item to verify for this task is the advertisement of the default route towards the private-wan domain. To verify that the vSRX is advertising the default route, we issue the following command:

```
jncie@vSRX> show route advertising-protocol bgp 172.16.56.26

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    I
```

Task 5.2: vSRX Policy Configuration

- The hosts in the dc-zone are allowed to access the private-wan zone. The hosts can also use HTTP, HTTPS and DNS to access any resource in the untrust zone. All other traffic should be dropped and logged locally on the device.
- From the private-wan zone into the dc-zone, the management subnet 192.168.123.0/24 should be allowed to use any of the following protocols:
    - SSH, HTTP, HTTPS, ICMP, DNS and TCP 3389

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- From the private-wan zone into the dc-zone, all hosts should be allowed the following protocols:
    - HTTP, HTTPS and DNS
- From the untrust zone, any host should be able to use HTTP and HTTPS to access the webserver configured with IP 81.125.65.3.
- You are not allowed to configure any address object for a specific security-zone.

Solution

vSRX:

Allowing hosts from the dc-zone access to the private-wan zone:

```
set security policies from-zone dc-zone to-zone private-wan policy 1 match source-address any
set security policies from-zone dc-zone to-zone private-wan policy 1 match destination-address
any
set security policies from-zone dc-zone to-zone private-wan policy 1 match application any
set security policies from-zone dc-zone to-zone private-wan policy 1 then permit
```

Allowing hosts from the dc-zone HTTP, HTTPS and DNS access to any resource in the untrust zone:

```
set security policies from-zone dc-zone to-zone untrust policy 1 match source-address any
set security policies from-zone dc-zone to-zone untrust policy 1 match destination-address any
set security policies from-zone dc-zone to-zone untrust policy 1 match application junos-http
set security policies from-zone dc-zone to-zone untrust policy 1 match application junos-https
set security policies from-zone dc-zone to-zone untrust policy 1 match application junos-dns-
tcp
set security policies from-zone dc-zone to-zone untrust policy 1 match application junos-dns-
udp
set security policies from-zone dc-zone to-zone untrust policy 1 then permit
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Denying other traffic from the dc-zone to the untrust zone:

```
set security policies from-zone dc-zone to-zone untrust policy 2 match source-address any
set security policies from-zone dc-zone to-zone untrust policy 2 match destination-address any
set security policies from-zone dc-zone to-zone untrust policy 2 match application any
set security policies from-zone dc-zone to-zone untrust policy 2 then deny
```

To have the vSRX log the traffic that was denied, we add the 'then log session-init' statement to the previously configured deny rule:

```
set security policies from-zone dc-zone to-zone untrust policy 2 then log session-init
```

The SRX will include the string 'RT_FLOW_SESSION' in every message that is logged. We configure a separate log called 'traffic-log' and make sure all messages, that contain the 'RT_FLOW_SESSION' string, are written to this file:

```
set system syslog file traffic-log any any
set system syslog file traffic-log match RT_FLOW_SESSION
```

To allow the required traffic from the private-wan zone into the dc-zone, we start out with the configuration of an address-object. Since we are not allowed to create an address object that is specific to any zone, we configure the address object in the global address book:

```
set security address-book global address remote-mgmt 192.168.123.0/24
```

We then create the policy that is required for this task:

```
set security policies from-zone private-wan to-zone dc-zone policy 1 match source-address
remote-mgmt
set security policies from-zone private-wan to-zone dc-zone policy 1 match destination-address
any
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
ssh
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
http
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
https
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
icmp-all
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
dns-udp
set security policies from-zone private-wan to-zone dc-zone policy 1 match application junos-
dns-tcp
set security policies from-zone private-wan to-zone dc-zone policy 1 match application
tcp_3389
set security policies from-zone private-wan to-zone dc-zone policy 1 then permit
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

We were not able to use any of the Junos OS' default applications to allow TCP port 3389. Therefore, we referenced application tcp_3389 in the newly created policy. For this to work, we will need to configure this application as well:

```
set applications application tcp_3389 protocol tcp
set applications application tcp_3389 destination-port 3389
```

We create address objects for the subnets in the private-wan domain:

```
set security address-book global address private-wan-120 192.168.120.0/24
set security address-book global address private-wan-121 192.168.121.0/24
set security address-book global address private-wan-122 192.168.122.0/24
set security address-book global address private-wan-123 192.168.123.0/24
set security address-book global address private-wan-124 192.168.124.0/24
```

We then assign the subnets to a single address-set:

```
set security address-book global address-set private-wan address private-wan-120
set security address-book global address-set private-wan address private-wan-121
set security address-book global address-set private-wan address private-wan-122
set security address-book global address-set private-wan address private-wan-123
set security address-book global address-set private-wan address private-wan-124
```

To allow the specified protocols for all hosts inside the private-wan zone into the dc-zone, we configure the following rule:

```
set security policies from-zone private-wan to-zone dc-zone policy 2 match source-address
private-wan
set security policies from-zone private-wan to-zone dc-zone policy 2 match destination-address
any
set security policies from-zone private-wan to-zone dc-zone policy 2 match application junos-
http
set security policies from-zone private-wan to-zone dc-zone policy 2 match application junos-
https
set security policies from-zone private-wan to-zone dc-zone policy 2 match application junos-
dns-udp
set security policies from-zone private-wan to-zone dc-zone policy 2 match application junos-
dns-tcp
set security policies from-zone private-wan to-zone dc-zone policy 2 then permit
```

Creating an address entry in the global address book that will enable us to create a policy that involved the 83.125.65.3/32 host:

```
set security address-book global address web-server 81.125.65.3/32
```

Allowing HTTP and HTTPS inbound to a specific host located in the dc-zone:

```
set security policies from-zone untrust to-zone dc-zone policy 1 match source-address any
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security policies from-zone untrust to-zone dc-zone policy 1 match destination-address
web-server
set security policies from-zone untrust to-zone dc-zone policy 1 match application junos-http
set security policies from-zone untrust to-zone dc-zone policy 1 match application junos-https
set security policies from-zone untrust to-zone dc-zone policy 1 then permit
```

Verification

We issue the following commands to verify the policies on the vSRX:

```
jncie@vSRX> show security policies
Default policy: deny-all
From zone: dc-zone, To zone: private-wan
  Policy: 1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
From zone: dc-zone, To zone: untrust
  Policy: 1, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Applications: junos-http, junos-https, junos-dns-tcp, junos-dns-udp
    Action: permit
  Policy: 2, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 2
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: deny, log
From zone: private-wan, To zone: dc-zone
  Policy: 1, State: enabled, Index: 7, Scope Policy: 0, Sequence number: 1
    Source addresses: remote-mgmt
    Destination addresses: any
    Applications: junos-ssh, junos-http, junos-https, junos-icmp-all, junos-dns-udp, junos-
dns-tcp, tcp_3389
    Action: permit
  Policy: 2, State: enabled, Index: 8, Scope Policy: 0, Sequence number: 2
    Source addresses: private-wan
    Destination addresses: any
    Applications: junos-http, junos-https, junos-dns-udp, junos-dns-tcp
    Action: permit
From zone: untrust, To zone: dc-zone
  Policy: 1, State: enabled, Index: 9, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: web-server
    Applications: junos-http, junos-https
    Action: permit
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Instead of verifying the rule by examining the output that is shown above, we can also consider a more thorough approach. An example would be generating traffic that conforms to one of the rules that we have configured. For example, from the private-wan domain, we have allowed SSH from subnet 192.168.123.0/24 into the dc-zone. What we can do is move over to vMX4 and use SSH to a resource that is reachable via the vSRX. In this example, we issue SSH from the routing-instance private-wan to an address that is located behind the firewall inside the dc-zone:

Note: The following command might not receive a complete response from to establish the session, but the session should be created through the vSRX and while the SSH command is trying to connect you can verify the session commands. If you are concerned with connectivity, you can use the ping utility instead of the SSH command.

```
jncie@vMX4> ssh routing-instance private-wan 81.125.65.2 source 192.168.123.1

The authenticity of host '81.125.65.2 (81.125.65.2)' can't be established.
ECDSA key fingerprint is 07:bc:81:b6:74:06:f2:bc:40:82:b2:7e:4e:f8:e4:20.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '81.125.65.2' (ECDSA) to the list of known hosts.
Password:
Last login: Tue May 16 14:10:03 2017 from 10.10.20.254
--- JUNOS 16.1R3-S2.1 Kernel 64-bit  JNPR-10.3-20161130.340898_build
```

In the previous output, we logged into the vMX4 from the vMX4. We started inside the routing-instance *private-wan*, traversed the firewall and ended in the routing-instance *dc-zone*.

As a result of this, we can see the following session on the vSRX:

```
jncie@vSRX> show security flow session
..
Session ID: 849, Policy name: 1/7, Timeout: 1790, Valid
  In: 192.168.123.1/50409 --> 81.125.65.2/22;tcp, If: ge-0/0/1.10, Pkts: 18, Bytes: 3951
  Out: 81.125.65.2/22 --> 192.168.123.1/50409;tcp, If: ge-0/0/2.100, Pkts: 16, Bytes: 3527
Total sessions: 3
```

It is also possible to examine the session in greater detail:

```
jncie@vSRX> show security flow session destination-prefix 81.125.65.2 extensive

Session ID: 849, Status: Normal
Flag: 0x40
Policy name: 1/7
Source NAT pool: Null, Application: junos-ssh/22
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
Maximum timeout: 1800, Current timeout: 1688
Session State: Valid
Start time: 348624, Duration: 577
   In: 192.168.123.1/50409 --> 81.125.65.2/22;tcp,
    Interface: ge-0/0/1.10,
    Session token: 0x6, Flag: 0x21
    Route: 0x110010, Gateway: 172.16.56.26, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 48, Bytes: 6211
   Out: 81.125.65.2/22 --> 192.168.123.1/50409;tcp,
    Interface: ge-0/0/2.100,
    Session token: 0x9, Flag: 0x20
    Route: 0x120010, Gateway: 81.125.65.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 39, Bytes: 6059
Total sessions: 1
```

Here we see all the characteristics related to this flow. Listed at the top is the policy that is permitting the traffic. The rest of the information (start time, amount of traffic, inbound and outbound interface) can be quite helpful in case you are troubleshooting.

The last thing we need to verify is the logging we configured for the deny rule from the dc-zone to untrust zone. To this end, we generate traffic to a destination currently not permitted by any policy on the vMX4 router:

```
jncie@vMX4> ping routing-instance dc-zone 5.5.5.5 count 1

PING 5.5.5.5 (5.5.5.5): 56 data bytes

--- 5.5.5.5 ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

jncie@vMX4>
```

These packets will be subjected to the policy that is configured to deny and log packets. Therefore, we should see a corresponding entry in the log we configured:

```
jncie@vSRX> show log traffic-log

May 16 14:10:48  vSRX RT_FLOW: RT_FLOW_SESSION_DENY: session denied 81.125.65.2/0->
5.5.5.5/40713 icmp 1(8) 2 dc-zone untrust UNKNOWN UNKNOWN N/A(N/A) ge-0/0/2.100 UNKNOWN policy
deny
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 5.3: vSRX NAT

- From the private-wan zone, all the 192.168.12x.0/24 subnets should be allowed to access resources inside the untrust zone. Inside the untrust zone, users should appear with the vSRX interface address.
- When users in the untrust zone try to contact 81.125.65.6, the vSRX should forward and allow that traffic towards host 192.168.123.123. The vSRX should automatically create a reverse mapping. Allow ICMP from the untrust to zone towards the private-wan zone. Allow ICMP and HTTPS from the private-wan to the untrust zone.
- When users in the untrust zone try to contact a webserver located at 81.122.31.2 using HTTPS or ICMP, traffic should be forwarded to 81.125.65.2.

Solution

vSRX:

NAT requires you to configure a NAT rule under the [security nat] configuration stanza. Apart from this, you may need to configure a policy to allow traffic.

```
For this task, all three assignments require the configuration of a NAT rule as well as a
policy.
```

Source NAT:

The first assignment task requires us to configure source NAT for a subnet located in the private-wan zone. Therefore, we configure the NAT rule from zone private-wan to the untrust zone:

```
set security nat source rule-set source-nat-interface from zone private-wan
set security nat source rule-set source-nat-interface to zone untrust
```

We narrow down the NAT rule so that only 1 subnet will qualify for translation:

```
set security nat source rule-set source-nat-interface rule src-nat-all match source-address
192.168.124.0/24
```

Since the users have to appear with the SRX interface address as translated source address, we specify 'then source nat interface' as the action for this NAT rule:

```
set security nat source rule-set source-nat-interface rule src-nat-all then source-nat
interface
```

After the configuration of the NAT rule, we configure an address object and a policy to allow traffic from the private-wan zone to access any resource inside the untrust zone:

```
set security address-book global address 192.168.124/24 192.168.124.0/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security policies from-zone private-wan to-zone untrust policy source-nat match source-
address 192.168.124/24
set security policies from-zone private-wan to-zone untrust policy source-nat match
destination-address any
set security policies from-zone private-wan to-zone untrust policy source-nat match
application any
set security policies from-zone private-wan to-zone untrust policy source-nat then permit
```

Note that inside the Junos OS flow module, source NAT takes place after the policy lookup. Because of this, we referenced the original subnet in our policy configuration.

Static NAT:

Since the vSRX has to automatically create a reverse mapping, we know that we have to configure static NAT. Packets with a destination IP address of 81.125.65.6 need to have that address translated to IP address 192.168.123.123 when they arrive on the vSRX from the untrust zone. We start out configuring a static NAT rule from the untrust zone, matching the destination IP address of 81.125.65.6:

```
set security nat static rule-set static-nat from zone untrust
set security nat static rule-set static-nat rule static-nat match destination-address
81.125.65.6/32
```

Next, we specify the address what packets inbound from the untrust zone should be translated to:

```
set security nat static rule-set static-nat rule static-nat then static-nat prefix
192.168.123.123/32
```

After the configuration of a NAT rule, we configure the required security policies and address objects to allow traffic:

```
set security address-book global address ws-private-wan 192.168.123.123/32
set security policies from-zone private-wan to-zone untrust policy static-nat match source-
address ws-private-wan
set security policies from-zone private-wan to-zone untrust policy static-nat match
destination-address any
set security policies from-zone private-wan to-zone untrust policy static-nat match
application junos-https
set security policies from-zone private-wan to-zone untrust policy static-nat match
application junos-icmp-all
set security policies from-zone private-wan to-zone untrust policy static-nat then permit

set security policies from-zone untrust to-zone private-wan policy static-nat match source-
address any
set security policies from-zone untrust to-zone private-wan policy static-nat match
destination-address ws-private-wan
set security policies from-zone untrust to-zone private-wan policy static-nat match
application junos-https
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security policies from-zone untrust to-zone private-wan policy static-nat match
application junos-icmp-all
set security policies from-zone untrust to-zone private-wan policy static-nat then permit
```

Note that inside the Junos OS flow module, static NAT takes place before the policy lookup. The rule is configured from the untrust to the private-wan zone. Therefore, we reference the address that is configured on the host in the private-wan zone.

Destination NAT:

Packets from the untrust zone destined to IP address 81.122.31.2 should be forwarded to 81.125.65.2. To this end, we configure a destination NAT-rule. Inside the rule, we reference the original destination IP address from packets coming from the untrust zone:

```
set security nat destination rule-set inbound-nat from zone untrust
set security nat destination rule-set inbound-nat rule dst-nat match destination-address
81.122.31.2/32
```

We configure the rule for destination NAT and reference a pool named dst-nat:

```
set security nat destination rule-set inbound-nat rule dst-nat then destination-nat pool dst-
nat
```

The pool previously referenced, dst-nat, still needs to be configured. We configure the pool with the address that packets, inbound from the untrust zone, need to have their destination IP translated to:

```
set security nat destination pool dst-nat address 81.125.65.2/32
```

Finally, we configure an address object and security policy that permits the traffic to the host inside the dc-zone:

```
set security address-book global address ws-dc-zone 81.125.65.2/32
set security policies from-zone untrust to-zone dc-zone policy destination-nat match source-
address any
set security policies from-zone untrust to-zone dc-zone policy destination-nat match
destination-address ws-dc-zone
set security policies from-zone untrust to-zone dc-zone policy destination-nat match
application junos-https
set security policies from-zone untrust to-zone dc-zone policy destination-nat match
application junos-icmp-all
set security policies from-zone untrust to-zone dc-zone policy destination-nat then permit
```

Note that inside the Junos OS flow module, destination NAT takes place before the policy lookup. Since incoming packets will have their destination IP translated before the policy lookup, we match the IP address that is configured on the host inside the dc-zone.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

For this task, we configured three different types of NAT. We will verify them one by one.

Source NAT:

The first assignment task required us to configure source NAT for the 192.168.12.0/24 subnet.

First, we check the newly configured security policy:

```
jncie@vSRX> show security policies from-zone private-wan to-zone untrust policy-name source-
nat

From zone: private-wan, To zone: untrust
  Policy: source-nat, State: enabled, Index: 14, Scope Policy: 0, Sequence number: 1
    Source addresses: 192.168.124/24
    Destination addresses: any
    Applications: any
    Action: permit
```

Then, we examine the summary for source NAT:

```
jncie@vSRX> show security nat source summary

Total port number usage for port translation pool: 0
Maximum port number for port translation pool: 33554432
Total pools: 0

Total rules: 1
Rule name          Rule set       From              To                Action
src-nat-all        source-nat-interface private-wan untrust           interface
```

This tells us 1 source NAT rule, named src-nat-all, is active. The rule is active from the private-wan zone to the untrust zone and the action is ‘interface’. This means that the egress interface's IP address of the vSRX will be used to perform source NAT. We can also examine the following command for additional output:

```
jncie@vSRX> show security nat source rule src-nat-all

source NAT rule: src-nat-all          Rule-set: source-nat-interface
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : private-wan
  To zone                    : untrust
  Match
    Source addresses         : 192.168.124.0   - 192.168.124.255
  Action                       : interface
    Persistent NAT type        : N/A
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
  Translation hits              : 0
    Successful sessions         : 0
    Failed sessions             : 0
  Number of sessions            : 0
```

Here we can see that the rule matches on a specific subnet. For additional analysis, we can setup an SSH session like this (keep in mind the SSH session may not complete, but you should still see a successful session):

```
jncie@vMX4> ssh routing-instance private-wan 95.0.0.1 source 192.168.124.1

Password:
Last login: Tue May 16 20:04:05 2017 from 81.125.66.141
--- JUNOS 16.1R3-S2.1 Kernel 64-bit  JNPR-10.3-20161130.340898_build
jncie@vMX4>
```

We have set up a TCP session that transits the firewall from the private-wan zone to the untrust zone. When we examine the rule again, we can see that several counters have been incremented:

```
jncie@vSRX> show security nat source rule src-nat-all

source NAT rule: src-nat-all          Rule-set: source-nat-interface
  Rule-Id                     : 1
  Rule position               : 1
  From zone                   : private-wan
  To zone                     : untrust
  Match
    Source addresses          : 192.168.124.0   - 192.168.124.255
  Action                        : interface
    Persistent NAT type         : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
  Translation hits              : 1
    Successful sessions         : 1
    Failed sessions             : 0
  Number of sessions            : 1
```

In addition to seeing several counters increment, we can also examine the session table. Using the following command, we can examine any active sessions towards the 95.0.0.1 address:

```
jncie@vSRX> show security flow session nat

Session ID: 1975, Policy name: 1/7, Timeout: 1790, Valid
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  In: 192.168.124.1/61118 --> 95.0.0.1/22;tcp, If: ge-0/0/1.10, Pkts: 187, Bytes: 161
  Out: 95.0.0.1/22 --> 81.125.66.141/23591;tcp, If: ge-0/0/0.20, Pkts: 159, Bytes: 34
Total sessions: 1
```

The 'In' is showing the incoming flow and the 'Out' is displaying the reverse flow. We can see that 192.168.124.1 is starting an SSH session with 95.0.0.1. The return flow is sourced from 95.0.0.1 to 81.125.66.141, which is the vSRX interface IP.

While on the vMX4 router, we can also see a login that is sourced from the vSRX:

```
jncie@vMX4> show system users | match 141
jncie    pts/2    81.125.66.141                 8:04PM      - -cli (cli)
```

Static NAT:

We first check the newly configured policies:

```
jncie@vSRX> show security policies from-zone untrust to-zone private-wan policy-name static-
nat detail

Policy: static-nat, action-type: permit, State: enabled, Index: 12, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: untrust, To zone: private-wan
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    ws-private-wan(global): 192.168.123.123/32
  Application: junos-https
    IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [443-443]
  Application: junos-icmp-all
    IP protocol: icmp, ALG: 0, Inactivity timeout: 60
      ICMP Information: type=255, code=0
  Per policy TCP Options: SYN check: No, SEQ check: No

jncie@vSRX> show security policies from-zone private-wan to-zone untrust policy-name static-
nat detail

Policy: static-nat, action-type: permit, State: enabled, Index: 11, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: private-wan, To zone: untrust
  Source addresses:
    ws-private-wan(global): 192.168.123.123/32
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  any-ipv6(global): ::/0
Application: junos-https
  IP protocol: tcp, ALG: 0, Inactivity timeout: 1800
    Source port range: [0-0]
    Destination port range: [443-443]
Application: junos-icmp-all
  IP protocol: icmp, ALG: 0, Inactivity timeout: 60
    ICMP Information: type=255, code=0
Per policy TCP Options: SYN check: No, SEQ check: No
```

Apart from the policy name, we see nothing that indicates static NAT will be applied. To see what static NAT rules are currently active, we issue the following command:

```
jncie@vSRX> show security nat static rule all

Total static-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Static NAT rule: static-nat           Rule-set: static-nat
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : untrust
  Destination addresses     : 81.125.65.6
  Host addresses            : 192.168.123.123
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 938
    Successful sessions     : 864
    Failed sessions         : 0
  Number of sessions        : 0
```

To be able to see the translation happening, we generate some packets from the vMX4 router:

```
jncie@vMX4> ping 81.125.65.6 routing-instance www count 1

PING 81.125.65.6 (81.125.65.6): 56 data bytes
64 bytes from 81.125.65.6: icmp_seq=0 ttl=63 time=175.852 ms

--- 81.125.65.6 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 175.852/175.852/175.852/0.000 ms
```

After this, we can examine the session table:

```
jncie@vSRX> show security flow session nat

Session ID: 1978, Policy name: static-nat/12, Timeout: 4, Valid
  In: 81.125.66.142/0 --> 81.125.65.6/17684;icmp, If: ge-0/0/0.20, Pkts: 1, Bytes: 84
  Out: 192.168.123.123/17684 --> 81.125.66.142/0;icmp, If: ge-0/0/1.10, Pkts: 1, Bytes: 84
Total sessions: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Destination NAT:

We verify the newly configured security policy:

```
jncie@vSRX> show security policies from-zone untrust to-zone dc-zone policy-name destination-
nat

From zone: untrust, To zone: dc-zone
  Policy: destination-nat, State: enabled, Index: 13, Scope Policy: 0, Sequence number: 2
    Source addresses: any
    Destination addresses: ws-dc-zone
    Applications: junos-https, junos-icmp-all
    Action: permit
```

To see what destination NAT rules are active, we issue the following command:

```
jncie@vSRX> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0

Destination NAT rule: dst-nat                Rule-set: inbound-nat
  Rule-Id                     : 1
  Rule position               : 1
  From zone                   : untrust
    Destination addresses     : 81.122.31.2     - 81.122.31.2
  Action                      : dst-nat
  Translation hits            : 8
    Successful sessions       : 8
    Failed sessions           : 0
  Number of sessions          : 0
```

We initiate a translation from vMX4:

```
jncie@vMX4> ping 81.122.31.2 routing-instance www count 1

PING 81.122.31.2 (81.122.31.2): 56 data bytes
64 bytes from 81.122.31.2: icmp_seq=0 ttl=63 time=10.031 ms

--- 81.122.31.2 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.031/10.031/10.031/0.000 ms
```

And finally, we verify the session table:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vSRX> show security flow session nat

Session ID: 1981, Policy name: destination-nat/13, Timeout: 2, Valid
  In: 81.125.66.142/0 --> 81.122.31.2/27668;icmp, If: ge-0/0/0.20, Pkts: 1, Bytes: 84
  Out: 81.125.65.2/27668 --> 81.125.66.142/0;icmp, If: ge-0/0/2.100, Pkts: 1, Bytes: 84
Total sessions: 1
```

Task 5.4: vSRX Screens

- Protect the session table from inbound traffic from the untrust zone to any other zone. There should be no more than 500 sessions per source and per destination IP.
- The firewall should also protect resources from an ICMP flood that is greater than 500 packets per second and an IP address sweeps that are greater than 10 ICMP packets arriving at a 500ms time interval.

Solution

vSRX:

Screens are configured using IDS-options. We can apply only a single screen to a security zone. To meet the objectives of this task, we need to configure several ids-options for a single screen.

We configure the screen to protect against an ICMP ip-sweep and set the threshold to 500000. This is because the parameter is defined in microseconds:

```
set security screen ids-option from-untrust icmp ip-sweep threshold 500000
```

We add protection from an ICMP flood to the screen:

```
set security screen ids-option from-untrust icmp flood threshold 500
```

After this, we add the ids-option that limits sessions for source and destination IP to the screen:

```
set security screen ids-option from-untrust limit-session source-ip-based 500
set security screen ids-option from-untrust limit-session destination-ip-based 500
```

Finally, we apply the screen to the untrust zone:

```
set security zones security-zone untrust screen from-untrust
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

We can verify the created ids-options by issuing the following command:

```
jncie@vSRX> show security screen ids-option from-untrust

Screen object status:

Name                                        Value
  ICMP flood threshold                      500
  ICMP address sweep threshold              500000
  Session source limit threshold            500
  Session destination limit threshold       500
```

To check what ids-options are associated with the untrust zone, we issue the following command:

```
jncie@vSRX> show security zones untrust

Security zone: untrust
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: from-untrust
  Interfaces bound: 1
  Interfaces:
    ge-0/0/0.20
```

To see if the configured screen has detected any attacks, we can issue the following command:

```
jncie@vSRX> show security screen statistics zone untrust

Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  UDP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
TCP no flag                                0
IP unknown protocol                        0
IP bad options                             0
IP record route option                     0
IP timestamp option                        0
IP security option                         0
IP loose source route option               0
IP strict source route option              0
IP stream option                           0
ICMP fragment                              0
ICMP large packet                          0
TCP SYN FIN                                0
TCP FIN no ACK                             0
  Source session limit                     0
  TCP SYN-ACK-ACK proxy                    0
  IP block fragment                        0
  Destination session limit                0
  IPv6 extension header                    0
  IPv6 extension hop by hop option         0
  IPv6 extension destination option        0
  IPv6 extension header limit              0
  IPv6 malformed header                    0
  ICMPv6 malformed packet                  0
```

Note that the previous command will output statistics for the entire zone. If we would have had multiple interfaces configured for the untrust zone, and if we wanted to check the statistics on a per interface basis, we could issue the following command:

```
jncie@vSRX> show security screen statistics interface ge-0/0/0.20

Screen statistics:
IDS attack type                            Statistics
  ICMP flood                               0
  UDP flood                                0
  TCP winnuke                              0
  TCP port scan                            0
  UDP port scan                            0
  ICMP address sweep                       0
  TCP sweep                                0
  UDP sweep                                0
  IP tear drop                             0
  TCP SYN flood                            0
  IP spoofing                              0
  ICMP ping of death                       0
  IP source route option                   0
  TCP land attack                          0
  TCP SYN fragment                         0
  TCP no flag                              0
  IP unknown protocol                      0
  IP bad options                           0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
IP record route option                    0
IP timestamp option                       0
IP security option                        0
IP loose source route option              0
IP strict source route option             0
IP stream option                          0
ICMP fragment                             0
ICMP large packet                         0
TCP SYN FIN                               0
TCP FIN no ACK                            0
Source session limit                      0
TCP SYN-ACK-ACK proxy                     0
IP block fragment                         0
Destination session limit                 0
IPv6 extension header                     0
IPv6 extension hop by hop option          0
IPv6 extension destination option         0
IPv6 extension header limit               0
IPv6 malformed header                     0
ICMPv6 malformed packet                   0
```

Task 5.5: vSRX RE Protection

- Configure RE protection on the SRX using a stateless firewall filter.
- The SRX should be able set up a BGP session.
- Configure a separate filter for the monitoring and management of the switch. Allow the following protocols: SSH, FTP, NTP, Radius, DNS, SNMP, traceroute and ICMP.
- Restrict SSH, Telnet, FTP, NTP, DNS, SNMP and Radius traffic to subnet 10.10.20.0/24 and 192.168.123.0/24.
- Rate limit both ICMP and traceroute to 1 Mbit/s. Rate limit both FTP and SSH to 5 Mbit/s.
- Discard all other traffic. A counter should be incremented for every discarded packet and the header of the discarded packet should be logged locally on the device.

Solution

vSRX:

Allowing the vSRX to setup BGP sessions:

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

Because several terms will need to allow the exact same subnets, we configure a prefix-list. This eases configuration and verification later on:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options prefix-list management 10.10.20.0/24
set policy-options prefix-list management 192.168.123.0/24
```

The first term in which we use the prefix-list is the term that will allow SSH:

```
set firewall family inet filter protect.re term ssh from source-prefix-list management
set firewall family inet filter protect.re term ssh from protocol tcp
set firewall family inet filter protect.re term ssh from port ssh
set firewall family inet filter protect.re term ssh then policer 5mbit
set firewall family inet filter protect.re term ssh then accept
```

Allowing FTP:

```
set firewall family inet filter protect.re term ftp from source-prefix-list management
set firewall family inet filter protect.re term ftp from protocol tcp
set firewall family inet filter protect.re term ftp from port ftp
set firewall family inet filter protect.re term ftp from port ftp-data
set firewall family inet filter protect.re term ftp then policer 5mbit
set firewall family inet filter protect.re term ftp then accept
```

Allowing NTP:

```
set firewall family inet filter protect.re term ntp from source-prefix-list management
set firewall family inet filter protect.re term ntp from protocol udp
set firewall family inet filter protect.re term ntp from port ntp
set firewall family inet filter protect.re term ntp then accept
```

Allowing Radius:

```
set firewall family inet filter protect.re term radius from source-prefix-list management
set firewall family inet filter protect.re term radius from protocol udp
set firewall family inet filter protect.re term radius from source-port radius
set firewall family inet filter protect.re term radius then accept
```

Allowing SNMP:

```
set firewall family inet filter protect.re term snmp from source-prefix-list management
set firewall family inet filter protect.re term snmp from protocol udp
set firewall family inet filter protect.re term snmp from destination-port snmp
set firewall family inet filter protect.re term snmp then accept
```

Allowing DNS for TCP as well as UDP:

```
set firewall family inet filter protect.re term dns from source-prefix-list management
set firewall family inet filter protect.re term dns from protocol tcp
set firewall family inet filter protect.re term dns from protocol udp
set firewall family inet filter protect.re term dns from source-port domain
set firewall family inet filter protect.re term dns then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Allowing traceroute:

```
set firewall family inet filter protect.re term traceroute from protocol udp
set firewall family inet filter protect.re term traceroute from destination-port 33434-33534
set firewall family inet filter protect.re term traceroute then policer 1mbit
set firewall family inet filter protect.re term traceroute then accept
```

Allowing ICMP:

```
set firewall family inet filter protect.re term icmp from protocol icmp
set firewall family inet filter protect.re term icmp then policer 1mbit
set firewall family inet filter protect.re term icmp then accept
```

Silently dropping all other traffic:

```
set firewall family inet filter protect.re term discard then discard
```

Having the counter ‘protect.re’ increment for every packet dropped:

```
set firewall family inet filter protect.re term discard then count protect.re
```

Logging the packet header of every discarded packet locally on the device:

```
set firewall family inet filter protect.re term discard then log
```

Creating the policers that were referenced in the previously configured firewall filter:

```
set firewall policer 1mbit if-exceeding bandwidth-limit 1m
set firewall policer 1mbit if-exceeding burst-size-limit 125k
set firewall policer 1mbit then discard
set firewall policer 5mbit if-exceeding bandwidth-limit 1m
set firewall policer 5mbit if-exceeding burst-size-limit 125k
set firewall policer 5mbit then discard
```

Applying the firewall filter:

```
set interfaces lo0 unit 0 family inet filter input protect.re
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

The content of a firewall filter can be inspected by verifying the configuration. To see if the firewall filter has been applied successfully, we can issue the following command:

```
jncie@vSRX> show interfaces lo0.0 detail | match Filter

      Input Filters: protect.re
```

Task 5.6: Limiting Bandwidth using a Stateless Firewall Filter

- Traffic from 192.168.120.1/32 towards 81.125.65.2/32 should bypass the flow engine.
- Configure a stateless firewall filter that permits both ICMP and TCP port 4567.
- TCP traffic should be rate-limited to 100m.

Solution

vSRX:

Under the [security forwarding-options family], we can put the SRX into packet mode for certain address families. Any configuration inside this stanza will affect the entire system. For this task, we have to isolate a single source and destination subnet. We can do this using a firewall filter. The action 'then packet-mode' will instruct the SRX to bypass the flow daemon for the traffic that the firewall filter matches.

We only need to police the TCP traffic. Therefore, we first create a term to match the ICMP traffic and turn on packet mode for this traffic:

```
set firewall family inet filter bypass-flowd term per-packet-icmp from source-address
192.168.120.1/32
set firewall family inet filter bypass-flowd term per-packet-icmp from destination-address
81.125.65.2/32
set firewall family inet filter bypass-flowd term per-packet-icmp from protocol icmp
set firewall family inet filter bypass-flowd term per-packet-icmp then packet-mode
set firewall family inet filter bypass-flowd term per-packet-icmp then accept
```

Then, we create a second term to match the TCP port 4567 traffic and apply both the policer and the packet-mode action to the traffic:

```
set firewall family inet filter bypass-flowd term per-packet-tcp from source-address
192.168.120.1/32
set firewall family inet filter bypass-flowd term per-packet-tcp from destination-address
81.125.65.2/32
set firewall family inet filter bypass-flowd term per-packet-tcp from protocol tcp
set firewall family inet filter bypass-flowd term per-packet-tcp from destination-port 4567
set firewall family inet filter bypass-flowd term per-packet-tcp then policer 100m
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set firewall family inet filter bypass-flowd term per-packet-tcp then packet-mode
set firewall family inet filter bypass-flowd term per-packet-tcp then accept
```

Firewall filters, by default, end with an implicit deny. This task did not instruct us to block any other traffic. Therefore, we create another term that accepts any other traffic:

```
set firewall family inet filter bypass-flowd term accept-rest then accept
```

We configure a 100Mbit policer:

```
set firewall policer 100m if-exceeding bandwidth-limit 100m
set firewall policer 100m if-exceeding burst-size-limit 625k
set firewall policer 100m then discard
```

Finally, we apply the rule on the interface that connects the firewall to the private-wan:

```
set interfaces ge-0/0/1 unit 10 family inet filter input bypass-flowd
```

Verification

Before the completion of this task, we would not have been able to exchange traffic:

```
jncie@vMX4> ping routing-instance private-wan source 192.168.120.1 81.125.65.2 rapid

PING 81.125.65.2 (81.125.65.2): 56 data bytes
.....
--- 81.125.65.2 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

After completing this task, we can exchange traffic between the two IP addresses:

```
jncie@vMX4> ping routing-instance private-wan source 192.168.120.1 81.125.65.2 rapid

PING 81.125.65.2 (81.125.65.2): 56 data bytes
!!!!!
--- 81.125.65.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.461/49.349/78.364/24.664 ms
```

When we check the policies in place between the private-wan and dc-zone, we can see that traffic between 192.168.120.1 and 81.125.65.2 is not permitted:

```
jncie@vSRX> show security policies from-zone private-wan to-zone dc-zone

From zone: private-wan, To zone: dc-zone
  Policy: ssh, State: enabled, Index: 10, Scope Policy: 0, Sequence number: 1
    Source addresses: remote-mgmt
    Destination addresses: any
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Applications: junos-ssh, junos-http, junos-https, junos-icmp-all, junos-dns-udp, junos-
dns-tcp, tcp_3389
    Action: permit
  Policy: 2, State: enabled, Index: 8, Scope Policy: 0, Sequence number: 2
    Source addresses: private-wan
    Destination addresses: any
    Applications: junos-http, junos-https, junos-dns-udp, junos-dns-tcp
    Action: permit

jncie@vSRX>
```

We have enabled traffic between the two hosts because of the firewall filter we applied to the interface:

```
jncie@vSRX> show interfaces ge-0/0/1.10 detail | match Filter

      Input Filters: bypass-flowd
```

This firewall filter ensures that the SRX flow daemon is bypassed.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Appendix – Chapter 6: Class of Service

Part 1: CoS Inside the Data Center

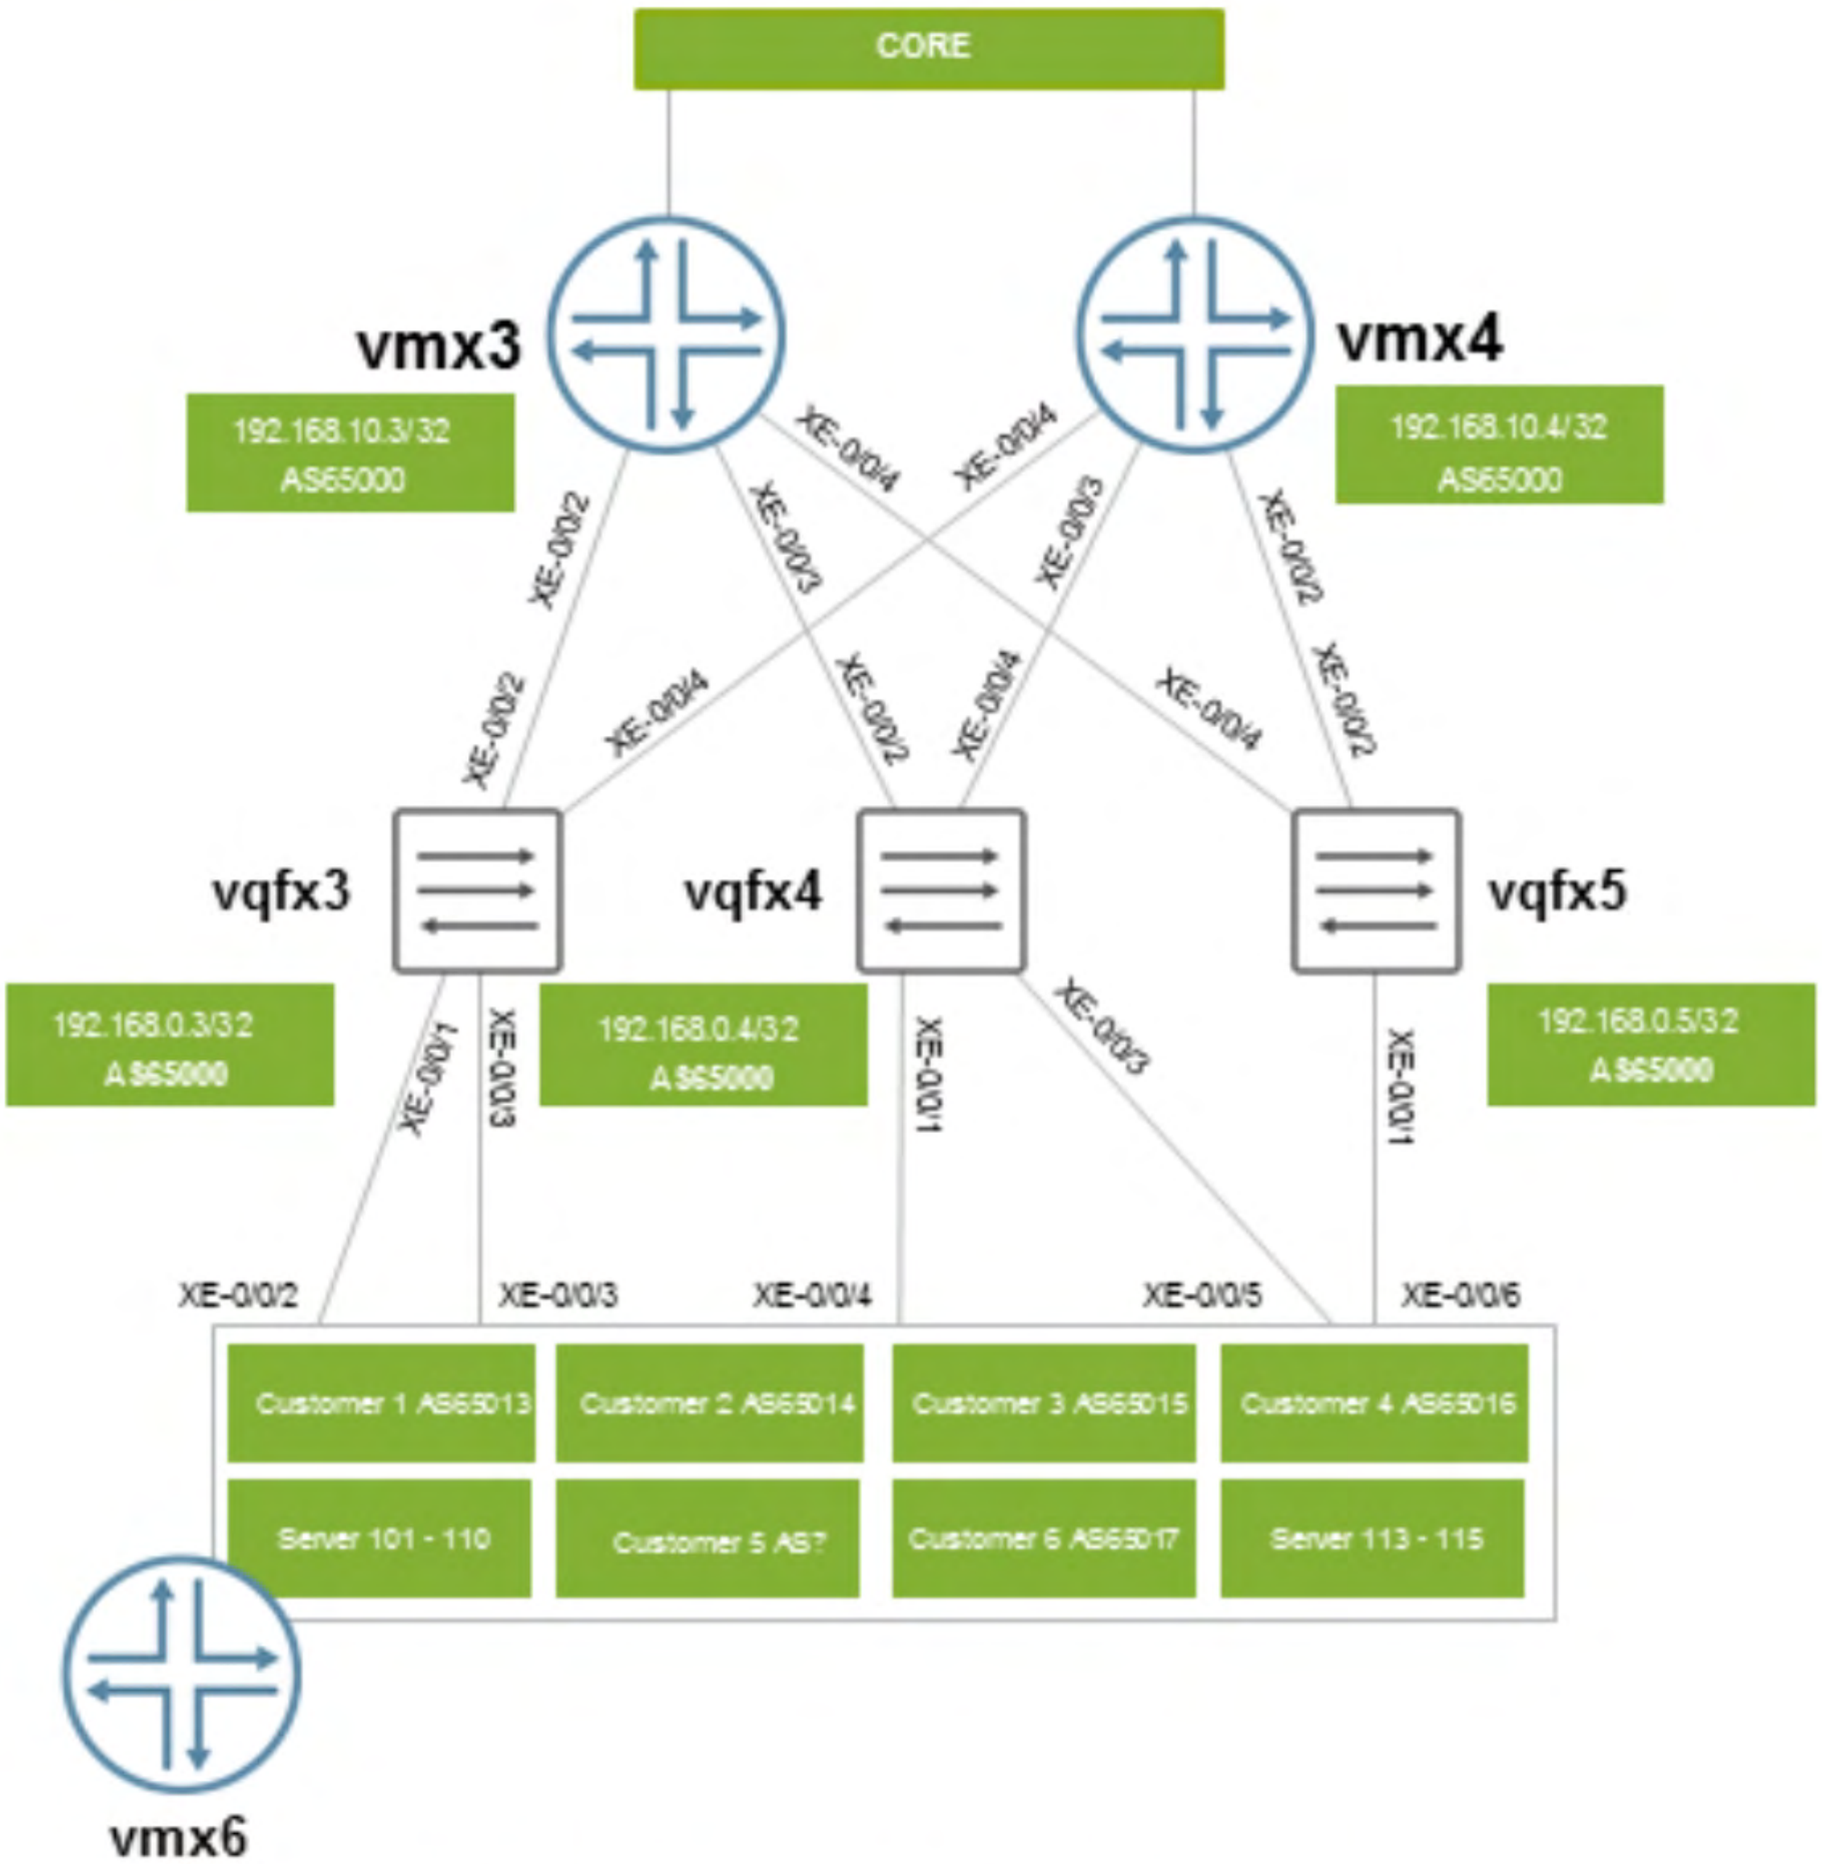

Figure 26 – DC network topology. The entire IP-Fabric has been preconfigured for you.

Note:

There are some differences in the way a vQFX and a QFX5100 handle CoS. After the solutions to task 6.6, you can find a special section called 'QFX5100 Class of Service' where these differences are highlighted.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 6.1: Forwarding Classes

- Configure the following forwarding classes on both vMX as well as vQFX:

Queue number	Class name
0	be
1	data
2	data-pri
3	rt
4	control

Solution

A forwarding class is a per-node setting. It is used to identify traffic that needs to receive similar treatment on a node.

To put the use of the forwarding class in perspective, the forwarding class itself does not really do anything. To assign traffic to different forwarding classes, we have to use classifiers (BA or MFC). And when traffic exits the node, the configured schedulers will determine how traffic inside a queue will be treated.

However, when we configure classifiers or schedulers later on, we will refer to the forwarding class name that we configure during this assignment.

During this assignment we only configure forwarding classes. This means we configure a name and tie that name to a queue. We configure the same forwarding classes on every node. Because the configuration for vMX and vQFX is exactly the same, we configure the following:

On every device:

```
set class-of-service forwarding-classes class be queue-num 0
set class-of-service forwarding-classes class data queue-num 1
set class-of-service forwarding-classes class data-pri queue-num 2
set class-of-service forwarding-classes class rt queue-num 3
set class-of-service forwarding-classes class control queue-num 4
```

Note: the assignment instructs us to configure a class named “be” and assign that to queue 0. We have to be very precise when we follow these instructions. If the class “be” would be configured as either “BE” or “Be”, we would fail the question on a JNCIE exam!

Verification

To verify what forwarding classes are active on a device, we issue the following command on vMX and vQFX:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
jncie@vMX3> show class-of-service forwarding-class

Forwarding class                  ID      Queue  Restricted queue  Fabric priority
  be                               0          0                 0             low
  data                             1          1                 1             low
  data-pri                         2          2                 2             low
  rt                               3          3                 3             low
  control                          4          4                 0             low

jncie@vQFX3> show class-of-service forwarding-class

Forwarding class                  ID      Queue  Policing priority  No-Loss
  be                              0          0             normal  Disabled
  rt                              1          3             normal  Disabled
  control                         2          4             normal  Disabled
  network-control                 3          7             normal  Disabled
  data                            4          1             normal  Disabled
  data-pri                        5          2             normal  Disabled
```

Task 6.2: Scheduling

- Configure the following schedulers on both vQFX as well as vMX:

Scheduler name	Transmit rate	Buffer-size	Priority
be	25	remainder	Low
data	30 percent	30 percent	low
data-pri	30 percent	30 percent	low
rt	10 percent	10 percent	Strict-high
control	5 percent	5 percent	low

- Apply the schedulers to all main interfaces.
- Ensure that it is impossible for a queue to get starved.

Solution

The schedulers control how traffic inside a single queue is treated. Every scheduler can be configured with a transmission rate, queue priority, buffer size and drop-profile. To complete this assignment, we configure the transmission rate, buffer-size and priority for the schedulers. The configuration takes place inside the [**class-of-service schedulers**] configuration stanza. Let's explain the different configuration options and how we can use them.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Transmit-rate:

With the '**transmit-rate rate *bps***' configuration command, we specify the rate in bits per second. If we use the '**transmit-rate percent *percent* '** configuration command, we allocate the scheduler a percentage of the available interface bandwidth.

It is important to note that these limits apply during times of congestion. When there is no congestion on an interface, schedulers configured with a transmit-rate will be able to use the maximum bandwidth.

On the MX, we can choose to add the optional **rate-limit** or **exact** configuration command to a scheduler. Adding **rate-limit** to a scheduler will ensure that the configured transmit rate will become the limit for this queue. Even when there is no congestion. The limit will be enforced through a policer action. When we configure the keyword **exact**, we also ensure that the configured transmit rate will become the limit for the queue in the absence of congestion. The difference when using **exact** is that this option will enforce the limit through a shaping action.

As an alternative to the '**transmit-rate rate**' and the '**transmit-rate percent**' statements, the MX also allows us to configure the '**transmit-rate remainder**' option. This simply means that the scheduler is assigned all bandwidth that is not yet configured to be allocated to other schedulers. This is not available on the QFX.

Buffer-size:

Packets are sent out 1 by 1. Packets that need to wait for another packet to be transmitted are stored in an output queue where they are buffered. Schedulers are used to allocate the buffer-size that should be available to a queue. We can allocate the buffer size as a percentage of the total buffer using the 'buffer-size percent *percent*' statement.

In addition to specifying the buffer size as a percentage, the MX can also be configured with temporal value. We can do this using the 'buffer-size temporal *microseconds*' to specify the buffer in microseconds.

Same as with the transmit rate, we can also choose to allocate the remaining buffer using the 'buffer-size remainder' configuration command. The main difference here is that this option is available on MX as well as QFX.

Priority:

Used to assign relative priorities to queues. We have different priorities available when comparing QFX to MX.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

On an MX, we can assign the following priorities:

- priority *low*
- priority *medium-low*
- priority *medium-high*
- priority *high*
- priority *strict-high*

On a QFX, we can assign the following priorities:

- priority *low*
- priority *strict-high*

Note that the strict high queue is always services before traffic in other queues. This might lead to traffic starvation for the other non-strict high queues. If you want to prevent starvation in other queues, you should rate-limit this queue.

On both vMX3 and vMX4:

```
set class-of-service schedulers be transmit-rate percent 25
set class-of-service schedulers be buffer-size remainder
set class-of-service schedulers be priority low

set class-of-service schedulers data transmit-rate percent 30
set class-of-service schedulers data buffer-size percent 30
set class-of-service schedulers data priority low

set class-of-service schedulers data-pri transmit-rate percent 30
set class-of-service schedulers data-pri buffer-size percent 30
set class-of-service schedulers data-pri priority low

set class-of-service schedulers rt transmit-rate percent 10
set class-of-service schedulers rt transmit-rate rate-limit
set class-of-service schedulers rt buffer-size percent 10
set class-of-service schedulers rt priority strict-high

set class-of-service schedulers control transmit-rate percent 5
set class-of-service schedulers control buffer-size percent 5
set class-of-service schedulers control priority low

set class-of-service scheduler-maps fabric forwarding-class be scheduler be
set class-of-service scheduler-maps fabric forwarding-class data scheduler data
set class-of-service scheduler-maps fabric forwarding-class data-pri scheduler data-pri
set class-of-service scheduler-maps fabric forwarding-class rt scheduler rt
set class-of-service scheduler-maps fabric forwarding-class control scheduler control

set class-of-service interfaces xe-0/0/* scheduler-map fabric
set class-of-service interfaces ae0 scheduler-map fabric
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

On every vQFX:

```
set class-of-service schedulers data transmit-rate percent 30
set class-of-service schedulers data buffer-size percent 30
set class-of-service schedulers data priority low

set class-of-service schedulers data-pri transmit-rate percent 30
set class-of-service schedulers data-pri buffer-size percent 30
set class-of-service schedulers data-pri priority low

set class-of-service schedulers rt shaping-rate percent 10
set class-of-service schedulers rt buffer-size percent 5
set class-of-service schedulers rt priority strict-high

set class-of-service schedulers control transmit-rate percent 5
set class-of-service schedulers control buffer-size percent 5
set class-of-service schedulers control priority low

set class-of-service schedulers be transmit-rate percent 25
set class-of-service schedulers be buffer-size remainder
set class-of-service schedulers be priority low

set class-of-service scheduler-maps fabric forwarding-class be scheduler be
set class-of-service scheduler-maps fabric forwarding-class data scheduler data
set class-of-service scheduler-maps fabric forwarding-class data-pri scheduler data-pri
set class-of-service scheduler-maps fabric forwarding-class rt scheduler rt
set class-of-service scheduler-maps fabric forwarding-class control scheduler control

set class-of-service interfaces xe-0/0/* scheduler-map fabric
```

Verification

Verifying the scheduler-map configuration:

```
jncie@vMX3> show class-of-service scheduler-map fabric

Scheduler map: fabric, Index: 3939

  Scheduler: be, Forwarding class: be, Index: 3109
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: data, Forwarding class: data, Index: 35538
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: data-pri, Forwarding class: data-pri, Index: 58134
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: rt, Forwarding class: rt, Index: 3636
    Transmit rate: 10 percent, Rate Limit: rate-limit, Buffer size: 10 percent, Buffer Limit:
none, Priority: strict-high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: control, Forwarding class: control, Index: 56655
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium low      any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>
```

To check what scheduler-map is applied to an interface, issue the following command:

```
jncie@vMX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 142
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Logical interface: xe-0/0/2.0, Index: 335
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

jncie@vQFX3> show class-of-service scheduler-map fabric

Scheduler map: fabric, Index: 3939

  Scheduler: be, Forwarding class: be, Index: 3109
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: rt, Forwarding class: rt, Index: 3636
    Transmit rate: unspecified, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 10 percent
    Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: control, Forwarding class: control, Index: 56655
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

  Scheduler: data, Forwarding class: data, Index: 35538
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  Scheduler: data-pri, Forwarding class: data-pri, Index: 58134
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any           1    <default-drop-profile>
      Medium high     any           1    <default-drop-profile>
      High            any           1    <default-drop-profile>

jncie@vQFX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 650
Maximum usable queues: 8, Queues in use: 6
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled

  Logical interface: xe-0/0/2.0, Index: 558
Object                  Name                    Type                    Index
Classifier              dscp-default            dscp                        7
```

Task 6.3: WRED

- Configure an interpolated aggressive drop profile. When the buffer is at 50% utilization, the drop probability should be 75%. When the buffer utilization is 75%, the drop probability should be 95%.
- Configure an interpolated non-aggressive drop profile. When the buffer is at 50% utilization, the drop probability should be 30%. When the buffer utilization is 75%, the drop probability should be 50%.
- Apply the aggressive drop profile to the ‘data’ scheduler and apply he non-aggressive profile to the ‘data-pri’ scheduler.

Solution

On every device:

Configuring, and applying, the aggressive drop profile:

```
set class-of-service drop-profiles aggressive interpolate fill-level [50 75]
set class-of-service drop-profiles aggressive interpolate drop-probability [75 95]

set class-of-service schedulers be drop-profile-map loss-priority low protocol any drop-
profile aggressive
```

Configuring, and applying, the non-aggressive drop profile:

```
set class-of-service drop-profiles non-aggressive interpolate fill-level [50 75]
set class-of-service drop-profiles non-aggressive interpolate drop-probability [30 50]
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service schedulers data drop-profile-map loss-priority low protocol any drop-
profile non-aggressive
```

Verification

Verify the configuration parameters for the drop-profiles:

```
jncie@vQFX3> show class-of-service drop-profile non-aggressive

Drop profile: non-aggressive, Type: interpolated, Index: 58603
  Fill level     Drop probability
          50                   30
          75                   50

{master:0}
jncie@vQFX3> show class-of-service drop-profile aggressive

Drop profile: aggressive, Type: interpolated, Index: 9093
  Fill level     Drop probability
          50                   75
          75                   95
```

Check the ‘be’ and ‘data’ class to see if the drop profile was applied correctly:

```
jncie@vMX3> show class-of-service scheduler-map fabric

Scheduler map: fabric, Index: 3939

  Scheduler: be, Forwarding class: be, Index: 3109
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any          9093    aggressive
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: data, Forwarding class: data, Index: 35538
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         58603    non-aggressive
      Medium low      any             1    <default-drop-profile>
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Medium high     any              1    <default-drop-profile>
      High            any              1    <default-drop-profile>

  Scheduler: data-pri, Forwarding class: data-pri, Index: 58134
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: rt, Forwarding class: rt, Index: 3636
    Transmit rate: 10 percent, Rate Limit: rate-limit, Buffer size: 10 percent, Buffer Limit:
none, Priority: strict-high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: control, Forwarding class: control, Index: 56655
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
```

```
jncie@vQFX3> show class-of-service scheduler-map fabric

Scheduler map: fabric, Index: 3939

  Scheduler: be, Forwarding class: be, Index: 3109
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any          9093    aggressive
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: rt, Forwarding class: rt, Index: 3636
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
    Transmit rate: unspecified, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 10 percent
    Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: control, Forwarding class: control, Index: 56655
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: data, Forwarding class: data, Index: 35538
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         58603    non-aggressive
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: data-pri, Forwarding class: data-pri, Index: 58134
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
```

Task 6.4: BA Classifiers

- Configure classifiers for IPv4.
- Apply the classifiers to all interfaces.
- Ensure that the following classes are mapped to the mentioned code-points:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Class name	Code-points
be	be
data	af21
data-pri	af31
rt	ef
control	af11

Solution

On both vMX3 and vMX4:

```
set class-of-service classifiers dscp dscp-classifier forwarding-class rt loss-priority low
code-points ef
set class-of-service classifiers dscp dscp-classifier forwarding-class data-pri loss-priority
low code-points af31
set class-of-service classifiers dscp dscp-classifier forwarding-class data loss-priority low
code-points af21
set class-of-service classifiers dscp dscp-classifier forwarding-class control loss-priority
low code-points af11
set class-of-service classifiers dscp dscp-classifier forwarding-class be loss-priority low
code-points be

set class-of-service interfaces xe-0/0/* unit * classifiers dscp dscp-classifier

set class-of-service interfaces ae0 unit * classifiers dscp dscp-classifier
```

On every vQFX:

```
set class-of-service classifiers dscp dscp-classifier forwarding-class rt loss-priority low
code-points ef
set class-of-service classifiers dscp dscp-classifier forwarding-class data-pri loss-priority
low code-points af31
set class-of-service classifiers dscp dscp-classifier forwarding-class data loss-priority low
code-points af21
set class-of-service classifiers dscp dscp-classifier forwarding-class control loss-priority
low code-points af11
set class-of-service classifiers dscp dscp-classifier forwarding-class be loss-priority low
code-points be

set class-of-service interfaces xe-0/0/* unit * classifiers dscp dscp-classifier
```

Verification

```
jncie@vMX3> show class-of-service classifier name dscp-classifier

Classifier: dscp-classifier, Code point type: dscp, Index: 51090
  Code point         Forwarding class                    Loss priority
  000000             be                                  low
  001010             control                             low
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  010010              data                                    low
  011010              data-pri                                low
  101110              rt                                      low

jncie@vMX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 142
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled

  Logical interface: xe-0/0/2.0, Index: 335
Object                  Name                    Type                    Index
Classifier              dscp-classifier         dscp                    51090

jncie@vQFX3> show class-of-service classifier name dscp-classifier

Classifier: dscp-classifier, Code point type: dscp, Index: 51090
  Code point          Forwarding class                    Loss priority
  000000              be                                  low
  001010              control                             low
  010010              data                                low
  011010              data-pri                            low
  101110              rt                                  low

{master:0}
jncie@vQFX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 650
Maximum usable queues: 8, Queues in use: 6
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled

  Logical interface: xe-0/0/2.0, Index: 558
Object                  Name                    Type                    Index
Classifier              dscp-classifier         dscp                    51090
```

Task 6.5: Rewrite Rules

- Make sure that when packets exit devices, their IPv4 headers have markings that correspond to the queues that serviced the packets.

Solution

On both vMX3 and vMX4:

```
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class be loss-priority low
code-point be
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data loss-priority low
code-point af21
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data-pri loss-priority
low code-point af31
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class rt loss-priority low
code-point ef
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class control loss-priority
low code-point af11

set class-of-service interfaces xe-0/0/* unit * rewrite-rules dscp dscp-rewriter
set class-of-service interfaces ae0 unit * rewrite-rules dscp dscp-rewriter
```

On every vQFX:

```
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class be loss-priority low
code-point be
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data loss-priority low
code-point af21
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data-pri loss-priority
low code-point af31
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class rt loss-priority low
code-point ef
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class control loss-priority
low code-point af11

set class-of-service interfaces xe-0/0/* unit * rewrite-rules dscp dscp-rewriter
```

Verification

```
jncie@vQFX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 650
Maximum usable queues: 8, Queues in use: 6
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled

  Logical interface: xe-0/0/2.0, Index: 558
Object                  Name                    Type                    Index
Rewrite-Output          dscp-rewriter           dscp                    20901
Classifier              dscp-classifier         dscp                    51090

jncie@vMX3> show class-of-service interface xe-0/0/2

Physical interface: xe-0/0/2, Index: 142
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: fabric, Index: 3939
  Congestion-notification: Disabled

  Logical interface: xe-0/0/2.0, Index: 335
Object                  Name                    Type                    Index
Rewrite-Output          dscp-rewriter           dscp                    20901
Classifier              dscp-classifier         dscp                    51090
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 6.6: Traffic Classification on vQFX3 and vQFX4

- On vQFX3, rate limit inbound traffic on xe-0/0/1.10 to 2000Mbit and make sure all traffic is serviced by the be queue.
- On vQFX4 , inbound UDP traffic on xe-0/0/1.20 should be placed into the rt queue. Inbound TCP traffic should be placed into the data-pri queue and rate-limited to 4000Mbit.
- On vQFX4, rate limit traffic inside VLAN 100 to 4Gbit and place traffic into the data queue. When future VLANs are added, they should not be rate-limited and traffic from those VLANs should be put into the be queue.

Solution

vQFX3:

```
set firewall policer 2gbit if-exceeding bandwidth-limit 2g
set firewall policer 2gbit if-exceeding burst-size-limit 50m
set firewall policer 2gbit then discard

set firewall filter 2gbit-filter term 1 then policer 2gbit
set firewall filter 2gbit-filter term 1 then loss-priority low
set firewall filter 2gbit-filter term 1 then forwarding-class be
set firewall filter 2gbit-filter term 1 then accept

set interfaces xe-0/0/1 unit 10 family inet filter input 2gbit-filter
```

vQFX4:

```
set firewall policer 4gbit if-exceeding bandwidth-limit 4g
set firewall policer 4gbit if-exceeding burst-size-limit 50m
set firewall policer 4gbit then discard

set firewall filter mfc-xe-0/0/1.20 term udp from protocol udp
set firewall filter mfc-xe-0/0/1.20 term udp then loss-priority low
set firewall filter mfc-xe-0/0/1.20 term udp then forwarding-class rt
set firewall filter mfc-xe-0/0/1.20 term tcp from protocol tcp
set firewall filter mfc-xe-0/0/1.20 term tcp then policer 4gbit
set firewall filter mfc-xe-0/0/1.20 term tcp then loss-priority low
set firewall filter mfc-xe-0/0/1.20 term tcp then forwarding-class data-pri

set interfaces xe-0/0/1.20 family inet filter input mfc-xe-0/0/1.20

set firewall family ethernet-switching filter mfc-xe-0/0/3 term vlan-100 from user-vlan-id 100
set firewall family ethernet-switching filter mfc-xe-0/0/3 term vlan-100 then forwarding-class
data
set firewall family ethernet-switching filter mfc-xe-0/0/3 term vlan-100 then loss-priority
low
set firewall family ethernet-switching filter mfc-xe-0/0/3 term vlan-100 then policer 4gbit
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set firewall family ethernet-switching filter mfc-xe-0/0/3 term accept-rest then forwarding-
class be
set firewall family ethernet-switching filter mfc-xe-0/0/3 term accept-rest then loss-priority
low

set interfaces xe-0/0/3 unit 0 family ethernet-switching filter input mfc-xe-0/0/3
```

Verification

Verify that the firewall filters were correctly applied:

```
jncie@vQFX3> show interfaces xe-0/0/1.10 detail | match "Input fil"
      Input Filters: 2gbit-filter

jncie@vQFX4> show interfaces xe-0/0/1 detail | match "Input fil"
      Input Filters: mfc-xe-0/0/1.20

jncie@vQFX4> show interfaces xe-0/0/3 detail | match "Input fil"
      Input Filters: mfc-xe-0/0/3
```

QFX5100 Class of Service.

In this section, we provide you with a QFX5100 CoS configuration that is based on the previous assignments.

Drop-profile:

The drop-profile configuration and possibilities differ between the vQFX and the QFX5100. On the QFX5100, we have to configure the first drop-probability as being a 0 and we are allowed to configure only 2 fill levels:

```
set class-of-service drop-profiles aggressive interpolate fill-level 50
set class-of-service drop-profiles aggressive interpolate fill-level 75
set class-of-service drop-profiles aggressive interpolate drop-probability 0
set class-of-service drop-profiles aggressive interpolate drop-probability 95

set class-of-service drop-profiles non-aggressive interpolate fill-level 50
set class-of-service drop-profiles non-aggressive interpolate fill-level 75
set class-of-service drop-profiles non-aggressive interpolate drop-probability 0
set class-of-service drop-profiles non-aggressive interpolate drop-probability 50
```

Forwarding-class-sets:

The forwarding-class configuration on the QFX5100 is the same:

```
set class-of-service forwarding-classes class be queue-num 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service forwarding-classes class data queue-num 1
set class-of-service forwarding-classes class data-pri queue-num 2
set class-of-service forwarding-classes class rt queue-num 3
set class-of-service forwarding-classes class control queue-num 4
```

However, in addition to forwarding classes, we also need to configure forwarding-class-sets.

```
set class-of-service forwarding-class-sets fabric class be
set class-of-service forwarding-class-sets fabric class data
set class-of-service forwarding-class-sets fabric class data-pri
set class-of-service forwarding-class-sets fabric class control
```

Because the 'rt' class contains a strict-high scheduler, we need to configure a separate forwarding-class-set for that class:

```
set class-of-service forwarding-class-sets fabric-rt class rt
```

Scheduler-map and traffic-control-profile:

Because we cannot mix strict-high and other priority values inside a scheduler-map, we also configure two different scheduler-maps:

```
set class-of-service scheduler-maps fabric forwarding-class be scheduler be
set class-of-service scheduler-maps fabric forwarding-class data scheduler data
set class-of-service scheduler-maps fabric forwarding-class data-pri scheduler data-pri
set class-of-service scheduler-maps fabric forwarding-class control scheduler control

set class-of-service scheduler-maps fabric-rt forwarding-class rt scheduler rt
```

We do not apply the scheduler-maps directly to the interface. We apply the scheduler-maps to the interfaces using traffic-control-profiles. Before we get to that part, we first need to first configure these profiles:

```
set class-of-service traffic-control-profiles fabric scheduler-map fabric
set class-of-service traffic-control-profiles fabric shaping-rate percent 90

set class-of-service traffic-control-profiles fabric-rt scheduler-map fabric-rt
set class-of-service traffic-control-profiles fabric-rt shaping-rate percent 10
```

Applying forwarding-class-sets and traffic-control-profiles:

After configuring the traffic-control-profiles, we can apply them to the interfaces of the QFX5100:

```
set class-of-service interfaces xe-0/0/* forwarding-class-set fabric output-traffic-control-
profile fabric
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service interfaces xe-0/0/* forwarding-class-set fabric-rt output-traffic-
control-profile fabric-rt
```

The final configuration for this chapter on a physical QFX5100 could have been something like this:

```
set class-of-service classifiers dscp dscp-classifier forwarding-class rt loss-priority low
code-points ef
set class-of-service classifiers dscp dscp-classifier forwarding-class data-pri loss-priority
low code-points af31
set class-of-service classifiers dscp dscp-classifier forwarding-class data loss-priority low
code-points af21
set class-of-service classifiers dscp dscp-classifier forwarding-class control loss-priority
low code-points af11
set class-of-service classifiers dscp dscp-classifier forwarding-class be loss-priority low
code-points be

set class-of-service drop-profiles aggressive interpolate fill-level 50
set class-of-service drop-profiles aggressive interpolate fill-level 75
set class-of-service drop-profiles aggressive interpolate drop-probability 0
set class-of-service drop-profiles aggressive interpolate drop-probability 95

set class-of-service drop-profiles non-aggressive interpolate fill-level 50
set class-of-service drop-profiles non-aggressive interpolate fill-level 75
set class-of-service drop-profiles non-aggressive interpolate drop-probability 0
set class-of-service drop-profiles non-aggressive interpolate drop-probability 50

set class-of-service forwarding-classes class be queue-num 0
set class-of-service forwarding-classes class data queue-num 1
set class-of-service forwarding-classes class data-pri queue-num 2
set class-of-service forwarding-classes class rt queue-num 3
set class-of-service forwarding-classes class control queue-num 4

set class-of-service traffic-control-profiles fabric scheduler-map fabric
set class-of-service traffic-control-profiles fabric shaping-rate percent 90
set class-of-service traffic-control-profiles fabric-rt scheduler-map fabric-rt
set class-of-service traffic-control-profiles fabric-rt shaping-rate percent 10

set class-of-service forwarding-class-sets fabric-rt class rt
set class-of-service forwarding-class-sets fabric class be
set class-of-service forwarding-class-sets fabric class data
set class-of-service forwarding-class-sets fabric class data-pri
set class-of-service forwarding-class-sets fabric class control

set class-of-service interfaces xe-0/0/* forwarding-class-set fabric-rt output-traffic-
control-profile fabric-rt
set class-of-service interfaces xe-0/0/* forwarding-class-set fabric output-traffic-control-
profile fabric

set class-of-service interfaces xe-0/0/* unit * classifiers dscp dscp-classifier
set class-of-service interfaces xe-0/0/* unit * rewrite-rules dscp dscp-rewriter
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service interfaces ae0 unit * classifiers dscp dscp-classifier
set class-of-service interfaces ae0 unit * rewrite-rules dscp dscp-rewriter
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class be loss-priority low
code-point be
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data loss-priority low
code-point af21
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class data-pri loss-priority
low code-point af31
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class rt loss-priority low
code-point ef
set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class control loss-priority
low code-point af11

set class-of-service scheduler-maps fabric forwarding-class be scheduler be
set class-of-service scheduler-maps fabric forwarding-class data scheduler data
set class-of-service scheduler-maps fabric forwarding-class data-pri scheduler data-pri
set class-of-service scheduler-maps fabric forwarding-class control scheduler control
set class-of-service scheduler-maps fabric-rt forwarding-class rt scheduler rt

set class-of-service schedulers be transmit-rate percent 25
set class-of-service schedulers be buffer-size remainder
set class-of-service schedulers be priority low
set class-of-service schedulers be drop-profile-map loss-priority low protocol any drop-
profile aggressive

set class-of-service schedulers data transmit-rate percent 30
set class-of-service schedulers data buffer-size percent 30
set class-of-service schedulers data priority low
set class-of-service schedulers data drop-profile-map loss-priority low protocol any drop-
profile non-aggressive

set class-of-service schedulers data-pri transmit-rate percent 30
set class-of-service schedulers data-pri buffer-size percent 30
set class-of-service schedulers data-pri priority low

set class-of-service schedulers rt buffer-size percent 10
set class-of-service schedulers rt priority strict-high

set class-of-service schedulers control transmit-rate percent 5
set class-of-service schedulers control buffer-size percent 5
set class-of-service schedulers control priority low
```

After applying the configuration to a QFX5100, the verification commands would have returned the following output:

```
lab@jncie-qfx5100> show class-of-service forwarding-class

Forwarding class                  ID      Queue  Policing priority  No-Loss
  be                              0         0          normal        Disabled
  rt                              1         3          normal        Disabled
  control                         2         4          normal        Disabled
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
  network-control                 3          7          normal       Disabled
  data                            4          1          normal       Disabled
  data-pri                        5          2          normal       Disabled
  mcast                           8          8          normal       Disabled

lab@jncie-qfx5100> show class-of-service forwarding-class-set

Forwarding class set: fabric, Type: normal-type, Forwarding class set index: 3939
  Forwarding class                          Index
  be                                          0
  data                                        4
  data-pri                                    5
  control                                     2

Forwarding class set: fabric-rt, Type: normal-type, Forwarding class set index: 15749
  Forwarding class                          Index
  rt                                          1

lab@jncie-qfx5100> show class-of-service drop-profile aggressive

Drop profile: aggressive, Type: interpolated, Index: 9093
  Fill level    Drop probability
          50                   0
          75                  95

lab@jncie-qfx5100> show class-of-service drop-profile non-aggressive

Drop profile: non-aggressive, Type: interpolated, Index: 58603
  Fill level    Drop probability
          50                   0
          75                  50

lab@jncie-qfx5100> show class-of-service scheduler-map fabric

Scheduler map: fabric, Index: 3939

  Scheduler: be, Forwarding class: be, Index: 3109
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any          9093    aggressive
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: control, Forwarding class: control, Index: 56655
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: data, Forwarding class: data, Index: 35538
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         58603    non-aggressive
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

  Scheduler: data-pri, Forwarding class: data-pri, Index: 58134
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 30 percent, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

lab@jncie-qfx5100> show class-of-service scheduler-map fabric-rt

Scheduler map: fabric-rt, Index: 15749

  Scheduler: rt, Forwarding class: rt, Index: 3636
    Transmit rate: unspecified, Rate Limit: none, Buffer size: 10 percent, Buffer Limit: none,
Priority: strict-high
    Excess Priority: unspecified, Explicit Congestion Notification: disable
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>

lab@jncie-qfx5100> show class-of-service traffic-control-profile fabric

Traffic control profile: fabric, Index: 3939
  Shaping rate: 90 percent
  Scheduler map: fabric

lab@jncie-qfx5100> show class-of-service traffic-control-profile fabric-rt

Traffic control profile: fabric-rt, Index: 15749
  Shaping rate: 10 percent
  Scheduler map: fabric-rt
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
lab@jncie-qfx5100> show class-of-service interface xe-0/0/0

Physical interface: xe-0/0/0, Index: 647
Queues supported: 12, Queues in use: 7
  Congestion-notification: Disabled
  Forwarding class set: fabric-rt, Index: 15749
    Output traffic control profile: fabric-rt, Index: 15749
  Forwarding class set: fabric, Index: 3939
    Output traffic control profile: fabric, Index: 3939
  Logical interface: xe-0/0/0.16386, Index: 555
```

Appendix – Chapter 7: Management

Part 1: Data Center Management

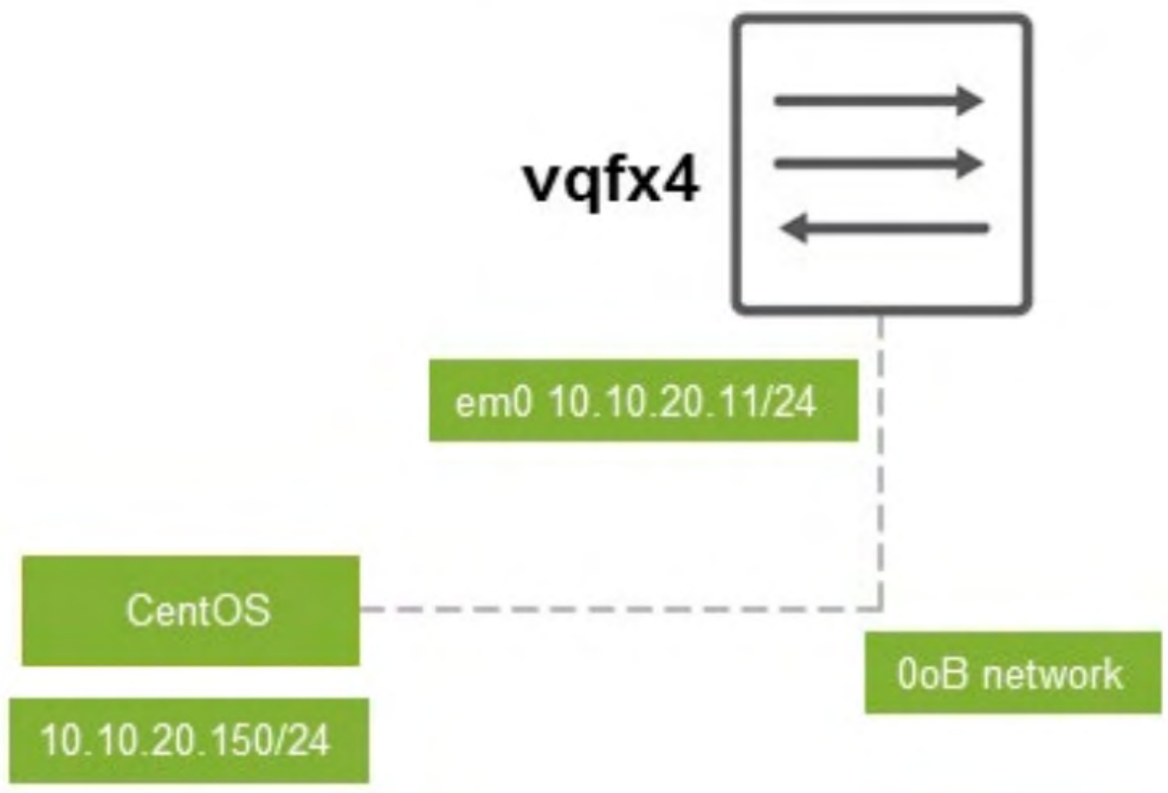

Figure 1 – ZTP and NETCONF.

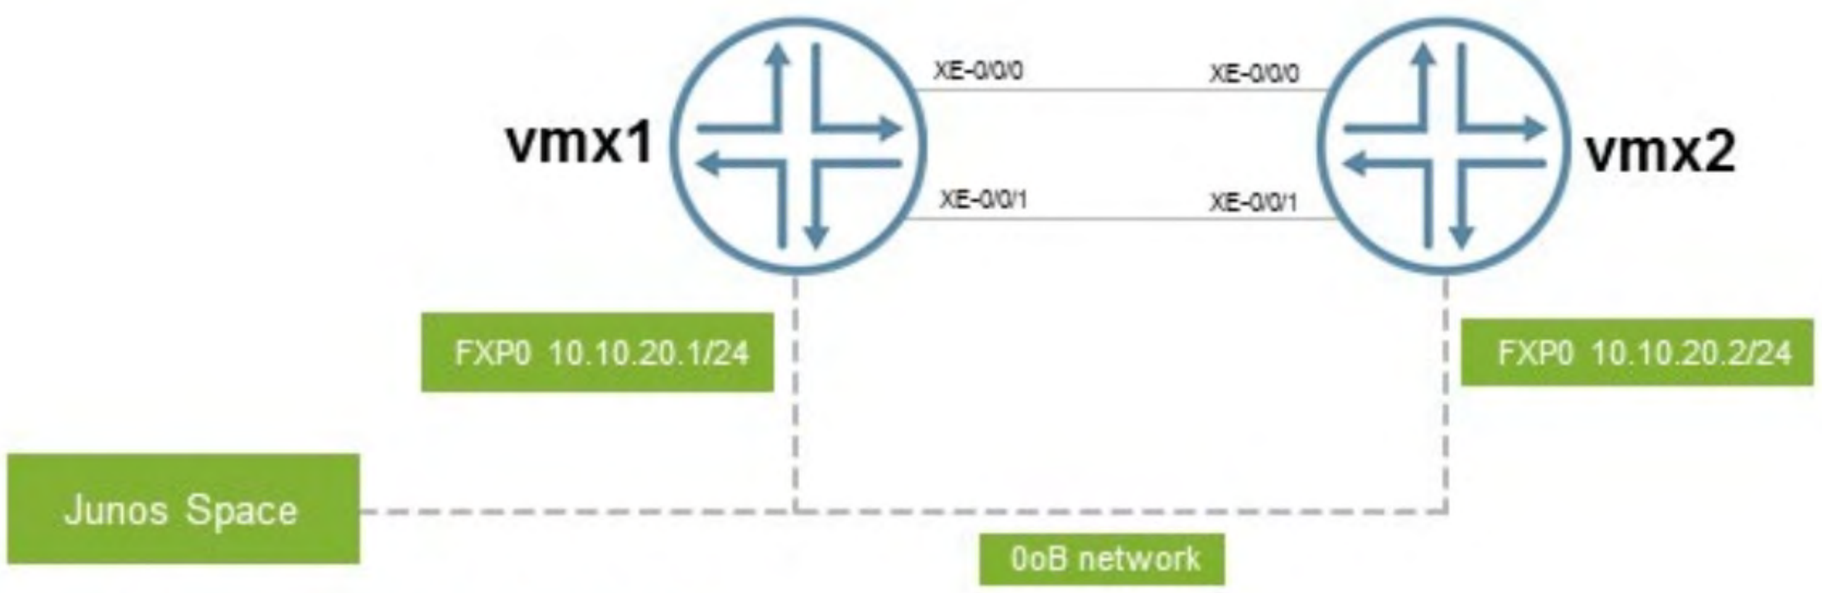

Figure 2 – Junos space and on box scripting.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 7.1: Zero Touch Provisioning (ZTP)

- The DHCP server should assign the 10.10.20.11 IP address based on the MAC address of vQFX4 em0 interface.

Note: When working with QFX hardware devices you should be using the MAC address of the chassis (incremented by one) for the DHCP configuration file definition.

```
{master:0}
jncie@vQFX4> show interfaces em0 | match current
  Current address: 00:50:56:a9:5b:2d, Hardware address: 00:50:56:a9:5b:2d
```

Edit the file /etc/dhcp/dhcpd.conf using the "vi" editor and add the following information using the jncie/jncie123 username and password for the student desktop.

```
subnet 10.10.20.0 netmask 255.255.255.0 {
option routers 10.10.20.254 ;
option subnet-mask 255.255.255.0 ;
default-lease-time 28800 ;
max-lease-time 32200 ;
range 10.10.20.50 10.10.20.55 ;
}
option option-150 code 150 = ip-address ;
option space ZTP ;
option ZTP.server-image code 4 = text ;
option ZTP.server-image code 0 = text ;
option ZTP.server-file  code 1 = text ;
option ZTP.image-file-type code 2 = text ;
option ZTP.transfer-mode code 3 = text ;
option ZTP-encapsulation code 43 = encapsulate ZTP ;

host vqfx4 {
hardware ethernet 00:50:56:a9:5b:2d ;
fixed-address 10.10.20.11 ;
option option-150 10.10.20.150 ;
}
```

Restart the DHCPD

```
[jncie@desktop /]$ sudo systemctl start dhcpd
```

- During ZTP, vQFX4 should automatically download the configuration file /var/ftp/pub/ztp/vqfx4 using FTP from the Centos host.

```
host vqfx4 {
hardware ethernet 00:50:56:a9:5b:2d ;
fixed-address 10.10.20.11 ;
option option-150 10.10.20.150 ;
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
option ZTP.server-file "/pub/ztp/vqfx4" ;
option ZTP.transfer-mode "ftp" ;
}
```

Restart the DHCPD

```
[jncie@desktop /]$ sudo systemctl start dhcpd
```

Warning: *Although the vqfx4 configuration file is located in /var/pub/ztp/vqfx4, the FTP server root directory is rooted at /var/ftp/. This means that the option ZTP.server-file path should be specified as "/pub/ztp/vqfx4" and not "/var/ftp/pub/ztp/vqfx4". You can test this by logging in to the FTP server using the anonymous username:*

```
#ftp 10.10.20.150
Connected to 10.10.20.150 (10.10.20.150).
220 (vsFTPd 3.0.2)
Name (10.10.20.150:root): anonymous
331 Please specify the password.
Password:
230 Login successful.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> ls
227 Entering Passive Mode (10,10,20,150,97,151).
150 Here comes the directory listing.
drwxr-xr-x    3 0        0              17 Mar 27 09:28 pub
226 Directory send OK.
ftp>
```

- Ensure that vQFX4 is able to perform ZTP using interface em0.

```
set chassis auto-image-upgrade
delete interfaces
set interfaces em0 unit 0 family inet dhcp vendor-id Juniper-qfx10002-72q
commit and-quit
```

If your correctly configured the DHCP configuration and enabled ZTP on vQFX4 you should see the following messages on the console output of vQFX4

```
Auto Image Upgrade: DHCP Client Bound interfaces: em0.0

Auto Image Upgrade: DHCP Client Unbound interfaces: xe-0/0/0.0 xe-0/0/1.0
xe-0/0/2.0 xe-0/0/3.0 xe-0/0/4.0 xe-0/0/5.0 xe-0/0/6.0 xe-0/0/7.0 xe-0/0/8.0
xe-0/0/9.0 xe-0/0/10.0 xe-0/0/11.0

Auto Image Upgrade: To stop, on CLI apply
"delete chassis auto-image-upgrade"  and commit

...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
...

Auto Image Upgrade: DHCP Options for client interface em0.0 ConfigFile:
/pub/ztp/vqfx4 Gateway: 10.10.20.254 DHCP Server: 10.10.20.150 File Server:
10.10.20.150 Options state:
Partial Options::Config File and File Server Set, Image File Not Set

Auto Image Upgrade: Active on client interface: em0.0

Auto Image Upgrade: Interface::    "em0"
Auto Image Upgrade: Server::       "10.10.20.150"
Auto Image Upgrade: Image File::   "NOT SPECIFIED"
Auto Image Upgrade: Server File:: "vqfx4"
Auto Image Upgrade: Gateway::      "10.10.20.254"
Auto Image Upgrade: Protocol::     "ftp"

Auto Image Upgrade: Start fetching vqfx4 file from server 10.10.20.150 through
em0 using ftp

Auto Image Upgrade: File vqfx4 fetched from server 10.10.20.150 through em0

Auto Image Upgrade: Applying vqfx4 file configuration fetched from server 10.10
.20.150 through em0
```

Task 7.2: Python using NETCONF

- On the CentOS host, create a Python script called “jncie.py” that retrieves the routing table of device vQFX4. You may adjust the following example script:

Example script:

```
from jnpr.junos import Device
from jnpr.junos.op.routes import RouteTable
from lxml import etree

username='root'
password='root1'
dev = Device(host='10.10.20.4', user=username, password=mypassword, gather_facts='false')
dev.open()

routes = RouteTable(dev).get()
print(routes)
dev.close()
```

Make the following changes:

```
Username='jncie'
password='jncie123'
dev = Device(host='10.10.20.11', user=username, password=password,
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Make sure that the Python script is able to execute using the following syntax:

To enable the jncie.py script to run by itself, you need to add a “shebang” statement as the first line of the script. Under Unix-like operating systems, in our case Centos, when a script with a shebang is run as a program, the program loader uses the first line of the script as the interpreter program (#!/usr/bin/python) and executes the rest of the script using this interpreter.

Your final script should look something like this:

```
#!/usr/bin/python
from jnpr.junos import Device
from jnpr.junos.op.routes import RouteTable
from lxml import etree

username='jncie'
password='jncie123'
dev = Device(host='10.10.20.11', user=username, password=password, gather_facts='false')
dev.open()

routes = RouteTable(dev).get()
print(routes)
dev.close()
quit()
```

Last but not least, make sure that the jncie.py script has executable rights

```
[jncie@desktop ~]$ chmod 750 jncie.py
```

Task 7.3: Junos Space

- Ensure that vMX1 and vMX2 are discovered using SNMP version 1/2c. Use SNMP community “jncie”
- Make sure that the configuration files of vMX1 and vMX2 are stored on the Junos Space server.

1. Configure the SNMP community on vMX1 and vMX2

```
root@vMX1# show snmp
community jncie;

root@vMX2# show snmp
community jncie;
```

2. In Junos Space, click Devices and then Device Discovery. Select “Discover Targets” and click on the plus, or 'add', sign.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

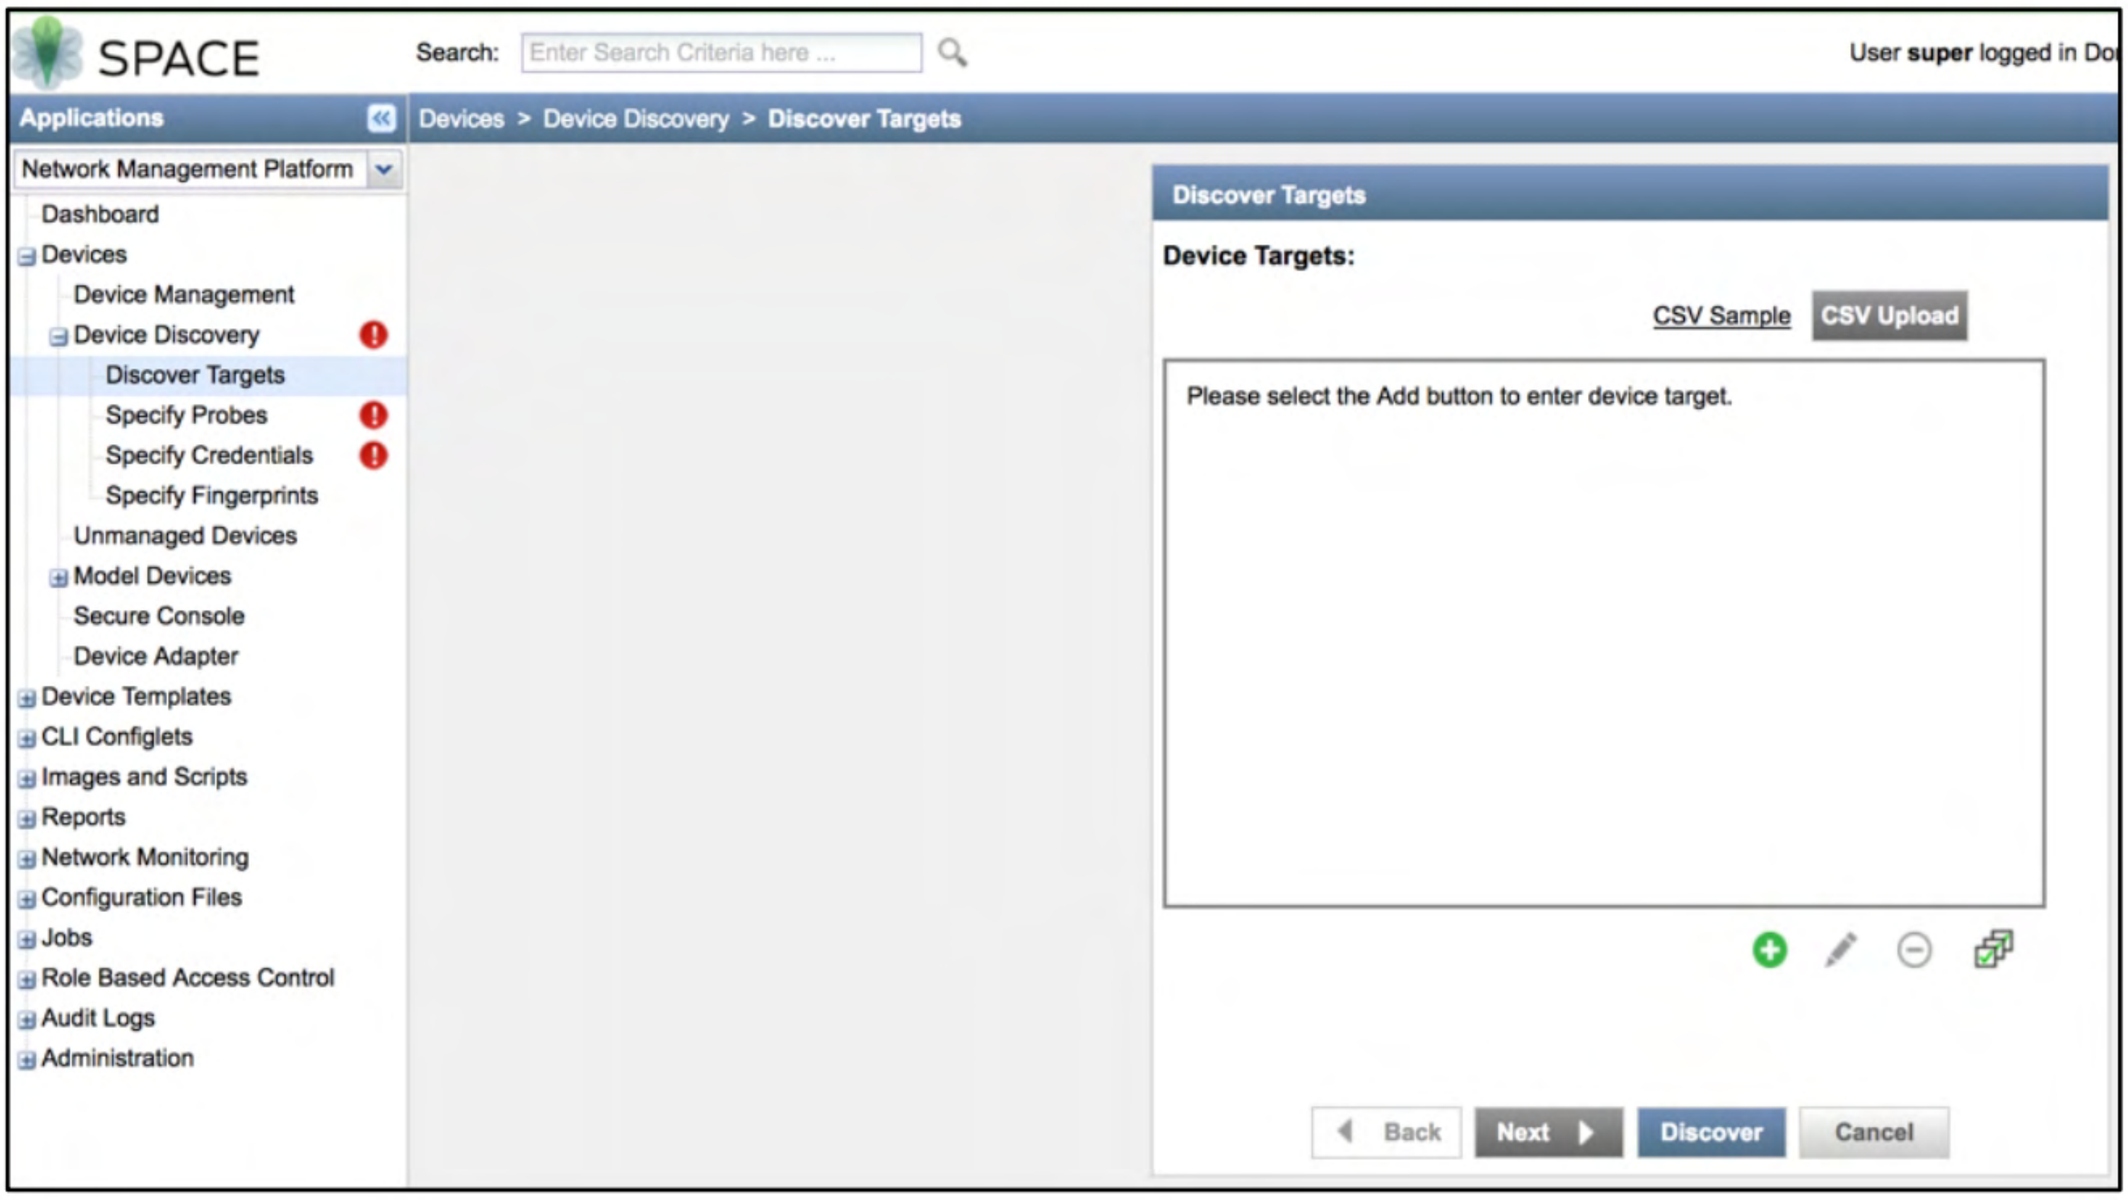

© 2020 Juniper Networks, Inc. All rights reserved.

3. Select IP range and enter 10.10.20.1 as the first IP and 10.10.20.2 as the last IP. Click “add”.

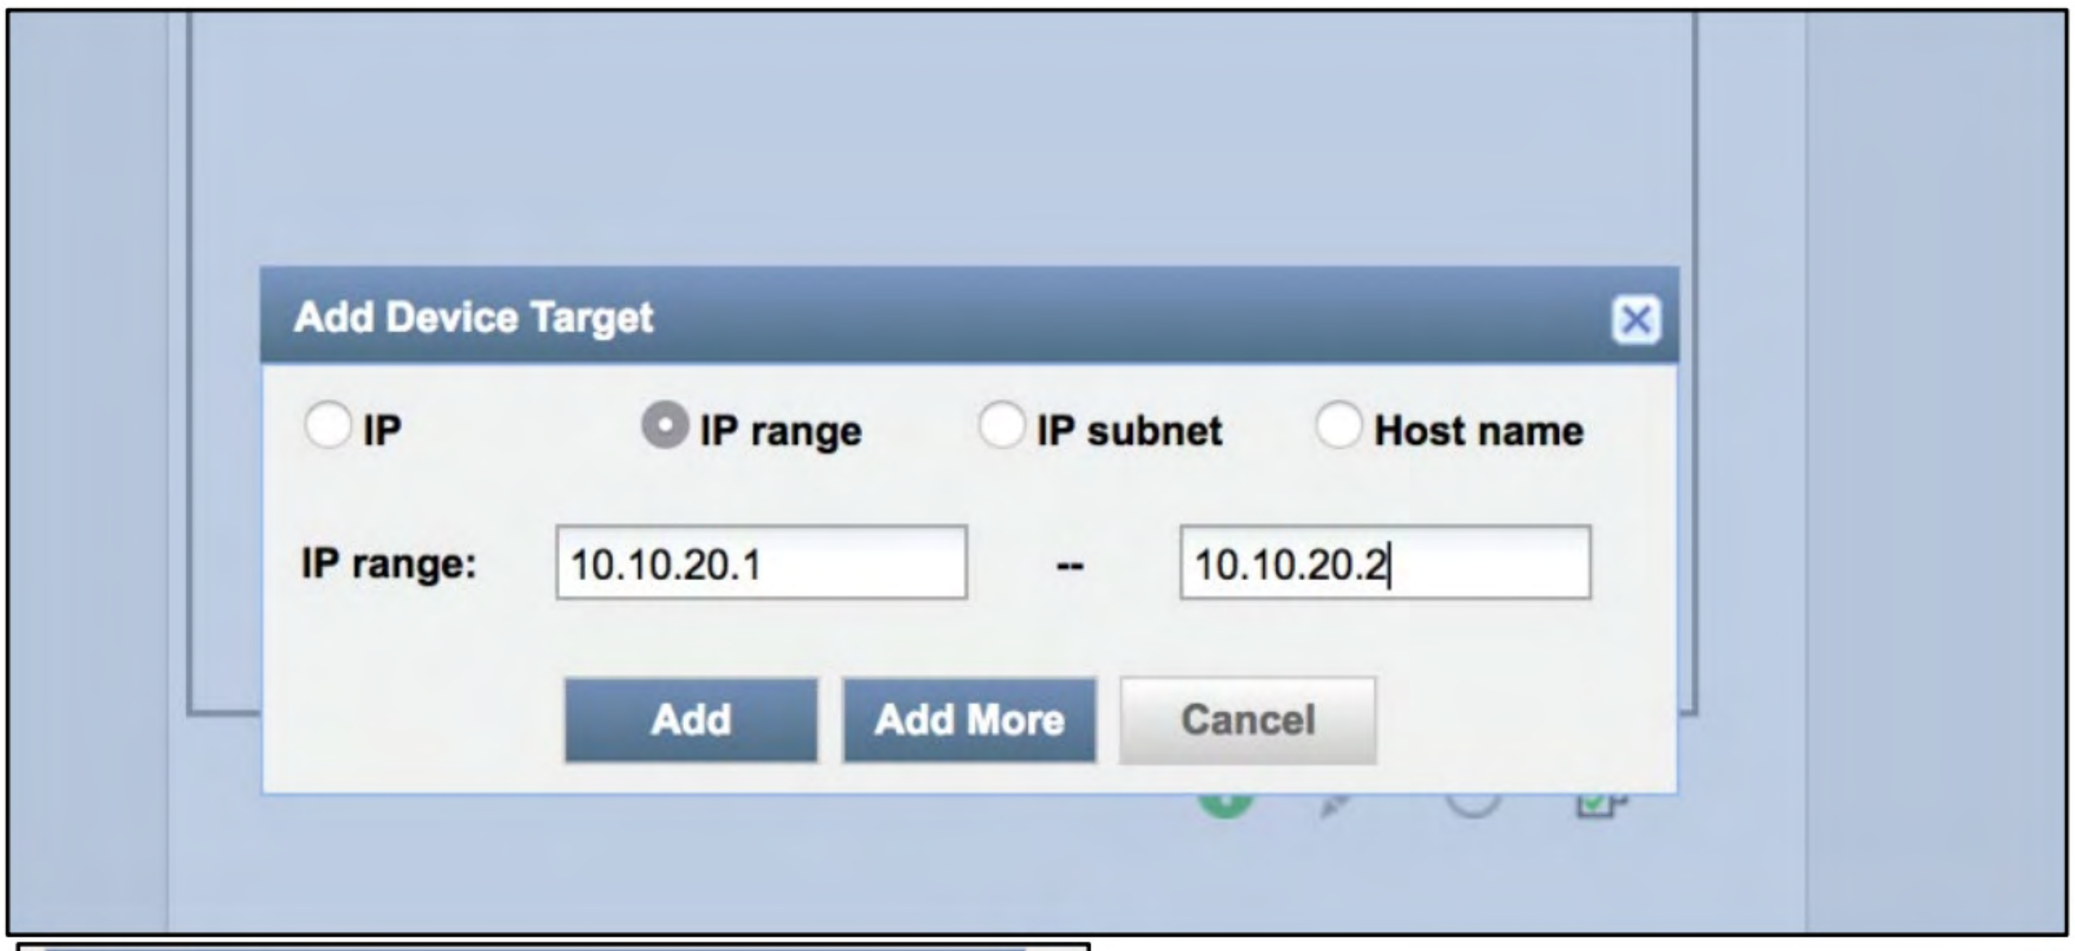

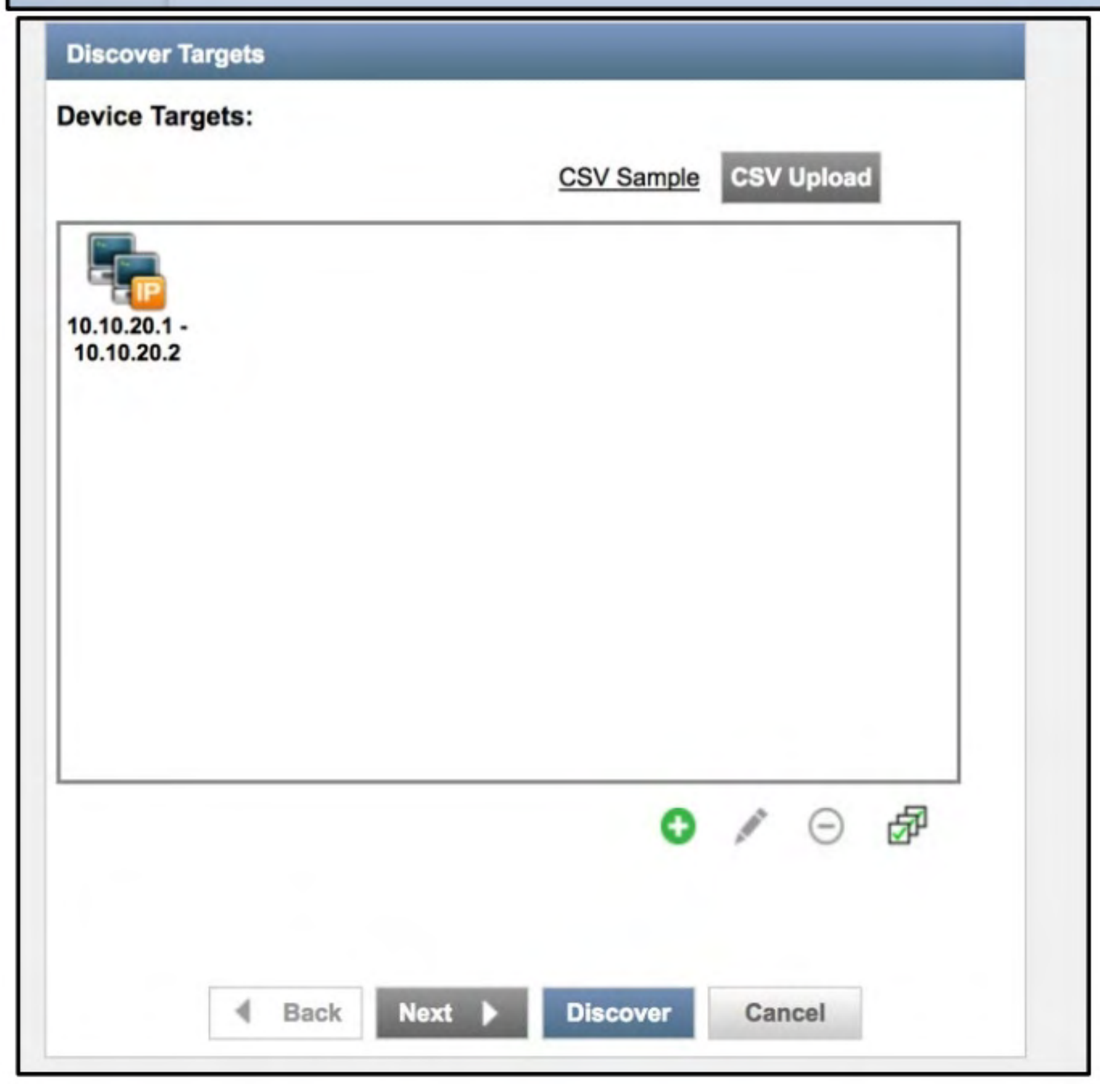

The devices 10.10.20.1 (vMX1) and 10.10.20.2 (vMX2) are both added as Discovery Targets.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

4. Uncheck the “Use Ping” option so only the “Use SNMP” option is selected. Highlight the public snmp community and click the edit “pencil” icon to modify the settings”.



5. Select "SNMP V1/V2C" and enter “jncie” as the community. Click “Modify”.



© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only



The configured SNMP settings should be visible now under "SNMP Settings".
Click “next”.

6. Configure username “jncie” and password “jncie123” to be used as Device Login Credentials. Click “add”.



© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

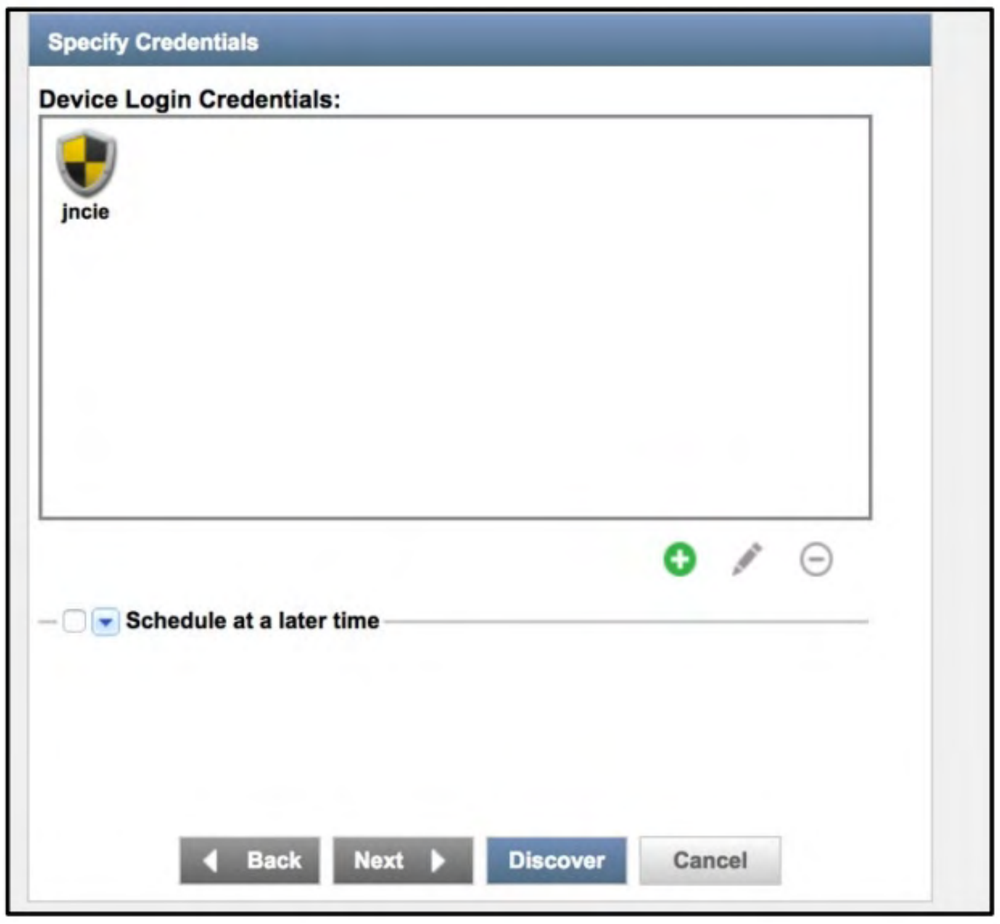

The Device login credentials should be displayed as shown above. Click “Discover”.

7. vMX1 and vMX2 should be discovered as shown below.

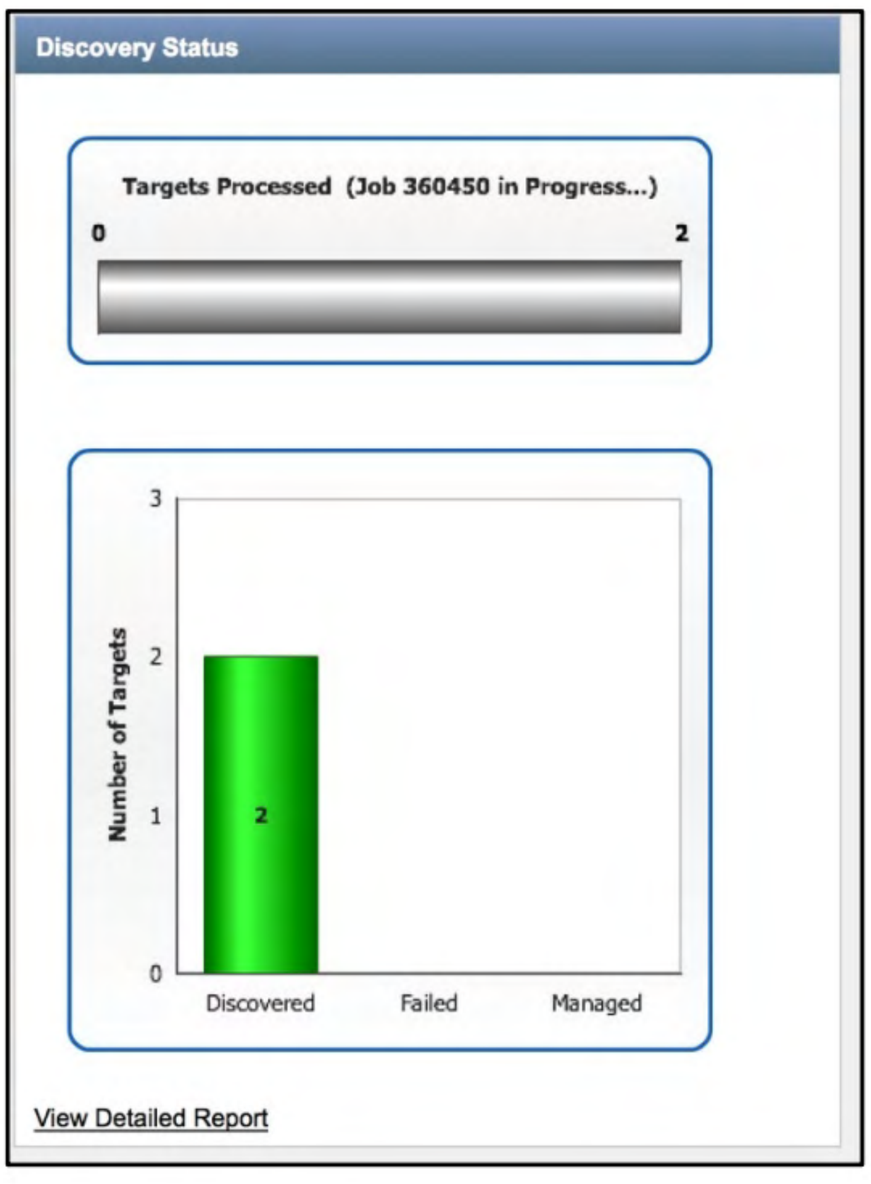

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only



If, for some reason, the devices are not discovered, make sure that vMX1 and vMX2 are able to reach the Junos Space server. Also make sure that the SNMP community “jncie” and username “jncie” are configured correctly on both vMX1 and vMX2.

8. In Junos Space, click on "Configuration Files". Select “click here to backup device configuration files”.

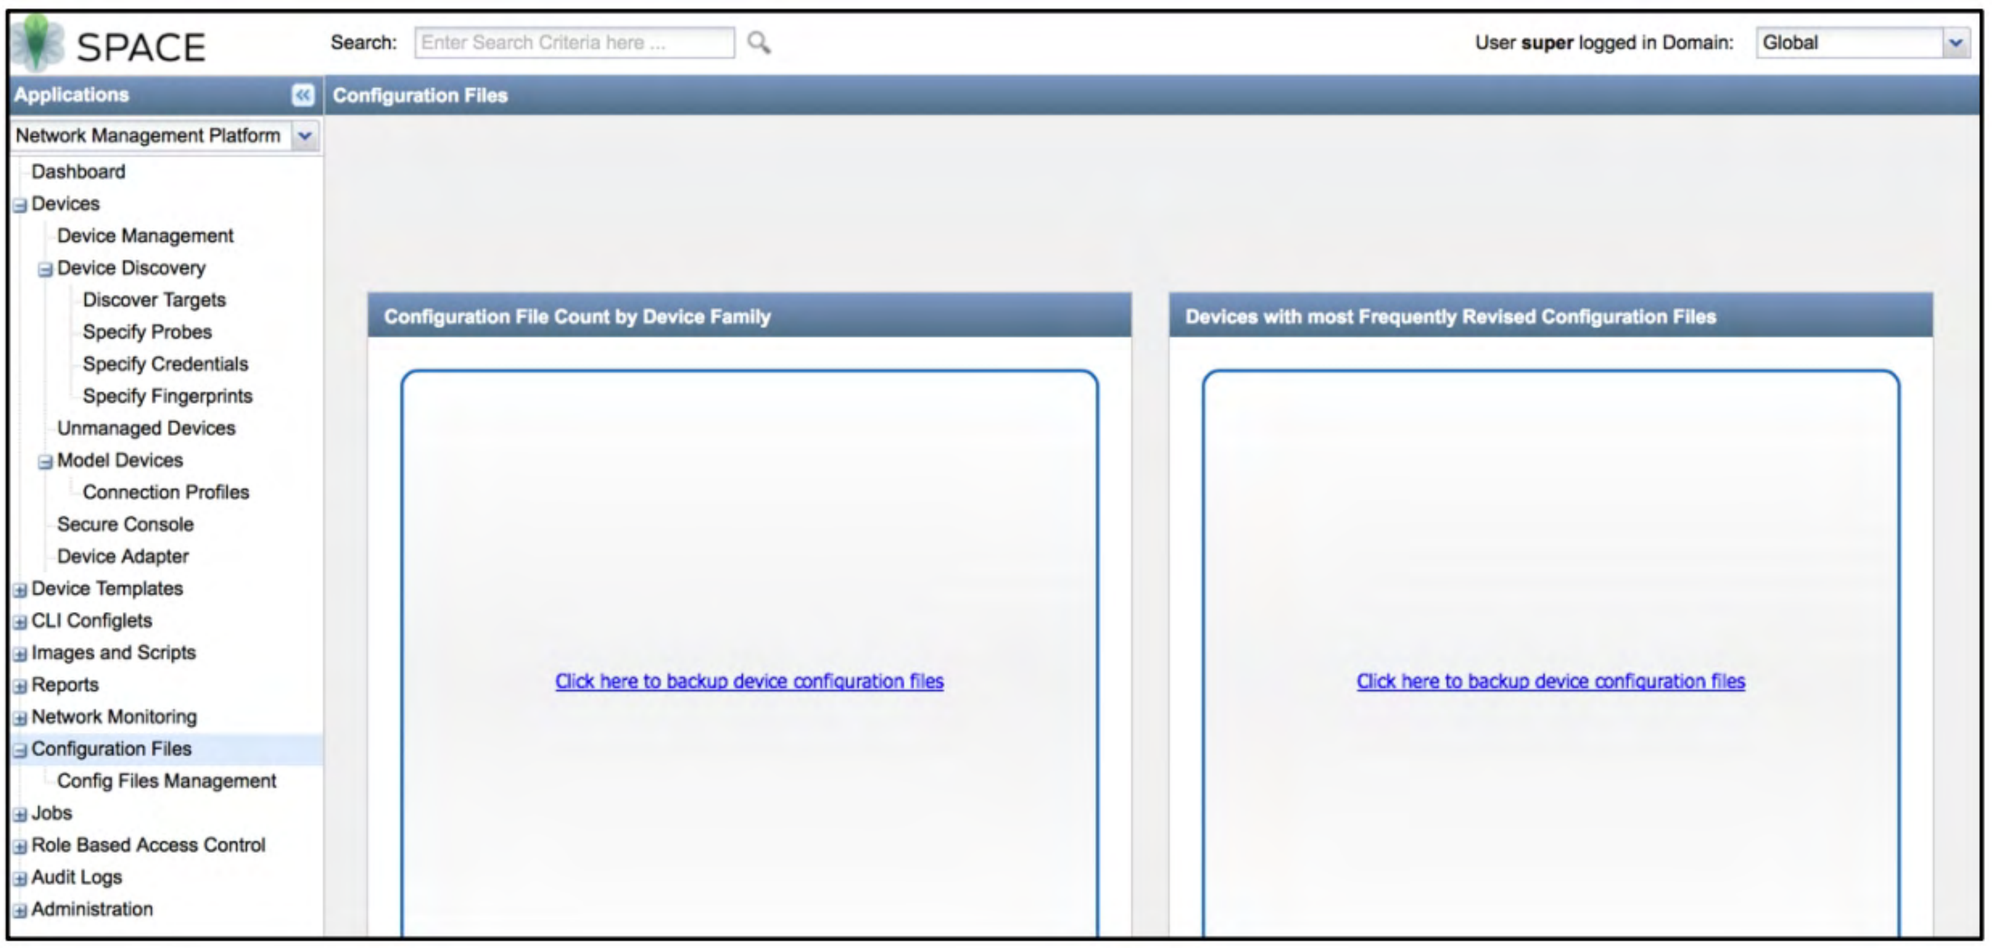

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

9. The Backup Configuration Files menu should display the configuration files of vMX1 and vMX2.

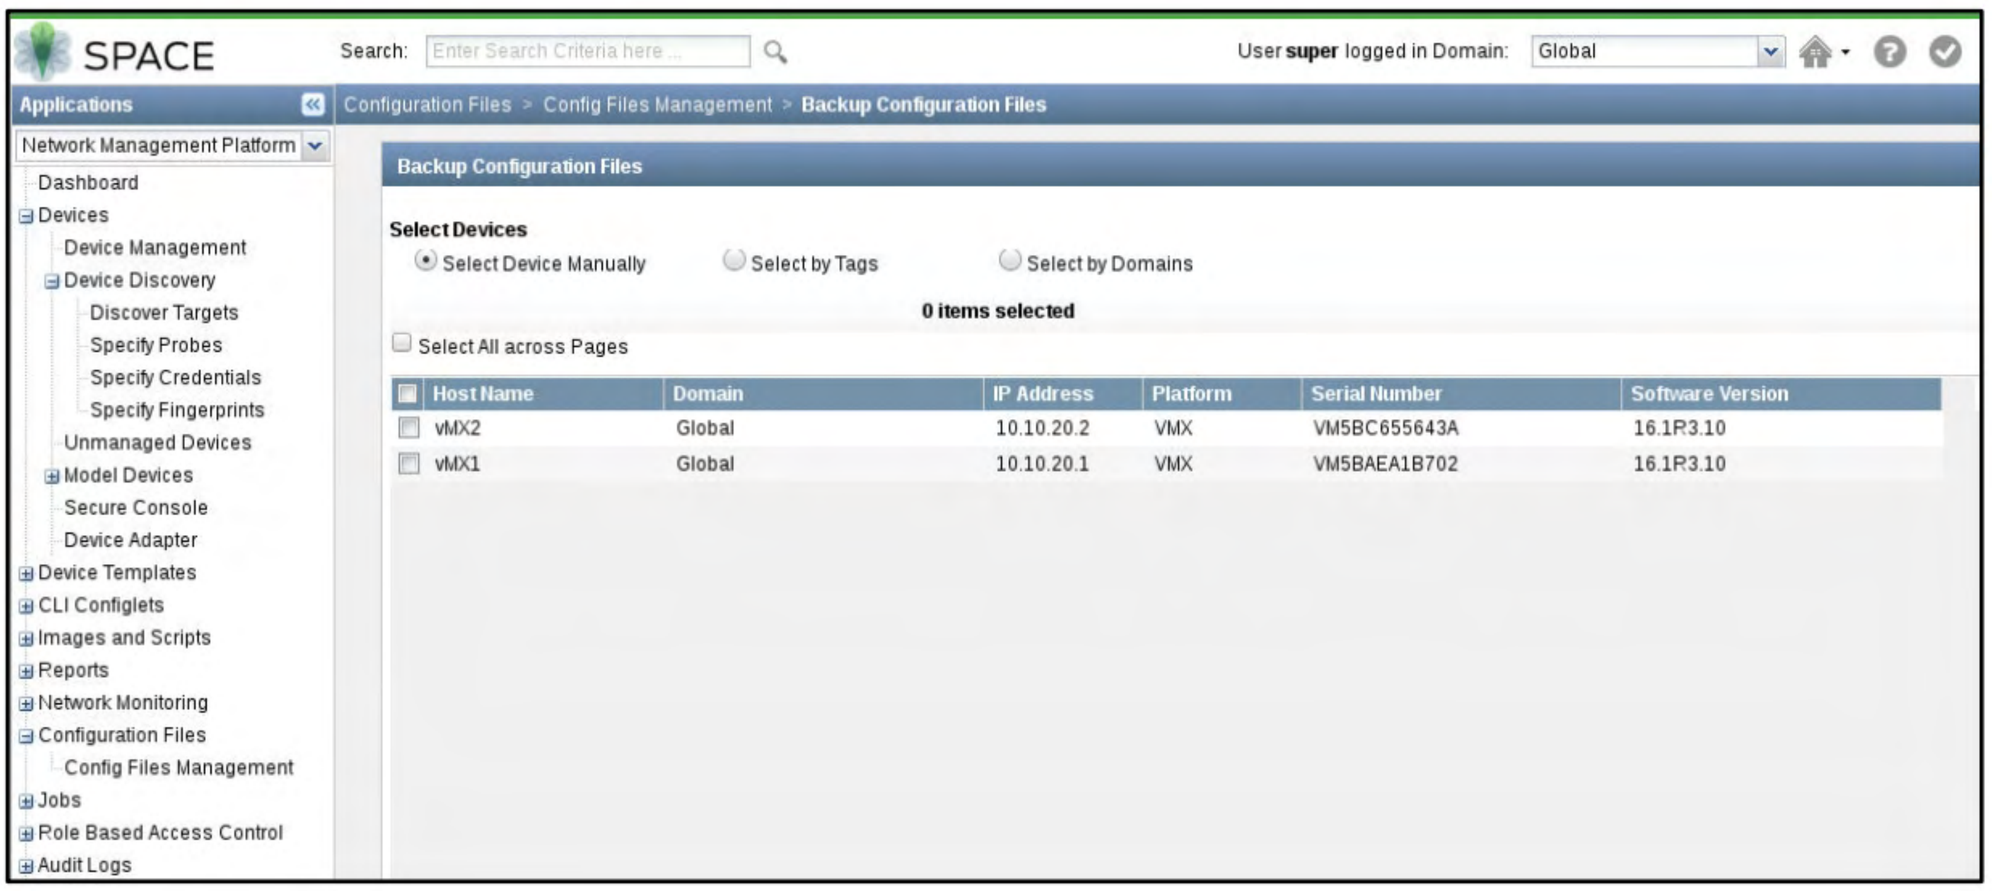

Task 7.4: On Box Scripting 1

- Configure an OSPF adjacency without a DR/BDR election between vMX1 and vMX2 on interface xe-0/0/0. Hello interval and dead interval should be set to the default values / 10. The vMX2 device should use IP address 192.168.12.2 on interface xe-0/0/0.0
- Ensure that interface xe-0/0/0 on vMX1 is disabled by the Junos OS when the OSPF adjacency on interface xe-0/0/0 fails two times in one minute.
- Ensure that, when the Junos OS disables link xe-0/0/0 in vMX1 due to OSPF failures, an XML file is created in the /tmp directory of vMX1. The file must list all active OSPF adjacencies.

Solution

vMX1:

```
set interfaces xe-0/0/0 unit 0 family inet address 192.168.12.1/24

set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 dead-interval 4

set event-options policy ospf events rpd_ospf_nbrdown
set event-options policy ospf within 60 trigger after
set event-options policy ospf within 60 trigger 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set event-options policy ospf within 60 events rpd_ospf_nbrup
set event-options policy ospf attributes-match rpd_ospf_nbrdown.neighbor-address equals
rpd_ospf_nbrup.neighbor-address
set event-options policy ospf attributes-match "{$$.neighbor-address}" matches 192.168.12.2
set event-options policy ospf then change-configuration commands "set interfaces xe-0/0/0
disable"
set event-options policy ospf then execute-commands commands "show ospf neighbor"
set event-options policy ospf then execute-commands output-filename jncielog
set event-options policy ospf then execute-commands destination jncielog
set event-options policy ospf then execute-commands output-format xml
set event-options destinations jncielog archive-sites /tmp
```

vMX2:

```
set interfaces xe-0/0/0 unit 0 family inet address 192.168.12.2/24

set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 hello-interval 1
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 dead-interval 4
```

Verification

Verify if the OSPF adjacency is currently established

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
192.168.12.2     xe-0/0/0.0             Full      10.10.20.2       128     3
```

Test the event-policy by clearing the OSPF adjacency

```
jncie@vMX1> clear ospf neighbor all
```

Verify if the OSPF adjacency is reestablished. This should be the case as the interface should only be disabled if the event occurs two times within one minute.

```
jncie@vMX1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
192.168.12.2     xe-0/0/0.0             Full      10.10.20.2       128     3
```

Clear the OSPF neighbor again. Make sure that you execute this command within the same minute to ensure that the event-policy is triggered.

```
[edit]
root@vMX1# run clear ospf neighbor
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verify if the OSPF neighbor is reestablished.

```
jncie@vMX1> show ospf neighbor
```

The OSPF adjacency is not established, great! Verify if the Junos OS indeed disabled interface xe-0/0/0 in the configuration.

```
jncie@vMX1> show configuration interfaces xe-0/0/0

disable;
unit 0 {
    family inet {
        address 192.168.12.1/24;
    }
}
```

Verify the creation of a file in the "/tmp" directory

```
jncie@vMX1> start shell
% pwd
/var/home/jncie
% ls -la /tmp/*jncie*
-rw-r-----  1 root  wheel  1011 Oct 29 16:13 /tmp/vMX1_20181029_161310_jncielog
-rw-r-----  1 root  wheel   717 Oct 29 16:13 /tmp/vMX1_20181029_161312_jncielog
% cat /tmp/vMX1_20181029_161312_jncielog
<?xml version="1.0" encoding="us-ascii"?>
<junoscript xmlns="http://xml.juniper.net/xnm/1.1/xnm"
xmlns:junos="http://xml.juniper.net/junos/16.1R3/junos"
schemaLocation="http://xml.juniper.net/junos/16.1R3/junos junos/16.1R3/junos.xsd" os="JUNOS"
release="16.1R3.10" hostname="vMX1" version="1.0">
<!-- session start at 2018-10-29 16:13:12 UTC -->
<!-- No zombies were killed during the creation of this user interface -->
<!-- user root, class super-user -->
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/16.1R3/junos">
<ospf-neighbor-information xmlns="http://xml.juniper.net/junos/16.1R3/junos-routing">
</ospf-neighbor-information>
</rpc-reply>
<!-- session end at 2018-10-29 16:13:12 UTC -->
</junoscript>
%
```

Task 7.5: On Box Scripting 2

- Configure an OSPF adjacency without DR/BDR election between vMX1 and vMX2 on interface xe-0/0/0 and xe-0/0/1. Hello interval and dead interval should be set to the default values / 10.
- Enable the LDP protocol on interface xe-0/0/0 and xe-0/0/1 for vMX1 and vMX2.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Ensure that an event-script called “jncie.slax” is executed on vMX2 when the LDP neighbor relationship on interface xe-0/0/0 or interface xe-0/0/1 goes down twice within one minute. The failed interface name must be passed to the jncie.slax script.
- Ensure that the jncie.slax script is paused for 1/2 second after each executed line when the CPU usage of vMX2 is 40% or more.

Note: *For this task you can use the predefined script named jncie.slax stored in the /var/db/scripts/events/ directory. The point for this task is to understand how to execute an event-script policy that calls a script.*

```
% pwd
/var/db/scripts/event
% cat jncie.slax
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "interface";
  }
}

param $interface;

match / {
  <event-script-results> {
    var $message = "The " _ $interface _ " interface is causing LDP to bounce more than two
times in the one minute";
    expr jcs:syslog("pfe.notice" , $message);
  }
}

%
```

Solution

vMX1:

```
delete interfaces xe-0/0/0 disable
delete protocols ospf

set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/1 unit 0 family inet address 192.168.22.1/24
set interfaces xe-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1.0 interface-type p2p

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/1.0

set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/1.0
set protocols ldp interface lo0.0
```

vMX2:

```
delete protocols ospf

set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/1 unit 0 family inet address 192.168.22.2/24
set interfaces xe-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 2.2.2.2/32

set protocols ospf area 0.0.0.0 interface xe-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface xe-0/0/1.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface lo0.0

set protocols mpls interface xe-0/0/0.0
set protocols mpls interface xe-0/0/1.0

set protocols ldp interface xe-0/0/0.0
set protocols ldp interface xe-0/0/1.0
set protocols ldp interface lo0.0

set event-options policy ldp events rpd_ldp_nbrdown
set event-options policy ldp within 60 trigger on
set event-options policy ldp within 60 trigger 2
set event-options policy ldp within 60 events rpd_ldp_nbrup
set event-options policy ldp attributes-match rpd_ldp_nbrdown.interface-name equals
rpd_ldp_nbrup.interface-name
set event-options policy ldp attributes-match "{$$.interface-name}" matches "xe\-0/0/[0-1].*"
set event-options policy ldp then event-script jncie.slax arguments interface "{$$.interface-
name}"

set event-options event-script dampen dampen-options cpu-factor 40
set event-options event-script dampen dampen-options line-interval 1
set event-options event-script dampen dampen-options time-interval 500000
set event-options event-script traceoptions flag all
set event-options event-script file jncie.slax
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Verification

Verify the LDP neighborships

```
jncie@vMX2> show ldp neighbor

Address                                 Interface       Label space ID      Hold time
192.168.12.1                            xe-0/0/0.0      1.1.1.1:0             10
192.168.22.1                            xe-0/0/1.0      1.1.1.1:0             12

[edit]
jncie@vMX2# deactivate protocols ldp interface xe-0/0/0

[edit]
jncie@vMX2# commit
commit complete

[edit]
jncie@vMX2# activate protocols ldp interface xe-0/0/0

[edit]
jncie@vMX2# commit
commit complete

[edit]
jncie@vMX2# deactivate protocols ldp interface xe-0/0/0

[edit]
jncie@vMX2# commit
commit complete

[edit]
jncie@vMX2# activate protocols ldp interface xe-0/0/0

[edit]
jncie@vMX2# commit
commit complete

jncie@vMX2> show log messages | match cscript

Oct 30 13:32:03  vMX2 cscript.crypto: dampen is set with time-interval: 500000, cpu-factor:
40, line-interval: 1
Oct 30 13:32:06  vMX2 cscript.crypto: The xe-0/0/0.0 interface is causing LDP to bounce more
than two times in the one minute
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 7.6: On Box Scripting 3

- Add the following commit script with name “jncie2.slax” to device vMX1
- Ensure that if “family mpls” is configured on interface xe-0/0/0 that the interface is disabled

Solution

vMX1:

Enter the shell on device vMX1 and go to the /var/db/scripts/commit directory

```
jncie@vMX1# run start shell

% cd /var/db/scripts/commit/
```

Use the vi editor and create file “jncie2.slax”. Insert the following script content as mentioned in the task requirements:

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
        for-each (interfaces/interface[name != 'xe-0/0/0']/unit) {
            if ( (family/iso)) {
                call jcs:emit-change () {
                        with $content= {
                            <disable>;
                        }
                    }
                }
            }
}
```

There are two errors in the above script. The for-each statement is executed for all interfaces except (**!=**) interface ‘xe-0/0/0’. The task requirement states that only interface xe-0/0/0 should be disabled. Also, the above script checks if “family iso” is configured on the interface and then disables it. The task requirement states that the interface should be disabled if “family mpls” is configured on interfaces “xe-0/0/0”. If you fix both errors your final script should look like:

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
        for-each (interfaces/interface[name = 'xe-0/0/0']/unit) {
            if ( (family/mpls)) {
                call jcs:emit-change () {
                        with $content= {
                            <disable>;
                        }
                    }
                }
            }
}
```

Save file “jncie2.slax” in the /var/db/scripts/commit directory of device vMX1.

Exit the shell and activate the commit script.

```
% exit
exit

[edit]
jncie@vMX1# set system scripts commit file jncie2.slax

jncie@vMX1# commit
commit complete
```

Check the configuration on interface xe-0/0/0. Remember that we already had MPLS enabled from the previous task (7.5) where we set up LDP between vMX1 and vMX2.

```
[edit]
jncie@vMX1# show interfaces xe-0/0/0
unit 0 {
    disable;
    family inet {
        address 192.168.12.1/24;
    }
    family mpls;
}
```

It seems the commit script works.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

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-DC 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: Device Management

Task 1.1: Zero Touch Provisioning (ZTP)

- Ensure that vQFX1 is able to perform ZTP using interface em0.
- The DHCP server that is running on the CentOS host should assign IP address 10.10.20.8 based on the MAC address of vQFX1 em0 interface.
- During ZTP, vQFX1 should automatically download the configuration file /var/pub/ztp/vqfx1 using FTP from the Centos host.
- Edit the */etc/dhcp/dhcpd.conf* file using the “vi” editor. Do not use another editor.

Task 1.2 : Junos Space

- Ensure that vMX1 is discovered in Junos Space using SNMP version 1/2c. Use SNMP community “jncie”
- Make sure that the configuration file of vMX1 is stored on the Junos Space server.

vMX1:

```
[edit]
jncie@vMX1# set snmp community jncie
```

Task 1.3: On Box Scripting

- Ensure that interface AE1 on vMX1 is disabled by the Junos OS when the OSPF adjacency with neighbor 10.20.121.3 fails 3 times within 2 minutes. Do not use an external script.
- Ensure that on-box Python scripts are allowed on all vMX devices in DC1 and DC2
- Ensure that Python commit scripts are allowed to be stored off-box. Other commit scripts should only by allowed to be stored on-box.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX1:

```
set event-options policy jncie events rpd_ospf_nbrdown
set event-options policy jncie within 120 trigger after
set event-options policy jncie within 120 trigger 2
set event-options policy jncie within 120 events rpd_ospf_nbrup
set event-options policy jncie attributes-match rpd_ospf_nbrdown.neighbor-address equals
rpd_ospf_nbrup.neighbor-address
set event-options policy jncie attributes-match "{$$.neighbor-address}" matches 10.20.121.3
set event-options policy jncie then change-configuration commands "set interfaces ae1 disable"
```

vMX1, vMX2, vMX3 and vMX4:

```
set system scripts op no-allow-url
set system scripts op allow-url-for-python
set system scripts language python
```

Task 1.4: NETCONF

- There is a Python script that uses a RFC 4741 compliant NETCONF protocol to manage the configuration of vMX1. Ensure that the Python script can connect to vMX1 over a secured shell session. Allow 6 concurrent connections and a maximum of 2 sessions per minute.

vMX1:

```
set system services netconf ssh connection-limit 6
set system services netconf ssh rate-limit 2
set system services netconf rfc-compliant
```

Task 2: Underlay

Task 2.1: Multi Chassis Link Aggregation (MC-LAG)

Configure an active/active Multi Chassis LAG on MX1 and MX2 according to the following requirements:

- Interface AE0 towards vQFX1 should use VLAN 121. vMX1 should use IP address 10.20.121.1/24, vMX2 should use IP address 10.20.121.2/24 and vQFX1 should use IP address 10.20.121.3/24
- Ensure that multi chassis protection is configured
- Configure a redundant ICL link. Use VLAN 12 for the ICCP connection. vMX1 should use IP address 10.21.12.1/30 and vMX2 should use IP address 10.21.12.2/30.
- Configure a redundant default gateway with IP address 10.20.121.254. vMX1 should always be the master, also when recovered from a failure. Use MD5 password “jncie”.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX1:

```
set bridge-domains iccp domain-type bridge
set bridge-domains iccp vlan-id 12
set bridge-domains iccp routing-interface irb.12
set bridge-domains mc-lag domain-type bridge
set bridge-domains mc-lag vlan-id 121
set bridge-domains mc-lag routing-interface irb.121

set interfaces xe-0/0/0 gigether-options 802.3ad ae2
set interfaces xe-0/0/1 gigether-options 802.3ad ae2
set interfaces xe-0/0/2 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 mac 00:12:12:12:12:12
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 01:01:01:01:01:01
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 multi-chassis-protection 10.21.12.2 interface ae2.0
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 121

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 aggregated-ether-options link-speed 10g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 encapsulation vlan-bridge
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 121
set interfaces ae2 unit 0 family bridge vlan-id-list 12

set interfaces irb unit 12 family inet address 10.21.12.1/30
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 virtual-address
10.20.121.254
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 priority 150
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 preempt
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 accept-data
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 authentication-
type md5
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 authentication-key
jncie

set protocols iccp local-ip-addr 10.21.12.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols iccp authentication-key jncie
set protocols iccp peer 10.21.12.2 redundancy-group-id-list 10
set protocols iccp peer 10.21.12.2 liveness-detection minimum-interval 3000
set chassis aggregated-devices ethernet device-count 5
set multi-chassis multi-chassis-protection 10.21.12.2 interface ae2
set switch-options service-id 1
```

vMX2:

```
set bridge-domains iccp domain-type bridge
set bridge-domains iccp vlan-id 12
set bridge-domains iccp routing-interface irb.12
set bridge-domains mc-lag domain-type bridge
set bridge-domains mc-lag vlan-id 121
set bridge-domains mc-lag routing-interface irb.121

set interfaces xe-0/0/0 gigether-options 802.3ad ae2
set interfaces xe-0/0/1 gigether-options 802.3ad ae2
set interfaces xe-0/0/3 gigether-options 802.3ad ae0

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 mac 00:12:12:12:12:12
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 00:01:01:01:01:01
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 multi-chassis-protection 10.21.12.1 interface ae2.0
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 121

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 aggregated-ether-options link-speed 10g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 encapsulation vlan-bridge
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 121
set interfaces ae2 unit 0 family bridge vlan-id-list 12

set interfaces irb unit 12 family inet address 10.21.12.2/30
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 virtual-address
10.20.121.254
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 accept-data
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 authentication-
type md5
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 authentication-key
jncie

set protocols iccp local-ip-addr 10.21.12.2
set protocols iccp authentication-key jncie
set protocols iccp peer 10.21.12.1 redundancy-group-id-list 10
set protocols iccp peer 10.21.12.1 liveness-detection minimum-interval 3000
set chassis aggregated-devices ethernet device-count 5
set multi-chassis multi-chassis-protection 10.21.12.1 interface ae2
set switch-options service-id 1
```

vQFX1:

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0

set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members jncielag

set interfaces irb unit 121 family inet address 10.20.121.3/24

set vlans jncielag vlan-id 121
set vlans jncielag l3-interface irb.121
```

Task 2.2: Layer 2 Underlay Configuration

- Configure the Layer 2 interfaces between the vQFX devices and the servers according to the following requirements:
- Configure interface xe-0/0/1 between S1 and vQFX2. VLAN 10, 20, and 30 should be enabled on this interface

vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl30

set vlans vl10 vlan-id 10
set vlans vl20 vlan-id 20
set vlans vl30 vlan-id 30
```

- Configure interface AE10 between S3 and vQFX3, vQFX4 and vQFX5. Ensure that vQFX3, vQFX4 and vQFX5 initiate LACP sessions. VLAN 10, 20 and 30 should be enabled on this interface

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3, VQFX4 and VQFX5:

```
set vlans vl10 vlan-id 10
set vlans vl20 vlan-id 20
set vlans vl30 vlan-id 30

set chassis aggregated-devices ethernet device-count 11
set interfaces xe-0/0/1 ether-options 802.3ad ae10

set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 00:11:11:11:11:22
set interfaces ae10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae10 unit 0 family ethernet-switching vlan members vl10
set interfaces ae10 unit 0 family ethernet-switching vlan members vl20
set interfaces ae10 unit 0 family ethernet-switching vlan members vl30
```

- Configure interface xe-0/0/3 on vQFX3, interfaces xe-0/0/3 on vQFX4 and interface xe-0/0/1 on vQFX6. VLAN 10, 20, and 30 should be enabled on these interfaces.

vQFX3 and VQFX4:

```
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl30
```

vQFX6:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl30

set vlans vl10 vlan-id 10
set vlans vl20 vlan-id 20
set vlans vl30 vlan-id 30
```

Task 2.3: Layer 2 Underlay Features

- Ensure traffic towards Server 10.150.10.5 is flooded within VLAN 10.

vQFX6:

```
set vlans vl10 switch-options no-mac-learning
```

- Configure VQFX1 VLAN 121 with a mac-address limit. Set the limit to the number of mac-addresses used by vMX1 and vMX2 in VLAN 121. Log a message and drop packets if the limit is reached.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX1:

```
set vlans jncielag switch-options interface-mac-limit 3
set vlans jncielag switch-options interface-mac-limit packet-action drop-and-log
```

Task 2.4: BUM Traffic Handling

- Ensure that only Broadcast and Unknown Unicast traffic is limited to 100Mbit on vQFX2
- Do not use the default profile
- Use an apply-group to activate the configuration for all 10g layer2 interfaces on vQFX2.

vQFX2:

```
set groups stormcontrol interfaces <xe-*> unit <*> family ethernet-switching storm-control
jncie
set interfaces apply-groups stormcontrol
set interfaces xe-0/0/0 apply-groups-except stormcontrol
set forwarding-options storm-control-profiles jncie all bandwidth-percentage 1
set forwarding-options storm-control-profiles jncie all no-multicast
```

Task 2.5: Layer 2 Redundancy

- Ensure that AE10 towards Server 3 keeps forwarding traffic for VLAN10, 20 and 30 if Server3 cannot exchange LACP PDUs during a server upgrade.

vQFX3, VQFX4, vQFX5:

```
set interfaces ae10 aggregated-ether-options lacp force-up
```

Task 2.6: Layer 3 Underlay MAIN DC

- Configure an untagged IP interface between vSRX and VQFX6. vSRX1 should use IP address 10.40.3.1/30 and vQFX6 should use IP address 10.40.3.2/30. Use only static routing between vQFX6 and vSRX.

vQFX6:

```
set interfaces xe-0/0/2 unit 0 family inet address 10.40.3.2/30
set routing-options static route 192.168.1.0/24 next-hop 10.40.3.1
```

vSRX1:

```
set interfaces ge-0/0/0 unit 0 family inet address 10.40.3.1/30
set routing-options static route 192.168.1.6/32 next-hop 10.40.3.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.7: Layer 3 Underlay Routing

- Configure OSPFv2 between vMX1, vMX2, vQFX1 according to the following requirements:
    - All interfaces should be configured in the backbone area
    - Ensure that OSPFv2 adjacencies between vQFX1, vMX1 and vMX2 are in the “full” state.

vMX1:

```
delete routing-options router-id
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 mac
00:05:86:e7:28:f0
set protocols ospf area 0.0.0.0 interface irb.121
```

vMX2:

```
delete routing-options router-id
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 mac
00:05:86:34:33:f0
set protocols ospf area 0.0.0.0 interface irb.121
```

vQFX1:

```
set protocols ospf area 0.0.0.0 interface irb.121
```

- Configure OSPFv2 between vQFX1 and vQFX2 according to the following requirements:
    - All interfaces should be configured in the backbone area
    - Ensure there is no DR/BDR election on interface between vQFX1 and vQFX2.
    - vQFX1 should use IP address 10.20.1.1/30 and vQFX2 should use IP address 10.20.1.2/30
    - Do not use VLANs

vQFX1:

```
set protocols ospf area 0 interface xe-0/0/0.0 interface-type p2p
set interfaces xe-0/0/0 unit 0 family inet address 10.20.12.1/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX2:

```
set protocols ospf area 0 interface xe-0/0/0.0 interface-type p2p
set interfaces xe-0/0/0 unit 0 family inet address 10.20.12.2/30
```

Task 2.8: EBGP Underlay

- Configure the IP fabric in DC2 according to the following requirements:

Device	**BGP AS**
vMX3	64523
vMX4	64524
vQFX3	64513
vQFX4	64514
vQFX5	64515

- Configure BFD on all EBGP sessions for all devices in DC2s IP fabric.
- Ensure that BFD sessions do no support BGP sessions that peer on loopback interfaces
- Enable BFD MD5 authentication with password “jncie”
- Ensure that traffic is load-balanced if there are equal cost paths towards the same destination.

vMX3:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64523
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.33.2 peer-as 64513
set protocols bgp group underlay neighbor 10.30.34.2 peer-as 64514
set protocols bgp group underlay neighbor 10.30.35.2 peer-as 64515
set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet
set routing-options forwarding-table export lbpp
```

vMX4:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64524
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.37.2 peer-as 64513
set protocols bgp group underlay neighbor 10.30.38.2 peer-as 64514
set protocols bgp group underlay neighbor 10.30.36.2 peer-as 64515

set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet
set routing-options forwarding-table export lbpp
```

vQFX3:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64513
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.33.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.37.1 peer-as 64524

set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet
set routing-options forwarding-table export lbpp
```

vQFX4:

```
delete interfaces xe-0/0/2
delete interfaces xe-0/0/4
set interfaces xe-0/0/2 unit 0 family inet address 10.30.34.2/30
set interfaces xe-0/0/4 unit 0 family inet address 10.30.38.2/30

set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64514
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.34.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.38.1 peer-as 64524

set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet

set routing-options forwarding-table export lbpp
```

vQFX5:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64515
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.35.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.36.1 peer-as 64524

set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet

set routing-options forwarding-table export lbpp
```

Task 3: Controllerless Overlay

Task 3.1: EVPN Overlay Configuration

Configure an EVPN overlay between DC1, DC2 and MAIN DC according to the following requirements:

- Use BGP AS: 65535
- vQFX2, vQFX3, vQFX4, vQFX5 and vQFX6 should only have EVPN overlay sessions with vMX3 and vMX4.
- vMX3 and vMX4 should **not** initiate BGP sessions used for the EVPN overlay. Do not use the “passive” command or a firewall filter to accomplish this task.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

NOTE: *You need to complete the DCI tasks before the overlay sessions between DC1, DC2 and MAIN DC can be established.*

vMX3:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.33
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 192.168.1.33
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay allow 192.168.1.0/29

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vMX4:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.44
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 192.168.1.44
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay allow 192.168.1.0/29

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX2:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set interfaces lo0 unit 0 family inet address 192.168.1.2/32
set protocols ospf area 0.0.0.0 interface lo0.0
```

vQFX3:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.3
set protocols bgp group overlay family evpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX4:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

delete interfaces lo0
set interfaces lo0 unit 0 family inet address 192.168.1.4

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX5:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.5
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

delete interfaces lo0
set interfaces lo0 unit 0 family inet address 192.168.1.5

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX6:

```
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.6
set protocols bgp group overlay family evpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44
```

Task 3.2: EVPN/VLAN Service Configuration

- Configure an EVPN service called "**vpn1**" for VLAN 10 and VLAN 20.
- Configure an EVPN service called "**vpn2**" for VLAN30.

EVPN/VLAN service	**VLAN**	**vMX3 gateway address**	**vMX4 gateway address**
vpn1	10	10.150.10.100 10.150.10.254	10.150.10.101 10.150.10.254
	20	10.150.20.100 10.150.20.254	10.150.20.101 10.150.20.254
vpn2	30	10.150.30.100 10.150.30.254	10.150.30.101 10.150.30.254

- Ensure that the following communities are used for each EVPN/VLAN service as shown in the following table:

EVPN/VXLAN service	**Community name**	**Community value**
EVPN1	evpn	1:1
	vpn1	1:1000
	vni-10	1:10
	vni-20	1:20
EVPN2	evpn	1:1
	vpn2	1:2000
	vni-30	1:30

vMX3:

```
set routing-instances vpn1 route-distinguisher 192.168.1.33:1
set routing-instances vpn2 route-distinguisher 192.168.1.33:2

set interfaces irb unit 10 family inet address 10.150.10.100/24 virtual-gateway-address
10.150.10.254
set interfaces irb unit 20 family inet address 10.150.20.100/24 virtual-gateway-address
10.150.20.254
set interfaces irb unit 30 family inet address 10.150.30.100/24 virtual-gateway-address
10.150.30.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set routing-instances vpn1 route-distinguisher 192.168.1.44:1
set routing-instances vpn2 route-distinguisher 192.168.1.44:2

set interfaces irb unit 10 family inet address 10.150.10.101/24 virtual-gateway-address
10.150.10.254
set interfaces irb unit 20 family inet address 10.150.20.101/24 virtual-gateway-address
10.150.20.254
set interfaces irb unit 30 family inet address 10.150.30.101/24 virtual-gateway-address
10.150.30.254
```

vMX3 and vMX4:

```
set routing-instances vpn1 vtep-source-interface lo0.0
set routing-instances vpn1 instance-type virtual-switch
set routing-instances vpn1 vrf-import vpn1-import
set routing-instances vpn1 vrf-target target:1:1000
set routing-instances vpn1 protocols evpn encapsulation vxlan
set routing-instances vpn1 protocols evpn extended-vni-list 10
set routing-instances vpn1 protocols evpn extended-vni-list 20
set routing-instances vpn1 protocols evpn vni-options vni 10 vrf-target target:1:10
set routing-instances vpn1 protocols evpn vni-options vni 20 vrf-target target:1:20
set routing-instances vpn1 protocols evpn multicast-mode ingress-replication
set routing-instances vpn1 bridge-domains ie-10 domain-type bridge
set routing-instances vpn1 bridge-domains ie-10 vlan-id 10
set routing-instances vpn1 bridge-domains ie-10 routing-interface irb.10
set routing-instances vpn1 bridge-domains ie-10 vxlan vni 10
set routing-instances vpn1 bridge-domains ie-10 vxlan ingress-node-replication
set routing-instances vpn1 bridge-domains ie-20 domain-type bridge
set routing-instances vpn1 bridge-domains ie-20 vlan-id 20
set routing-instances vpn1 bridge-domains ie-20 routing-interface irb.20
set routing-instances vpn1 bridge-domains ie-20 vxlan vni 20
set routing-instances vpn1 bridge-domains ie-20 vxlan ingress-node-replication

set routing-instances vpn2 vtep-source-interface lo0.0
set routing-instances vpn2 instance-type virtual-switch
set routing-instances vpn2 vrf-import vpn2-import
set routing-instances vpn2 vrf-target target:1:2000
set routing-instances vpn2 protocols evpn encapsulation vxlan
set routing-instances vpn2 protocols evpn extended-vni-list 30
set routing-instances vpn2 protocols evpn vni-options vni 30 vrf-target target:1:30
set routing-instances vpn2 protocols evpn multicast-mode ingress-replication
set routing-instances vpn2 bridge-domains ie-30 domain-type bridge
set routing-instances vpn2 bridge-domains ie-30 vlan-id 30
set routing-instances vpn2 bridge-domains ie-30 routing-interface irb.30
set routing-instances vpn2 bridge-domains ie-30 vxlan vni 30
set routing-instances vpn2 bridge-domains ie-30 vxlan ingress-node-replication

set policy-options policy-statement vpn1-import term 1 from community evpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options policy-statement vpn1-import term 1 then accept
set policy-options policy-statement vpn1-import term 2 from community vpn1
set policy-options policy-statement vpn1-import term 2 then accept
set policy-options policy-statement vpn1-import term 3 from community vni-10
set policy-options policy-statement vpn1-import term 3 then accept
set policy-options policy-statement vpn1-import term 4 from community vni-20
set policy-options policy-statement vpn1-import term 4 then accept

set policy-options policy-statement vpn2-import term 1 from community evpn
set policy-options policy-statement vpn2-import term 1 then accept
set policy-options policy-statement vpn2-import term 2 from community vpn2
set policy-options policy-statement vpn2-import term 2 then accept
set policy-options policy-statement vpn2-import term 3 from community vni-30
set policy-options policy-statement vpn2-import term 3 then accept

set policy-options community evpn members target:1:1
set policy-options community vni-10 members target:1:10
set policy-options community vni-20 members target:1:20
set policy-options community vni-30 members target:1:30
set policy-options community vpn1 members target:1:1000
set policy-options community vpn2 members target:1:2000
```

vQFX2:

```
set switch-options route-distinguisher 192.168.1.2:1
set routing-options router-id 192.168.1.2
```

vQFX3:

```
set switch-options route-distinguisher 192.168.1.3:1
set routing-options router-id 192.168.1.3
```

vQFX4:

```
set switch-options route-distinguisher 192.168.1.4:1
set routing-options router-id 192.168.1.4
```

vQFX5:

```
set switch-options route-distinguisher 192.168.1.5:1
set routing-options router-id 192.168.1.5
```

vQFX6:

```
set switch-options route-distinguisher 192.168.1.6:1
set routing-options router-id 192.168.1.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX2, vQFX3, vQFX4, vQFX5 and vQFX6:

```
set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list 10
set protocols evpn extended-vni-list 20
set protocols evpn extended-vni-list 30
set protocols evpn multicast-mode ingress-replication
set protocols evpn vni-options vni 10 vrf-target export target:1:10
set protocols evpn vni-options vni 20 vrf-target export target:1:20
set protocols evpn vni-options vni 30 vrf-target export target:1:30

set switch-options vtep-source-interface lo0.0
set switch-options vrf-import vpn-import
set switch-options vrf-target target:1:1

set vlans vl10 vxlan vni 10
set vlans vl10 vxlan ingress-node-replication
set vlans vl20 vxlan vni 20
set vlans vl20 vxlan ingress-node-replication
set vlans vl30 vxlan vni 30
set vlans vl30 vxlan ingress-node-replication

set policy-options policy-statement vpn-import term 1 from community evpn
set policy-options policy-statement vpn-import term 1 then accept
set policy-options policy-statement vpn-import term 2 from community vpn1
set policy-options policy-statement vpn-import term 2 then accept
set policy-options policy-statement vpn-import term 3 from community vpn2
set policy-options policy-statement vpn-import term 3 then accept
set policy-options policy-statement vpn-import term 4 from community vni-10
set policy-options policy-statement vpn-import term 4 then accept
set policy-options policy-statement vpn-import term 5 from community vni-20
set policy-options policy-statement vpn-import term 5 then accept
set policy-options policy-statement vpn-import term 6 from community vni-30
set policy-options policy-statement vpn-import term 6 then accept

set policy-options community evpn members target:1:1
set policy-options community vni-10 members target:1:10
set policy-options community vni-20 members target:1:20
set policy-options community vni-30 members target:1:30
set policy-options community vpn1 members target:1:1000
set policy-options community vpn2 members target:1:2000
```

Task 3.3: EVPN/VXLAN Service Tuning

- Configure vQFX6 to age out the address of remote VTEPs after 6 minutes, if all the MAC addresses learned from that VTEP age out. Your solution must also account for VNIs that might be added in the future.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX6:

```
set groups unreachable vlans <*> vxlan unreachable-vtep-aging-timer 360

set vlans apply-groups unreachable
set vlans default apply-groups-except unreachable
```

Task 3.4: EVPN Signaling

- Ensure that vQFX2 receives the following EVPN type 1 routes:

1:192.168.1.3:1::11111111111111122::0/304

1:192.168.1.4:1::11111111111111122::0/304

1:192.168.1.5:1::11111111111111122::0/304

vQFX3, VQFX4 and vQFX5:

```
set interfaces ae10 esi 00:11:11:11:11:11:11:11:22 all-active
```

- Ensure that vQFX2 receives the following EVPN type 2 routes:

2:192.168.1.33:2::30::00:99:00:00:00:10/304

2:192.168.1.33:1::20::00:99:00:00:00:20/304

2:192.168.1.33:1::10::00:99:00:00:00:30/304

vMX3 and vMX4:

```
set routing-instances vpn1 route-distinguisher 192.168.1.33:1
set interfaces irb unit 10 mac 00:99:00:00:00:30
set interfaces irb unit 20 mac 00:99:00:00:00:20

set routing-instances vpn2 route-distinguisher 192.168.1.33:2
set interfaces irb unit 30 mac 00:99:00:00:00:10
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4: Data Center Interconnect

Task 4.1: DCI Configuration

- Configure two L3VPN services according to the following information:
    - The first L3VPN service should be named “underlay” on vMX5 and is used for the underlay between DC1, DC2 and MAIN DC.

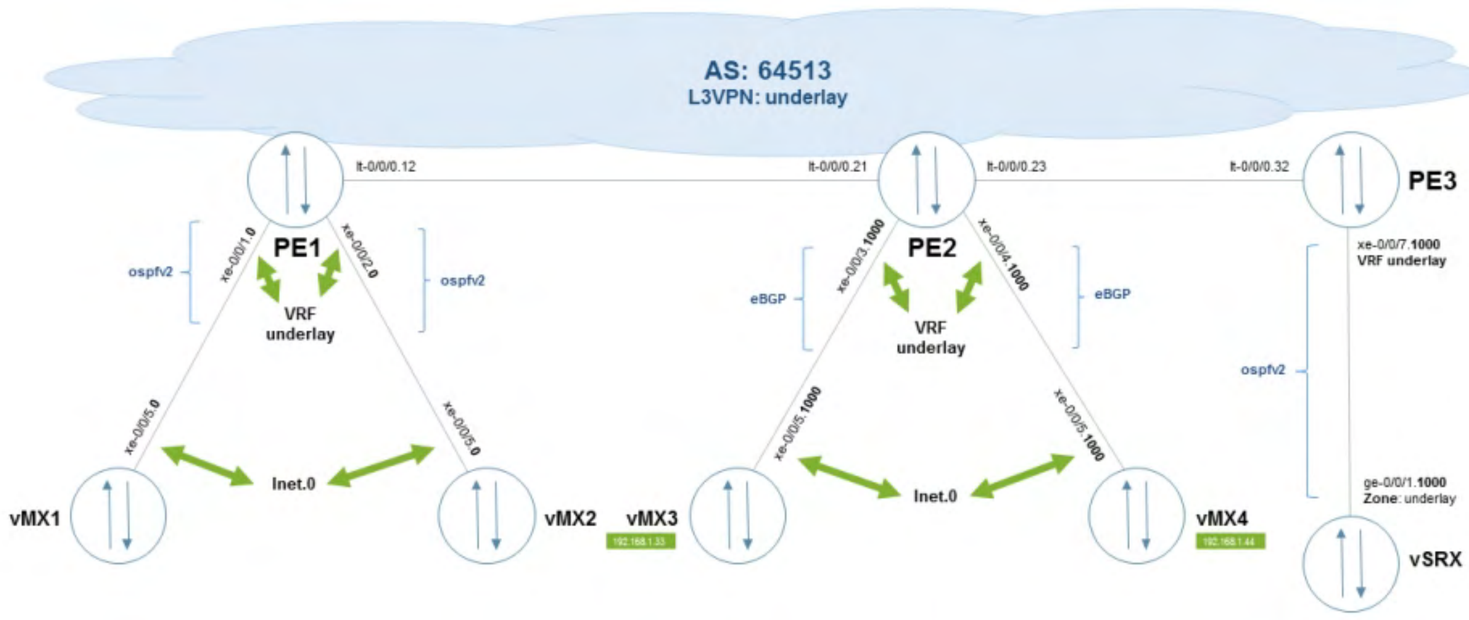

vMX5 (PE1):

```
set logical-systems PE1 interfaces lt-0/0/0 unit 12
set logical-systems PE1 interfaces xe-0/0/1 unit 0
set logical-systems PE1 interfaces xe-0/0/2 unit 0
set logical-systems PE1 interfaces lo0 unit 1

set logical-systems PE1 protocols bgp group l3vpn type internal
set logical-systems PE1 protocols bgp group l3vpn local-address 172.17.1.1
set logical-systems PE1 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE1 protocols bgp group l3vpn local-as 64513
set logical-systems PE1 protocols bgp group l3vpn neighbor 172.17.1.2
set logical-systems PE1 protocols bgp group l3vpn neighbor 172.17.1.3

set logical-systems PE1 protocols ospf area 0.0.0.0 interface all hello-interval 1
set logical-systems PE1 protocols ospf area 0.0.0.0 interface all dead-interval 4

set logical-systems PE1 protocols ldp interface all

set logical-systems PE1 protocols mpls interface all

set logical-systems PE1 policy-options policy-statement 2ospf from protocol bgp
set logical-systems PE1 policy-options policy-statement 2ospf then accept

set logical-systems PE1 policy-options policy-statement lbpp then load-balance per-packet
set logical-systems PE1 policy-options policy-statement lbpp then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set logical-systems PE1 routing-instances underlay instance-type vrf
set logical-systems PE1 routing-instances underlay interface xe-0/0/1.0
set logical-systems PE1 routing-instances underlay interface xe-0/0/2.0
set logical-systems PE1 routing-instances underlay route-distinguisher 1:1
set logical-systems PE1 routing-instances underlay vrf-target target:64513:1
set logical-systems PE1 routing-instances underlay vrf-table-label
set logical-systems PE1 routing-instances underlay protocols ospf export 2ospf
set logical-systems PE1 routing-instances underlay protocols ospf area 0.0.0.0 interface xe-
0/0/1.0
set logical-systems PE1 routing-options forwarding-table export lbpp

set interfaces lo0 unit 1 family inet address 172.17.1.1/32
```

vMX5 (PE2):

```
set logical-systems PE2 interfaces lt-0/0/0 unit 21
set logical-systems PE2 interfaces lt-0/0/0 unit 23
set logical-systems PE2 interfaces xe-0/0/3 unit 1000
set logical-systems PE2 interfaces xe-0/0/4 unit 1000
set logical-systems PE2 interfaces lo0 unit 2

set logical-systems PE2 protocols bgp group l3vpn type internal
set logical-systems PE2 protocols bgp group l3vpn local-address 172.17.1.2
set logical-systems PE2 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE2 protocols bgp group l3vpn local-as 64513
set logical-systems PE2 protocols bgp group l3vpn neighbor 172.17.1.1
set logical-systems PE2 protocols bgp group l3vpn neighbor 172.17.1.3

set logical-systems PE2 protocols ospf area 0.0.0.0 interface all hello-interval 1
set logical-systems PE2 protocols ospf area 0.0.0.0 interface all dead-interval 4

set logical-systems PE2 protocols ldp interface all

set logical-systems PE2 protocols mpls interface all

set logical-systems PE2 policy-options policy-statement lbpp then load-balance per-packet
set logical-systems PE2 policy-options policy-statement lbpp then accept

set logical-systems PE2 routing-instances underlay instance-type vrf
set logical-systems PE2 routing-instances underlay interface xe-0/0/3.1000
set logical-systems PE2 routing-instances underlay interface xe-0/0/4.1000
set logical-systems PE2 routing-instances underlay route-distinguisher 1:2
set logical-systems PE2 routing-instances underlay vrf-target target:64513:1
set logical-systems PE2 routing-instances underlay vrf-table-label
set logical-systems PE2 routing-instances underlay protocols bgp group underlay type external
set logical-systems PE2 routing-instances underlay protocols bgp group underlay local-as 64513
set logical-systems PE2 routing-instances underlay protocols bgp group underlay neighbor
10.30.3.2 peer-as 64523
set logical-systems PE2 routing-instances underlay protocols bgp group underlay neighbor
10.30.4.2 peer-as 64524
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set logical-systems PE2 routing-options forwarding-table export lbpp

set interfaces lo0 unit 2 family inet address 172.17.1.2/32
```

vMX5 (PE3):

```
set logical-systems PE3 interfaces lt-0/0/0 unit 32
set logical-systems PE3 interfaces xe-0/0/7 unit 2000
set logical-systems PE3 interfaces lo0 unit 3

set logical-systems PE3 protocols bgp group l3vpn type internal
set logical-systems PE3 protocols bgp group l3vpn local-address 172.17.1.3
set logical-systems PE3 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE3 protocols bgp group l3vpn local-as 64513
set logical-systems PE3 protocols bgp group l3vpn neighbor 172.17.1.1
set logical-systems PE3 protocols bgp group l3vpn neighbor 172.17.1.2

set logical-systems PE3 protocols ospf area 0.0.0.0 interface all hello-interval 1
set logical-systems PE3 protocols ospf area 0.0.0.0 interface all dead-interval 4

set logical-systems PE3 protocols ldp interface all

set logical-systems PE3 protocols mpls interface all

set logical-systems PE3 policy-options policy-statement 2ospf from protocol bgp
set logical-systems PE3 policy-options policy-statement 2ospf then accept

set logical-systems PE3 routing-instances underlay instance-type vrf
set logical-systems PE3 routing-instances underlay interface xe-0/0/7.1000
set logical-systems PE3 routing-instances underlay route-distinguisher 1:3
set logical-systems PE3 routing-instances underlay vrf-target target:64513:1
set logical-systems PE3 routing-instances underlay vrf-table-label
set logical-systems PE3 routing-instances underlay protocols ospf export 2ospf
set logical-systems PE3 routing-instances underlay protocols ospf area 0.0.0.0 interface xe-
0/0/7.1000

set logical-systems PE3 routing-options autonomous-system 64513

set interfaces lo0 unit 3 family inet address 172.17.1.3/32
```

vMX1 and vMX2:

```
set protocols ospf area 0.0.0.0 interface xe-0/0/5.0
set protocols ospf area 0.0.0.0 interface irb.121
```

vMX3:

```
set protocols bgp group underlay neighbor 10.30.3.1 peer-as 64513
set protocols bgp group underlay neighbor 10.30.3.1 local-as 64523
set protocols bgp group underlay neighbor 10.30.3.1 as-override
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay neighbor 10.30.33.2 as-override
set protocols bgp group underlay export myloopback

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vMX4:

```
set protocols bgp group underlay neighbor 10.30.4.1 peer-as 64513
set protocols bgp group underlay neighbor 10.30.4.1 as-override
set protocols bgp group underlay neighbor 10.30.37.2 as-override
set protocols bgp group underlay export myloopback

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vSRX:

```
set interfaces ge-0/0/1 unit 1000 description underlay
set interfaces ge-0/0/1 unit 1000 vlan-id 1000
set interfaces ge-0/0/1 unit 1000 family inet address 10.40.1.2/30

set security zones security-zone underlay host-inbound-traffic system-services ping
set security zones security-zone underlay host-inbound-traffic protocols ospf
set security zones security-zone underlay interfaces ge-0/0/1.1000
set security zones security-zone underlay interfaces ge-0/0/0.0

set policy-options policy-statement 2ospf term 1 from protocol static
set policy-options policy-statement 2ospf term 1 from route-filter 192.168.1.6/32 exact
set policy-options policy-statement 2ospf term 1 then accept

set protocols ospf export 2ospf
set protocols ospf area 0.0.0.0 interface ge-0/0/1.1000

delete routing-options router-id
set security address-book global address 192.168.1/24 192.168.1.0/24
set security policies from-zone underlay to-zone underlay policy jncie1 match source-address
192.168.1/24
set security policies from-zone underlay to-zone underlay policy jncie1 match destination-
address 192.168.1/24
set security policies from-zone underlay to-zone underlay policy jncie1 match application
junos-bgp
set security policies from-zone underlay to-zone underlay policy jncie1 match application
junos-icmp-ping
set security policies from-zone underlay to-zone underlay policy jncie1 match application
vxlan
set security policies from-zone underlay to-zone underlay policy jncie1 then permit
set applications application vxlan protocol udp
set applications application vxlan destination-port 4789
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- The second L3VPN should be named “dc-vpns” on vMX5 and is used for L3 routing for EVPN service “vpn1” and “vpn2” between DC2 and MAIN DC.

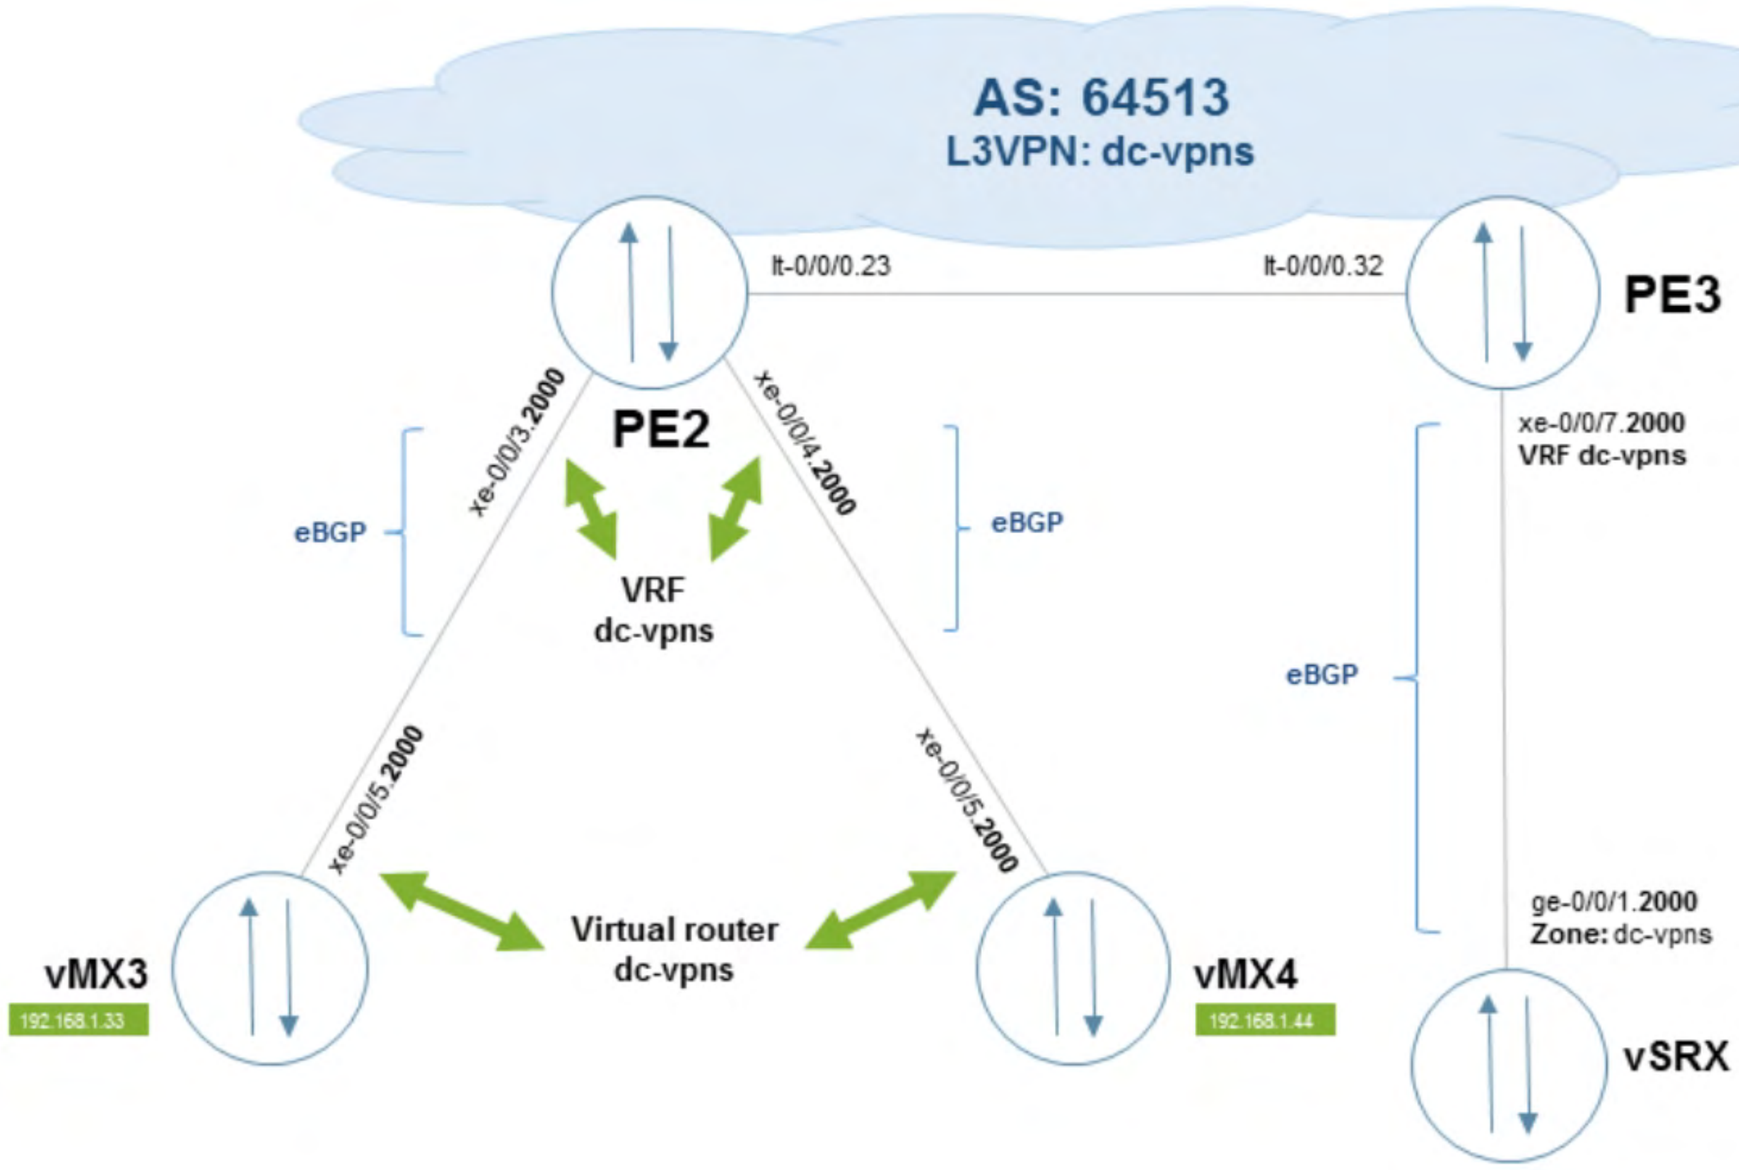

vMX5 (PE2):

```
set logical-systems PE2 routing-instances dc-vpns instance-type vrf
set logical-systems PE2 routing-instances dc-vpns interface xe-0/0/3.2000
set logical-systems PE2 routing-instances dc-vpns interface xe-0/0/4.2000
set logical-systems PE2 routing-instances dc-vpns route-distinguisher 2:2
set logical-systems PE2 routing-instances dc-vpns vrf-target target:64513:2
set logical-systems PE2 routing-instances dc-vpns vrf-table-label
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns type external
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns local-as 64513
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns neighbor
10.30.3.6 peer-as 64523
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns neighbor
10.30.4.6 peer-as 64524
```

vMX5 (PE3):

```
set logical-systems PE3 routing-instances dc-vpns instance-type vrf
set logical-systems PE3 routing-instances dc-vpns interface xe-0/0/7.2000
set logical-systems PE3 routing-instances dc-vpns route-distinguisher 2:3
set logical-systems PE3 routing-instances dc-vpns vrf-target target:64513:2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set logical-systems PE3 routing-instances dc-vpns vrf-table-label
set logical-systems PE3 routing-instances dc-vpns protocols bgp group external type external
set logical-systems PE3 routing-instances dc-vpns protocols bgp group external neighbor
10.40.2.2 peer-as 64516
```

vMX3:

```
set routing-instances dc-vpns instance-type virtual-router
set routing-instances dc-vpns interface xe-0/0/5.2000
set routing-instances dc-vpns interface irb.10
set routing-instances dc-vpns interface irb.20
set routing-instances dc-vpns interface irb.30
set routing-instances dc-vpns protocols bgp group dc-vpns type external
set routing-instances dc-vpns protocols bgp group dc-vpns export direct-2-bgp
set routing-instances dc-vpns protocols bgp group dc-vpns local-as 64523
set routing-instances dc-vpns protocols bgp group dc-vpns neighbor 10.30.3.5 peer-as 64513
set policy-options policy-statement direct-2-bgp term 1 from protocol direct
set policy-options policy-statement direct-2-bgp term 1 then accept
```

vMX4:

```
set routing-instances dc-vpns instance-type virtual-router
set routing-instances dc-vpns interface xe-0/0/5.2000
set routing-instances dc-vpns interface irb.10
set routing-instances dc-vpns interface irb.20
set routing-instances dc-vpns interface irb.30
set routing-instances dc-vpns protocols bgp group dc-vpns type external
set routing-instances dc-vpns protocols bgp group dc-vpns export direct-2-bgp
set routing-instances dc-vpns protocols bgp group dc-vpns local-as 64524
set routing-instances dc-vpns protocols bgp group dc-vpns neighbor 10.30.4.5 peer-as 64513
set policy-options policy-statement direct-2-bgp term 1 from protocol direct
set policy-options policy-statement direct-2-bgp term 1 then accept
```

vSRX:

```
set protocols bgp group external type external
set protocols bgp group external export default-2-bgp
set protocols bgp group external neighbor 10.40.2.1 peer-as 64513

set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1

set policy-options policy-statement default-2-bgp from protocol static
set policy-options policy-statement default-2-bgp from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-2-bgp then accept

set interfaces ge-0/0/1 unit 2000 vlan-id 2000
set interfaces ge-0/0/1 unit 2000 family inet address 10.40.2.2/30

delete interfaces ge-0/0/0.0 family inet address 10.40.2.1/31
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic system-
services ping
set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic
protocols bgp
```

Task 4.2: DC1 Redundancy

- Ensure that the link between PE1 and vMX1 is the primary path for traffic to and from DC1. Ensure that the link between PE1 and vMX2 is only used for traffic to and from DC1 if the primary path is down.

vMX2:

```
set protocols ospf area 0.0.0.0 interface xe-0/0/5.0 metric 1000
```

vMX5 (PE1):

```
set logical-systems PE1 routing-instances underlay protocols ospf area 0.0.0.0 interface xe-
0/0/2.0 metric 1000
```

Task 4.3: (Inter) Data Center Connectivity

- Ensure that the Servers in DC1, DC2 and MAIN DC have full connectivity between all VLANs.

Use the test addresses shown in the topology for verification purposes

Task 4.4: Data Center Connectivity to External Server

- Ensure that the Servers in DC1, DC2 and MAIN DC can ping server (8.8.8.8). The server is preconfigured on vMX5.
  - Traffic between the servers in DC1, DC2 , MAIN DC and the server should always transit the “dc-vpns” and “untrust” zones of the vSRX firewall.
  - The server should use destination IP address 1.1.1.2 if it sends ICMP echo replies to Servers in DC1, DC2 and MAIN DC.

vSRX:

```
set protocols bgp group external export default-2-bgp

set policy-options policy-statement default-2-bgp from protocol static
set policy-options policy-statement default-2-bgp from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-2-bgp then accept

set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security zones security-zone untrust interfaces ge-0/0/1.3000

set security policies from-zone dc-vpns to-zone untrust policy 1 match source-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match destination-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match application junos-icmp-
ping
set security policies from-zone dc-vpns to-zone untrust policy 1 then permit

set security nat source rule-set jncie from zone dc-vpns
set security nat source rule-set jncie to zone untrust
set security nat source rule-set jncie rule ie-1 match source-address 0.0.0.0/0
set security nat source rule-set jncie rule ie-1 match destination-address 0.0.0.0/0
set security nat source rule-set jncie rule ie-1 match application any
set security nat source rule-set jncie rule ie-1 then source-nat interface
```

Task 5: Security

Task 5.1: User Authentication and Authorization

- Configure a user called “jncie-1” on vSRX. Helpful information should be shown when user “jncie-1” logs in to the device. User “jncie-1” should only be allowed to display the device configuration and view system and interfaces statistics

vSRX:

```
set system login class jncie login-tip
set system login class jncie permissions view
set system login class jncie permissions view-configuration
set system login user jncie-1 uid 2000
set system login user jncie-1 class jncie
set system login user jncie-1 authentication encrypted-password "
$1$ZWoFbJTI$XL0oPIMXfUpBLmTZav5H50"
```

Task 5.2: SRX Configuration

- Configure the vSRX according to the following requirements:

  See answers in DCI underlay and dc-vpn sections

  - Interface ge-0/0/0.0 and ge-0/0/1.1000 should be in the “underlay” zone
  - Interface ge-0/0/1.2000 should be in the “dc-vpns” zone
  - Interface ge-0/0/1.3000 should be in the “untrust” zone
  - Configure a static default route towards the next-hop IP address (1.1.1.1) of AS 64513 in the untrust zone.
  - Allow ICMP PING traffic towards vSRX interfaces, except for interfaces in the “untrust” zone
  - Allow ICMP traffic between devices in the “underlay” zone

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Do not configure routing-instances on vSRX
- Only allow traffic if required to complete other tasks in this superlab. Do not allow other traffic

Task 5.3: RE Protection

- Configure control-plane protection on vQFX6 with the following requirements:
  - Allow incoming telnet and ssh sessions
  - Allow BGP sessions to vMX3 and vMX4. Make sure that if new BGP neighbors are configured in the "protocols bgp" stanza, you do not need to change your RE protection configuration.
  - Allow ICMP
  - Allow specific traffic needed by other tasks in this superlab.
  - Deny and count all other traffic. The control plane should not respond to denied traffic.

vQFX6:

```
set interfaces lo0 unit 0 family inet filter input re-protect

set policy-options prefix-list bgp apply-path "protocols bgp group <*> neighbor <*>"

set firewall family inet filter re-protect term 1 from source-prefix-list bgp
set firewall family inet filter re-protect term 1 from protocol tcp
set firewall family inet filter re-protect term 1 from source-port bgp
set firewall family inet filter re-protect term 1 then accept
set firewall family inet filter re-protect term 2 from protocol icmp
set firewall family inet filter re-protect term 2 then accept
set firewall family inet filter re-protect term 3 from protocol tcp
set firewall family inet filter re-protect term 3 from destination-port ssh
set firewall family inet filter re-protect term 3 from destination-port telnet
set firewall family inet filter re-protect term 3 then accept
set firewall family inet filter re-protect term 5 then count rejected
set firewall family inet filter re-protect term 5 then discard
```

Task 6: Class of Service

Task 6.1: Forwarding Classes

Configure the following forwarding classes and forwarding class sets on all vQFX devices.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Forwarding Class	Queue ID	No-Loss
BE	0	Disabled
FCOE	3	Enabled
NOLOSS	4	Enabled
NC	7	Disabled

Forwarding Class Set	Forwarding Class
LOSS	BE
	NC
NOLOSS	FCOE
	NOLOSS

vQFX devices:

```
set class-of-service forwarding-classes class BE queue-num 0
set class-of-service forwarding-classes class FCOE queue-num 3
set class-of-service forwarding-classes class FCOE no-loss
set class-of-service forwarding-classes class NOLOSS queue-num 4
set class-of-service forwarding-classes class NOLOSS no-loss
set class-of-service forwarding-classes class NC queue-num 7

set class-of-service forwarding-class-sets LOSS class BE
set class-of-service forwarding-class-sets LOSS class NC
set class-of-service forwarding-class-sets NOLOSS class FCOE
set class-of-service forwarding-class-sets NOLOSS class NOLOSS
```

Task 6.2: Schedulers and Traffic Profiles

- Configure a scheduler mapping named “LOSS” that includes the BE and NC schedulers on all vQFX devices.
- Configure a scheduler mapping named “NOLOSS” that includes the FCOE and NOLOSS schedulers on all vQFX devices.
- Configure traffic profiles for each scheduler

Scheduler name	Transmit rate	Buffer size	Priority
BE	55 %	remainder	low
FCOE	25 %	25 %	low
NOLOSS	19 %	19 %	low
NC	1 %	1 %	low

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX devices:

```
set class-of-service schedulers BE transmit-rate percent 55
set class-of-service schedulers BE buffer-size remainder
set class-of-service schedulers BE priority low

set class-of-service schedulers FCOE transmit-rate percent 25
set class-of-service schedulers FCOE buffer-size percent 29
set class-of-service schedulers FCOE priority low

set class-of-service schedulers NOLOSS transmit-rate percent 19
set class-of-service schedulers NOLOSS buffer-size percent 19
set class-of-service schedulers NOLOSS priority low

set class-of-service schedulers NC transmit-rate percent 1
set class-of-service schedulers NC buffer-size percent 1
set class-of-service schedulers NC priority low

set class-of-service scheduler-maps NOLOSS forwarding-class NOLOSS scheduler NOLOSS
set class-of-service scheduler-maps NOLOSS forwarding-class FCOE scheduler FCOE
set class-of-service scheduler-maps LOSS forwarding-class BE scheduler BE
set class-of-service scheduler-maps LOSS forwarding-class NC scheduler NC

set class-of-service traffic-control-profiles NOLOSS scheduler-map NOLOSS
set class-of-service traffic-control-profiles LOSS scheduler-map LOSS
```

Note: *vQFX does not support the usage of forwarding-class-sets and traffic profiles on xe-interfaces. For this reason, we do not ask you to activate them on the interfaces. Please make sure you understand how to apply them on physical QFX5100 devices interfaces.*

See the following example:

```
set class-of-service interfaces xe-* forwarding-class-set LOSS output-traffic-control-profile
LOSS
set class-of-service interfaces xe-* forwarding-class-set NOLOSS output-traffic-control-
profile NOLOSS
```

Task 6.3: Policing

- Configure vQFX2 so that FTP transfers from Server 1 are policed to 1g and should use the BE forwarding class with low loss-priority.

vQFX2:

```
set firewall family ethernet-switching filter jncie term 1 from port ftp-data
set firewall family ethernet-switching filter jncie term 1 from port ftp
set firewall family ethernet-switching filter jncie term 1 then accept
set firewall family ethernet-switching filter jncie term 1 then forwarding-class BE
set firewall family ethernet-switching filter jncie term 1 then loss-priority low
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set firewall family ethernet-switching filter jncie term 1 then policer 1g
set firewall family ethernet-switching filter jncie term 2 then accept

set firewall policer 1g if-exceeding bandwidth-limit 1g
set firewall policer 1g if-exceeding burst-size-limit 100m
set firewall policer 1g then discard

set interfaces xe-0/0/1 unit 0 family ethernet-switching filter input jncie
```

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-DC 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: Device Management

Task 1.1: Zero Touch Provisioning (ZTP)

- Ensure that vQFX1 is able to perform ZTP using interface em0.
- The DHCP server that is running on the CentOS host should assign IP address 10.10.20.8 based on the MAC address of vQFX1 em0 interface.
- During ZTP, vQFX1 should automatically download the configuration file /var/pub/ztp/vqfx1 using FTP from the Centos host.
- Edit the */etc/dhcp/dhcpd.conf* file using the “vi” editor.

Task 1.2 : Junos Space

- Ensure that all vMX devices in DC1 and DC2 are discovered in Junos Space using SNMP version 1/2c. Use SNMP community “jncie”.
- Make sure that the configuration file vMX devices are stored on the Junos Space server.

vMX1, vMX2, vMX3 and vXM4:

```
set snmp community jncie
set snmp trap-group space targets 10.10.20.101
```

Task 1.3: On Box Scripting

- Ensure that interface xe-0/0/5 on vMX1, vMX2, vMX3 and vMX4 is disabled by the Junos OS when it goes down 2 times within 1 minutes. Do not use an external script
- Ensure that event-policies are not allowed to be executed simultaneously.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Ensure that on-box Python scripts are allowed on all vMX devices in DC1 and DC2
- 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, vMX2, vMX3 and vMX4:

```
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 xe-0/0/5
set event-options policy jncie then change-configuration commands "set interfaces xe-0/0/5
disable"
```

vMX1, vMX2, vMX3 and vMX4:

```
set system scripts op no-allow-url
set system scripts op allow-url-for-python
set system scripts language python
```

Task 1.4: NETCONF

- There is a Python script that uses an RFC 4741 compliant NETCONF protocol to manage the configuration of vMX1. Ensure that the Python script can connect to vMX1 over a secured shell session. Allow 6 concurrent connections and a maximum of 2 sessions per minute.

vMX1:

```
set system services netconf ssh connection-limit 6
set system services netconf ssh rate-limit 2
set system services netconf rfc-compliant
```

Task 2: Underlay

Task 2.1: Layer 2 Underlay Configuration

Configure the Layer 2 interfaces between the vQFX devices and the servers according to the following requirements:

- Configure interface xe-0/0/1 between S1 and vQFX2. VLAN 10, 20 and 30 should be enabled on this interface
- Configure interface AE0 and xe-0/0/0 on vQFX1 as trunk interfaces for VLAN 121

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX1:

```
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members jncielag
set interfaces xe-0/0/0 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members jncielag
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0
set vlans jncielag vlan-id 121
set chassis aggregated-devices ethernet device-count 1
```

vQFX2:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl30

set vlans vl10 vlan-id 10
set vlans vl20 vlan-id 20
set vlans vl30 vlan-id 30
```

- Configure interface AE10 between S3 and vQFX3, vQFX4 and vQFX5. Ensure that vQFX3, vQFX4 and vQFX5 initiate LACP sessions. VLAN 10, 20 and 30 should be enabled on this interface.

vQFX3

```
set interfaces xe-0/0/3 ether-options 802.3ad ae10
```

VQFX4 and VQFX5:

```
set interfaces xe-0/0/1 ether-options 802.3ad ae10
```

vQFX3, VQFX4 and VQFX5:

```
set vlans vl10 vlan-id 10
set vlans vl20 vlan-id 20
set vlans vl30 vlan-id 30


set interfaces ae10 aggregated-ether-options lacp active
set interfaces ae10 aggregated-ether-options lacp system-id 00:11:11:11:11:22
set interfaces ae10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae10 unit 0 family ethernet-switching vlan members vl10
set interfaces ae10 unit 0 family ethernet-switching vlan members vl20
set interfaces ae10 unit 0 family ethernet-switching vlan members vl30
set chassis aggregated-devices ethernet device-count 11
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- Configure interface xe-0/0/1 on vQFX3, interfaces xe-0/0/3 on vQFX4. VLAN 10, 20 and 30 should be enabled on these interfaces.

vQFX3:

```
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members vl30
```

VQFX4:

```
set interfaces xe-0/0/3 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl10
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl20
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members vl30
```

Task 2.2: Multi Chassis Link Aggregation (MC-LAG)

Configure an active/active Multi Chassis LAG on MX1 and MX2 according to the following requirements:

- Interface AE0 towards vQFX1 should use VLAN 121. vMX1 should use IP address 10.20.121.1/24, vMX2 should use IP address 10.20.121.2/24 and vQFX2 should use IP address 10.20.121.3/24 on interface xe-0/0/0
- Ensure that multi chassis protection is configured. Do not configure multi chassis protection under the interface hierarchy.
- Configure a redundant ICL link. Use VLAN 12 for the ICCP connection. vMX1 should use IP address 10.21.12.1/30 and vMX2 should use IP address 10.21.12.2/30.
- Configure a redundant default gateway with IP address 10.20.121.254. vMX1 should always be the master, also when recovered from a failure. Use MD5 password “jncie”.

vMX1:

```
set bridge-domains iccp domain-type bridge
set bridge-domains iccp vlan-id 12
set bridge-domains iccp routing-interface irb.12
set bridge-domains mc-lag domain-type bridge
set bridge-domains mc-lag vlan-id 121
set bridge-domains mc-lag routing-interface irb.121

set interfaces xe-0/0/0 gigether-options 802.3ad ae2
set interfaces xe-0/0/1 gigether-options 802.3ad ae2
set interfaces xe-0/0/2 gigether-options 802.3ad ae0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 mac 00:12:12:12:12:12
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 01:01:01:01:01:01
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 multi-chassis-protection 10.21.12.2 interface ae2.0
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 121

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 aggregated-ether-options link-speed 10g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 encapsulation vlan-bridge
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 121
set interfaces ae2 unit 0 family bridge vlan-id-list 12

set interfaces irb unit 12 family inet address 10.21.12.1/30
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 mac
00:05:86:e7:28:f0
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 virtual-address
10.20.121.254
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 priority 150
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 preempt
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 accept-data
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 authentication-
type md5
set interfaces irb unit 121 family inet address 10.20.121.1/24 vrrp-group 1 authentication-key
jncie

set protocols iccp local-ip-addr 10.21.12.1
set protocols iccp authentication-key jncie
set protocols iccp peer 10.21.12.2 redundancy-group-id-list 10
set protocols iccp peer 10.21.12.2 liveness-detection minimum-interval 3000
set chassis aggregated-devices ethernet device-count 5
set multi-chassis multi-chassis-protection 10.21.12.2 interface ae2
set switch-options service-id 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX2:

```
set bridge-domains iccp domain-type bridge
set bridge-domains iccp vlan-id 12
set bridge-domains iccp routing-interface irb.12
set bridge-domains mc-lag domain-type bridge
set bridge-domains mc-lag vlan-id 121
set bridge-domains mc-lag routing-interface irb.121

set interfaces xe-0/0/0 gigether-options 802.3ad ae2
set interfaces xe-0/0/1 gigether-options 802.3ad ae2
set interfaces xe-0/0/3 gigether-options 802.3ad ae0

set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 mac 00:12:12:12:12:12
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-id 01:01:01:01:01:01
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 1
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control standby
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 multi-chassis-protection 10.21.12.1 interface ae2.0
set interfaces ae0 unit 0 family bridge interface-mode trunk
set interfaces ae0 unit 0 family bridge vlan-id-list 121

set interfaces ae2 flexible-vlan-tagging
set interfaces ae2 encapsulation flexible-ethernet-services
set interfaces ae2 aggregated-ether-options minimum-links 1
set interfaces ae2 aggregated-ether-options link-speed 10g
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 unit 0 encapsulation vlan-bridge
set interfaces ae2 unit 0 family bridge interface-mode trunk
set interfaces ae2 unit 0 family bridge vlan-id-list 121
set interfaces ae2 unit 0 family bridge vlan-id-list 12


set interfaces irb unit 12 family inet address 10.21.12.2/30
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 mac
00:05:86:34:33:f0
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 virtual-address
10.20.121.254
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 accept-data
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 authentication-
type md5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set interfaces irb unit 121 family inet address 10.20.121.2/24 vrrp-group 1 authentication-key
jncie

set protocols iccp local-ip-addr 10.21.12.2
set protocols iccp authentication-key jncie
set protocols iccp peer 10.21.12.1 redundancy-group-id-list 10
set protocols iccp peer 10.21.12.1 liveness-detection minimum-interval 3000
set chassis aggregated-devices ethernet device-count 5
set multi-chassis multi-chassis-protection 10.21.12.1 interface ae2
set switch-options service-id 1
```

Task 2.3: Layer 2 Underlay Features

- Log a message if there are more than 10 mac-flaps in VLAN 121 on vQFX3.
- Limit the total number of mac-addresses for VLAN 121 on VQFX1 to 256. Drop packets if the threshold is exceeded.

vQFX3:

```
set vlans vl10 switch-options mac-move-limit 10
set vlans vl10 switch-options mac-move-limit packet-action log
```

vQFX1:

```
set vlans jncielag switch-options mac-table-size 256
set vlans jncielag switch-options mac-table-size packet-action drop
set vlans jncielag switch-options interface-mac-limit 256
```

Task 2.4: BUM Traffic Handling

- Ensure that only Unknown Unicast traffic is limited to 150 kbps on vQFX2.
- Do not use the default profile.
- Enable this filter for all existing 10G layer 2 interfaces and automatically for 10G Layer 2 interfaces that are added in the future.

vQFX2:

```
set groups stormcontrol interfaces <xe-*> unit <*> family ethernet-switching storm-control
jnciedc
set interfaces apply-groups stormcontrol
set interfaces xe-0/0/0 apply-groups-except stormcontrol
set forwarding-options storm-control-profiles jnciedc all bandwidth-level 150
set forwarding-options storm-control-profiles jnciedc all no-broadcast
set forwarding-options storm-control-profiles jnciedc all no-multicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.5: Layer 2 Redundancy

Ensure that AE10 towards Server 3 keeps forwarding traffic for VLAN10, 20 and 30 if Server3 cannot exchange LACP PDUs during a server upgrade.

vQFX3, VQFX4, vQFX5:

```
set interfaces ae10 aggregated-ether-options lacp force-up
```

Task 2.6: Layer 3 Underlay MAIN DC

Configure interface ge-0/0/0 on vSRX for VLAN 50, VLAN 51 and VLAN 52. Use the last possible address within the /24 subnet as the default gateway address for each VLAN.

vSRX1:

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 50 vlan-id 50
set interfaces ge-0/0/0 unit 50 family inet address 10.150.50.254/24
set interfaces ge-0/0/0 unit 51 vlan-id 51
set interfaces ge-0/0/0 unit 51 family inet address 10.150.51.254/24
set interfaces ge-0/0/0 unit 52 vlan-id 52
set interfaces ge-0/0/0 unit 52 family inet address 10.150.52.254/24
```

Task 2.7: Layer 3 Underlay Routing DC1

Configure OSPFv2 between vMX1, vMX2, vQFX2 according to the following requirements:

- All interfaces should be configured in the backbone area.
- Ensure that OSPFv2 adjacencies between vQFX2, vMX1 and vMX2 are in the “full” state.

vMX1 & vMX2:

```
delete routing-options router-id
set protocols ospf area 0.0.0.0 interface irb.121
```

vQFX2:

```
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0
```

NOTE: This task requires that the MC-LAG between vMX1, vMX2 is configured correctly.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.8: Layer 2 Underlay Switching vQFX1

Ensure the Ethernet switching table of vQFX1 has exactly the same output as shown below:

```
Ethernet switching table : 4 entries, 4 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical                NH
RTR
    name                address             flags              interface              Index
ID
    jncielag            00:00:5e:00:01:fe   D                -   ae0.0                  0
0
    jncielag            00:05:01:01:01:01   D                -   ae0.0                  0
0
    jncielag            00:05:02:02:02:02   D                -   ae0.0                  0
0
    jncielag            00:05:03:03:03:03   D                -   xe-0/0/0.0             0
0
```

vMX1:

```
set interfaces irb unit 121 mac 00:05:01:01:01:01
rename interfaces irb.121 family inet address 10.20.121.1/24 vrrp-group 1 to vrrp-group 254
delete interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.1/24 arp 10.20.121.2 mac
00:05:02:02:02:02
```

vMX2:

```
set interfaces irb unit 121 mac 00:05:02:02:02:02
rename interfaces irb.121 family inet address 10.20.121.2/24 vrrp-group 1 to vrrp-group 254
delete interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 l2-interface
ae2.0
set interfaces irb unit 121 family inet address 10.20.121.2/24 arp 10.20.121.1 mac
00:05:01:01:01:01
```

vQFX2:

```
set interfaces xe-0/0/0 mac 00:05:03:03:03:03
```

vQFX1:

```
clear ethernet-switching table
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 2.9: EBGP Underlay DC2

- Configure the IP fabric in DC1 and DC2 according to the following requirements:

Device	BGP AS
vMX1	64512
vMX2	64512
vMX3	64523
vMX4	64524
vQFX3	64513
vQFX4	64514
vQFX5	64515

Configure BFD on all EBGP sessions for all devices in DC2s IP fabric.

- Ensure that BFD sessions do no support BGP sessions that peer on loopback interfaces.
- Enable BFD MD5 authentication with password “jncie”.
- Ensure that traffic is load-balanced if there are equal cost paths towards the same destination.

vMX3:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64523
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.33.2 peer-as 64513
set protocols bgp group underlay neighbor 10.30.34.2 peer-as 64514
set protocols bgp group underlay neighbor 10.30.35.2 peer-as 64515
set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet
set routing-options forwarding-table export lbpp
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vMX4:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64524
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.37.2 peer-as 64513
set protocols bgp group underlay neighbor 10.30.38.2 peer-as 64514
set protocols bgp group underlay neighbor 10.30.36.2 peer-as 64515
set protocols bgp group underlay neighbor 10.30.4.1 peer-as 64513

set security authentication-key-chains key-chain underlay key 0 secret jncie
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet
set routing-options forwarding-table export lbpp
```

vQFX3:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64513
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.33.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.37.1 peer-as 64524
set security authentication-key-chains key-chain underlay key 0 secret "$9$yWneMXaJDkqfVwoG"
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet

set routing-options forwarding-table export lbpp
```

vQFX4:

```
delete interfaces xe-0/0/2
delete interfaces xe-0/0/4
set interfaces xe-0/0/2 unit 0 family inet address 10.30.34.2/30
set interfaces xe-0/0/4 unit 0 family inet address 10.30.38.2/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64514
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.34.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.38.1 peer-as 64524

set security authentication-key-chains key-chain underlay key 0 secret "$9$yWneMXaJDkqfVwoG"
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet

set routing-options forwarding-table export lbpp
```

vQFX5:

```
set protocols bgp group underlay type external
set protocols bgp group underlay local-as 64515
set protocols bgp group underlay bfd-liveness-detection minimum-interval 3000
set protocols bgp group underlay bfd-liveness-detection multiplier 3
set protocols bgp group underlay bfd-liveness-detection authentication key-chain underlay
set protocols bgp group underlay bfd-liveness-detection authentication algorithm keyed-md5
set protocols bgp group underlay bfd-liveness-detection session-mode single-hop
set protocols bgp group underlay multipath multiple-as
set protocols bgp group underlay neighbor 10.30.35.1 peer-as 64523
set protocols bgp group underlay neighbor 10.30.36.1 peer-as 64524

set security authentication-key-chains key-chain underlay key 0 secret "$9$yWneMXaJDkqfVwoG"
set security authentication-key-chains key-chain underlay key 0 start-time "2017-1-1.00:00:00
+0000"

set policy-options policy-statement lbpp then load-balance per-packet

set routing-options forwarding-table export lbpp
```

Task 3: Controllerless Overlay

Task 3.1: EVPN Overlay Configuration

Configure an EVPN overlay between DC1, DC2 and MAIN DC according to the following requirements:

- Use BGP AS: 65535

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

- vQFX2, vQFX3, vQFX4, vQFX5 and vQFX6 should only have EVPN overlay sessions with vMX3 and vMX4.
- vMX3 and vMX4 should **not** initiate BGP sessions used for the EVPN overlay. Do not use the "passive" command or a firewall filter to accomplish this task.

NOTE: You must complete the DCI tasks before the overlay sessions between DC1, DC2 and MAIN DC can be established.

vMX3:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.33
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 192.168.1.33
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay allow 192.168.1.0/29

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vMX4:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.44
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay cluster 192.168.1.44
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay allow 192.168.1.0/29

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX2:

```
set routing-options router-id 192.168.1.2
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.2
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set protocols ospf area 0.0.0.0 interface lo0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.3
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
```

vQFX4:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.4
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept

rename interfaces lo0.0 family inet address 192.168.1.3/32 to address 192.168.1.4/32
```

vQFX5:

```
set protocols bgp group underlay export myloopback
set protocols bgp group overlay type internal
set protocols bgp group overlay local-address 192.168.1.5
set protocols bgp group overlay family evpn signaling
set protocols bgp group overlay local-as 65535
set protocols bgp group overlay multipath
set protocols bgp group overlay neighbor 192.168.1.33
set protocols bgp group overlay neighbor 192.168.1.44

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept
rename interfaces lo0.0 family inet address 192.168.1.3/32 to address 192.168.1.5/32
```

Task 3.2: EVPN/VLAN Service Configuration

- Configure an EVPN service called “vpn1” for VLAN 10, VLAN 20 and VLAN30.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

EVPN/VLAN service	VLAN	vMX3 gateway address	vMX4 gateway address
vpn1	10	10.150.10.100 10.150.10.254	10.150.10.101 10.150.10.254
	20	10.150.20.100 10.150.20.254	10.150.20.101 10.150.20.254
	30	10.150.30.100 10.150.30.254	10.150.30.101 10.150.30.254

Ensure that the following communities are used for each EVPN/VLAN service as shown in the following table:

EVPN/VXLAN service	Community name	Community value
vpn1	evpn	1:1
	vpn1	1:1000
	vni-10	1:10
	vni-20	1:20
	vni-30	1:30

vMX3:

```
set routing-instances vpn1 route-distinguisher 192.168.1.33:1

set interfaces irb unit 10 family inet address 10.150.10.100/24 virtual-gateway-address
10.150.10.254
set interfaces irb unit 20 family inet address 10.150.20.100/24 virtual-gateway-address
10.150.20.254
set interfaces irb unit 30 family inet address 10.150.30.100/24 virtual-gateway-address
10.150.30.254
```

vMX4:

```
set routing-instances vpn1 route-distinguisher 192.168.1.44:1

set interfaces irb unit 10 family inet address 10.150.10.101/24 virtual-gateway-address
10.150.10.254
set interfaces irb unit 20 family inet address 10.150.20.101/24 virtual-gateway-address
10.150.20.254
set interfaces irb unit 30 family inet address 10.150.30.101/24 virtual-gateway-address
10.150.30.254
```

vMX3 and vMX4:

```
set routing-instances vpn1 vtep-source-interface lo0.0
set routing-instances vpn1 instance-type virtual-switch
set routing-instances vpn1 vrf-import vpn1-import
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set routing-instances vpn1 vrf-target target:1:1000
set routing-instances vpn1 protocols evpn encapsulation vxlan
set routing-instances vpn1 protocols evpn extended-vni-list 10
set routing-instances vpn1 protocols evpn extended-vni-list 20
set routing-instances vpn1 protocols evpn extended-vni-list 30
set routing-instances vpn1 protocols evpn vni-options vni 10 vrf-target target:1:10
set routing-instances vpn1 protocols evpn vni-options vni 20 vrf-target target:1:20
set routing-instances vpn1 protocols evpn vni-options vni 30 vrf-target target:1:30
set routing-instances vpn1 protocols evpn multicast-mode ingress-replication
set routing-instances vpn1 bridge-domains ie-10 domain-type bridge
set routing-instances vpn1 bridge-domains ie-10 vlan-id 10
set routing-instances vpn1 bridge-domains ie-10 routing-interface irb.10
set routing-instances vpn1 bridge-domains ie-10 vxlan vni 10
set routing-instances vpn1 bridge-domains ie-10 vxlan ingress-node-replication
set routing-instances vpn1 bridge-domains ie-20 domain-type bridge
set routing-instances vpn1 bridge-domains ie-20 vlan-id 20
set routing-instances vpn1 bridge-domains ie-20 routing-interface irb.20
set routing-instances vpn1 bridge-domains ie-20 vxlan vni 20
set routing-instances vpn1 bridge-domains ie-20 vxlan ingress-node-replication
set routing-instances vpn1 bridge-domains ie-30 domain-type bridge
set routing-instances vpn1 bridge-domains ie-30 vlan-id 30
set routing-instances vpn1 bridge-domains ie-30 routing-interface irb.30
set routing-instances vpn1 bridge-domains ie-30 vxlan vni 30
set routing-instances vpn1 bridge-domains ie-30 vxlan ingress-node-replication

set policy-options policy-statement vpn1-import term 1 from community evpn
set policy-options policy-statement vpn1-import term 1 then accept
set policy-options policy-statement vpn1-import term 2 from community vpn1
set policy-options policy-statement vpn1-import term 2 then accept
set policy-options policy-statement vpn1-import term 3 from community vni-10
set policy-options policy-statement vpn1-import term 3 then accept
set policy-options policy-statement vpn1-import term 4 from community vni-20
set policy-options policy-statement vpn1-import term 4 from community vni-10
set policy-options policy-statement vpn1-import term 4 then accept
set policy-options policy-statement vpn1-import term 5 from community vni-20
set policy-options policy-statement vpn1-import term 5 then accept
set policy-options policy-statement vpn1-import term 6 from community vni-30
set policy-options policy-statement vpn1-import term 6 then accept
set policy-options community evpn members target:1:1
set policy-options community vni-10 members target:1:10
set policy-options community vni-20 members target:1:20
set policy-options community vni-30 members target:1:30
set policy-options community vpn1 members target:1:1000
```

vQFX2:

```
set switch-options route-distinguisher 192.168.1.2:1
set routing-options router-id 192.168.1.2
set interfaces lo0.0 family inet address 192.168.1.2/32 primary preferred
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

vQFX3:

```
set switch-options route-distinguisher 192.168.1.3:1
set routing-options router-id 192.168.1.3
```

vQFX4:

```
set switch-options route-distinguisher 192.168.1.4:1
set routing-options router-id 192.168.1.4
```

vQFX5:

```
set switch-options route-distinguisher 192.168.1.5:1
set routing-options router-id 192.168.1.5
```

vQFX2, vQFX3, vQFX4, vQFX5:

```
set protocols evpn encapsulation vxlan
set protocols evpn extended-vni-list 10
set protocols evpn extended-vni-list 20
set protocols evpn extended-vni-list 30
set protocols evpn multicast-mode ingress-replication
set protocols evpn vni-options vni 10 vrf-target export target:1:10
set protocols evpn vni-options vni 20 vrf-target export target:1:20
set protocols evpn vni-options vni 30 vrf-target export target:1:30

set switch-options vtep-source-interface lo0.0
set switch-options vrf-import vpn-import
set switch-options vrf-target target:1:1

set vlans vl10 vxlan vni 10
set vlans vl10 vxlan ingress-node-replication
set vlans vl20 vxlan vni 20
set vlans vl20 vxlan ingress-node-replication
set vlans vl30 vxlan vni 30
set vlans vl30 vxlan ingress-node-replication

set interfaces ae10 esi 00:11:11:11:11:11:11:11:22
set interfaces ae10 esi all-active

set policy-options policy-statement vpn-import term 1 from community evpn
set policy-options policy-statement vpn-import term 1 then accept
set policy-options policy-statement vpn-import term 2 from community vpn1
set policy-options policy-statement vpn-import term 2 then accept
set policy-options policy-statement vpn-import term 3 from community vni-10
set policy-options policy-statement vpn-import term 3 then accept
set policy-options policy-statement vpn-import term 4 from community vni-20
set policy-options policy-statement vpn-import term 4 then accept
set policy-options policy-statement vpn-import term 5 from community vni-30
set policy-options policy-statement vpn-import term 5 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set policy-options community evpn members target:1:1
set policy-options community vni-10 members target:1:10
set policy-options community vni-20 members target:1:20
set policy-options community vni-30 members target:1:30
set policy-options community vpn1 members target:1:1000
```

Task 3.3: EVPN/VXLAN Service Tuning

- Configure vQFX2 to age out the address of the remote VTEP after 320 seconds, if all the MAC addresses learned from that VTEP age out. Your solution must also account for VNIs that might be added in the future.

vQFX2:

```
set groups unreachable vlans <*> vxlan unreachable-vtep-aging-timer 320

set vlans apply-groups unreachable
set vlans default apply-groups-except unreachable
```

Task 3.4: EVPN signaling

- Ensure that vQFX2, vQFX3, vQFX4 and vQFX5 receive the following EVPN type 2 routes:

2:192.168.1.33:1::**30**::00:99:00:00:00:**10**/304
2:192.168.1.33:1::20::00:99:00:00:00:20/304
2:192.168.1.33:1::**10**::00:99:00:00:00:**30**/304

vMX3 and vMX4:

```
set interfaces irb unit 10 mac 00:99:00:00:00:30
set interfaces irb unit 20 mac 00:99:00:00:00:20
set interfaces irb unit 30 mac 00:99:00:00:00:10
```

Task 4: Data Center Interconnect

Task 4.1: DCI Configuration

Configure two L3VPN services according to the information:

- The first L3VPN service should be named “underlay” on vMX5 and is used for the underlay between DC1 and DC2.

The following table shows the devices and interface configurations needed for the DCI. Some of the interfaces are already preconfigured.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

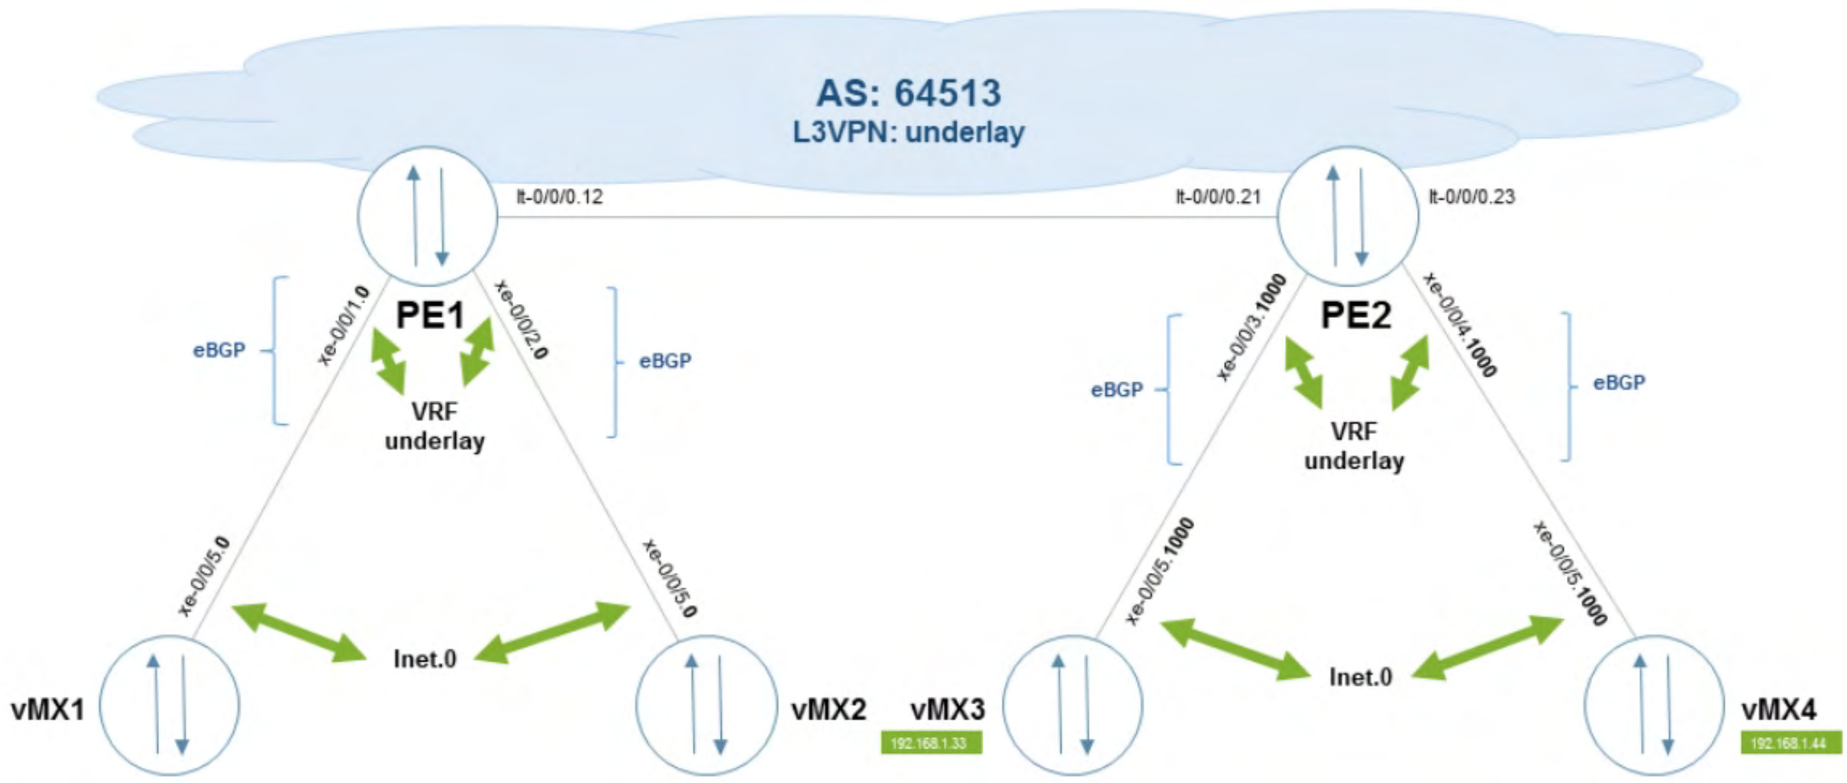

vMX5 (PE1):

```
set logical-systems PE1 protocols mpls interface all
set logical-systems PE1 protocols bgp group l3vpn type internal
set logical-systems PE1 protocols bgp group l3vpn local-address 172.17.1.1
set logical-systems PE1 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE1 protocols bgp group l3vpn local-as 64513
set logical-systems PE1 protocols bgp group l3vpn neighbor 172.17.1.2
set logical-systems PE1 protocols bgp group l3vpn neighbor 172.17.1.3
set logical-systems PE1 protocols ospf area 0.0.0.0 interface lt-0/0/0.12 hello-interval 1
set logical-systems PE1 protocols ospf area 0.0.0.0 interface lt-0/0/0.12 dead-interval 4
set logical-systems PE1 protocols ospf area 0.0.0.0 interface lo0.1
set logical-systems PE1 protocols ldp interface all
set logical-systems PE1 policy-options policy-statement lbpp then load-balance per-packet
set logical-systems PE1 policy-options policy-statement lbpp then accept
set logical-systems PE1 routing-instances underlay instance-type vrf
set logical-systems PE1 routing-instances underlay interface xe-0/0/1.0
set logical-systems PE1 routing-instances underlay interface xe-0/0/2.0
set logical-systems PE1 routing-instances underlay route-distinguisher 1:1
set logical-systems PE1 routing-instances underlay vrf-target target:64513:1
set logical-systems PE1 routing-instances underlay vrf-table-label
set logical-systems PE1 routing-instances underlay protocols bgp group dc1 type external
set logical-systems PE1 routing-instances underlay protocols bgp group dc1 peer-as 64512
set logical-systems PE1 routing-instances underlay protocols bgp group dc1 local-as 64513
set logical-systems PE1 routing-instances underlay protocols bgp group dc1 neighbor 10.20.1.2
set logical-systems PE1 routing-instances underlay protocols bgp group dc1 neighbor 10.20.2.2
set logical-systems PE1 routing-options forwarding-table export lbpp
```

vMX5 (PE2):

```
set logical-systems PE2 protocols mpls interface all
set logical-systems PE2 protocols bgp group l3vpn type internal
set logical-systems PE2 protocols bgp group l3vpn local-address 172.17.1.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set logical-systems PE2 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE2 protocols bgp group l3vpn local-as 64513
set logical-systems PE2 protocols bgp group l3vpn neighbor 172.17.1.1
set logical-systems PE2 protocols bgp group l3vpn neighbor 172.17.1.3

set logical-systems PE2 protocols ospf area 0.0.0.0 interface all hello-interval 1
set logical-systems PE2 protocols ospf area 0.0.0.0 interface all dead-interval 4

set logical-systems PE2 protocols ldp interface all

set logical-systems PE1 protocols mpls interface all

set logical-systems PE2 policy-options policy-statement lbpp then load-balance per-packet
set logical-systems PE2 policy-options policy-statement lbpp then accept

set logical-systems PE2 routing-instances underlay instance-type vrf
set logical-systems PE2 routing-instances underlay interface xe-0/0/3.1000
set logical-systems PE2 routing-instances underlay interface xe-0/0/4.1000
set logical-systems PE2 routing-instances underlay route-distinguisher 1:2
set logical-systems PE2 routing-instances underlay vrf-target target:64513:1
set logical-systems PE2 routing-instances underlay vrf-table-label
set logical-systems PE2 routing-instances underlay protocols bgp group underlay type external
set logical-systems PE2 routing-instances underlay protocols bgp group underlay local-as 64513
set logical-systems PE2 routing-instances underlay protocols bgp group underlay neighbor
10.30.3.2 peer-as 64523
set logical-systems PE2 routing-instances underlay protocols bgp group underlay neighbor
10.30.4.2 peer-as 64524

set logical-systems PE2 routing-options forwarding-table export lbpp
```

vMX5 (PE3):

```
set logical-systems PE3 protocols mpls interface all
set logical-systems PE3 protocols bgp group l3vpn type internal
set logical-systems PE3 protocols bgp group l3vpn local-address 172.17.1.3
set logical-systems PE3 protocols bgp group l3vpn family inet-vpn unicast
set logical-systems PE3 protocols bgp group l3vpn local-as 64513
set logical-systems PE3 protocols bgp group l3vpn neighbor 172.17.1.2
set logical-systems PE3 protocols bgp group l3vpn neighbor 172.17.1.1
set logical-systems PE3 protocols ospf area 0.0.0.0 interface all hello-interval 1
set logical-systems PE3 protocols ospf area 0.0.0.0 interface all dead-interval 4
set logical-systems PE3 protocols ldp interface all
```

vMX1:

```
set protocols bgp group dci neighbor 10.20.1.1 peer-as 64513
set protocols bgp group dci neighbor 10.20.1.1 local-as 64512
```


© 2020 Juniper Networks, Inc. All rights reserved.

vMX2:

```
set protocols bgp group dci neighbor 10.20.2.1 peer-as 64513
set protocols bgp group dci neighbor 10.20.2.1 local-as 64512
```

vMX3:

```
set protocols bgp group underlay neighbor 10.30.3.1 peer-as 64513
set protocols bgp group underlay neighbor 10.30.3.1 local-as 64523
set protocols bgp group underlay neighbor 10.30.3.1 as-override
set protocols bgp group underlay export myloopback

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept

set protocols bgp group underlay neighbor 10.30.33.2 as-override
```

vMX4:

```
set protocols bgp group underlay neighbor 10.30.4.1 peer-as 64513
set protocols bgp group underlay neighbor 10.30.4.1 as-override
set protocols bgp group underlay export myloopback

set policy-options policy-statement myloopback term 1 from interface lo0.0
set policy-options policy-statement myloopback term 1 then accept

set protocols bgp group underlay neighbor 10.30.37.2 as-override
```

- The second L3VPN should be named “dc-vpns” on vMX5 and is used for L3 routing for EVPN service “vpn1” between DC2 and MAIN DC.

© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

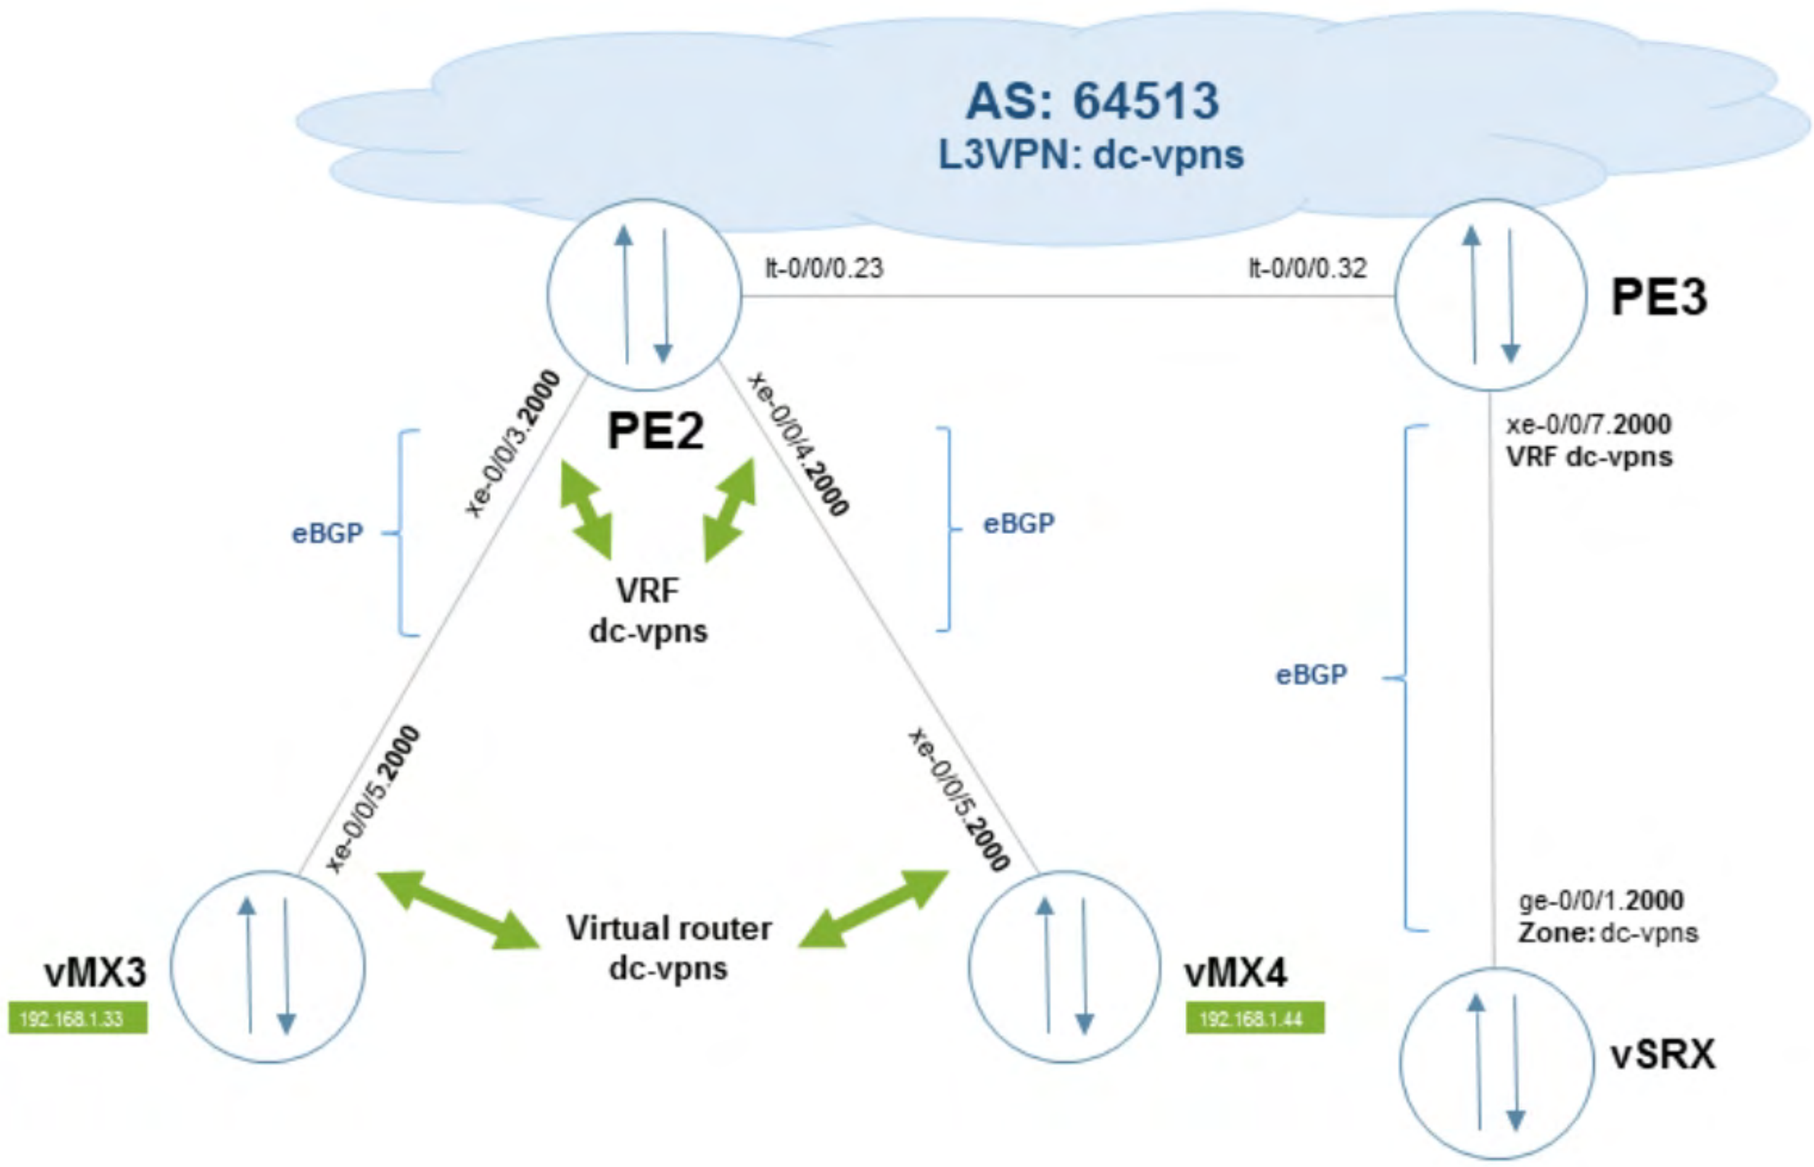

vMX5 (PE2):

```
set logical-systems PE2 routing-instances dc-vpns instance-type vrf
set logical-systems PE2 routing-instances dc-vpns interface xe-0/0/3.2000
set logical-systems PE2 routing-instances dc-vpns interface xe-0/0/4.2000
set logical-systems PE2 routing-instances dc-vpns route-distinguisher 2:2
set logical-systems PE2 routing-instances dc-vpns vrf-target target:64513:2
set logical-systems PE2 routing-instances dc-vpns vrf-table-label
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns type external
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns local-as 64513
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns neighbor
10.30.3.6 peer-as 64523
set logical-systems PE2 routing-instances dc-vpns protocols bgp group dc-vpns neighbor
10.30.4.6 peer-as 64524
```

vMX5 (PE3):

```
set logical-systems PE3 routing-instances dc-vpns instance-type vrf
set logical-systems PE3 routing-instances dc-vpns interface xe-0/0/7.2000
set logical-systems PE3 routing-instances dc-vpns route-distinguisher 2:3
set logical-systems PE3 routing-instances dc-vpns vrf-target target:64513:2
set logical-systems PE3 routing-instances dc-vpns vrf-table-label
set logical-systems PE3 routing-instances dc-vpns protocols bgp group external type external
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set logical-systems PE3 routing-instances dc-vpns protocols bgp group external neighbor
10.40.2.2 peer-as 64516
set logical-systems PE3 routing-options autonomous-system 64513
```

vMX3:

```
set routing-instances dc-vpns instance-type virtual-router
set routing-instances dc-vpns interface xe-0/0/5.2000
set routing-instances dc-vpns interface irb.10
set routing-instances dc-vpns interface irb.20
set routing-instances dc-vpns interface irb.30
set routing-instances dc-vpns protocols bgp group dc-vpns type external
set routing-instances dc-vpns protocols bgp group dc-vpns export direct-2-bgp
set routing-instances dc-vpns protocols bgp group dc-vpns local-as 64523
set routing-instances dc-vpns protocols bgp group dc-vpns neighbor 10.30.3.5 peer-as 64513
set policy-options policy-statement direct-2-bgp term 1 from protocol direct
set policy-options policy-statement direct-2-bgp term 1 then accept
```

vMX4:

```
set routing-instances dc-vpns instance-type virtual-router
set routing-instances dc-vpns interface xe-0/0/5.2000
set routing-instances dc-vpns interface irb.10
set routing-instances dc-vpns interface irb.20
set routing-instances dc-vpns interface irb.30
set routing-instances dc-vpns protocols bgp group dc-vpns type external
set routing-instances dc-vpns protocols bgp group dc-vpns export direct-2-bgp
set routing-instances dc-vpns protocols bgp group dc-vpns local-as 64524
set routing-instances dc-vpns protocols bgp group dc-vpns neighbor 10.30.4.5 peer-as 64513
set policy-options policy-statement direct-2-bgp term 1 from protocol direct
set policy-options policy-statement direct-2-bgp term 1 then accept
```

vSRX:

```
set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1

set protocols bgp export 2bgp
set protocols bgp group external type external
set protocols bgp group external neighbor 10.40.2.1 peer-as 64513

set policy-options policy-statement 2bgp from route-filter 0.0.0.0/0 exact
set policy-options policy-statement 2bgp then accept

set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic system-
services ping
set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic
protocols bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.2: DC1 Redundancy

vMX1 and vMX2 use BGP AS 64512 to exchange routes with PE1. Ensure that vMX1 and vMX2 always prefers the shortest path towards PE1. Ensure that vMX1 and vMX2 redistribute DCI underlay routes into OSPFv2. Ensure that the vQFX2 VTEP interface is learned in DC2.

vMX1 and vMX2:

```
set policy-options policy-statement bgp2ospf term 1 from protocol bgp
set policy-options policy-statement bgp2ospf term 1 from external
set policy-options policy-statement bgp2ospf term 1 then accept

set policy-options policy-statement ospf2bgp term 1 from protocol ospf
set policy-options policy-statement ospf2bgp term 1 from route-filter 192.168.1.2/32 exact
set policy-options policy-statement ospf2bgp term 1 then accept

set protocols bgp group dci export ospf2bgp
set protocols ospf export bgp2ospf
set protocols ospf external-preference 180
```

Task 4.3: (Inter) Data Center Connectivity

- Ensure that the Servers in DC1, DC2 and MAIN DC have full connectivity between all VLANs.

Use the test addresses shown in the topology for verification purposes.

vSRX:

```
set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic system-
services ping
set security zones security-zone dc-vpns interfaces ge-0/0/1.2000 host-inbound-traffic
protocols bgp
set security zones security-zone servers host-inbound-traffic system-services ping
set security zones security-zone servers interfaces ge-0/0/0.50
set security zones security-zone servers interfaces ge-0/0/0.51
set security zones security-zone servers interfaces ge-0/0/0.52
set security zones security-zone untrust interfaces ge-0/0/1.3000
set security policies from-zone dc-vpns to-zone untrust policy 1 match source-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match destination-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match application junos-icmp-
ping
set security policies from-zone dc-vpns to-zone untrust policy 1 then permit
set security policies from-zone dc-vpns to-zone servers policy 1 match source-address any
set security policies from-zone dc-vpns to-zone servers policy 1 match destination-address any
set security policies from-zone dc-vpns to-zone servers policy 1 match application junos-icmp-
ping
set security policies from-zone dc-vpns to-zone servers policy 1 then permit
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set security policies from-zone servers to-zone servers policy 1 match source-address any
set security policies from-zone servers to-zone servers policy 1 match destination-address any
set security policies from-zone servers to-zone servers policy 1 match application junos-icmp-
ping
set security policies from-zone servers to-zone servers policy 1 then permit
set security policies from-zone servers to-zone dc-vpns policy 1 match source-address any
set security policies from-zone servers to-zone dc-vpns policy 1 match destination-address any
set security policies from-zone servers to-zone dc-vpns policy 1 match application junos-icmp-
ping
set security policies from-zone servers to-zone dc-vpns policy 1 then permit
set security policies from-zone servers to-zone untrust policy 1 match source-address any
set security policies from-zone servers to-zone untrust policy 1 match destination-address any
set security policies from-zone servers to-zone untrust policy 1 match application junos-icmp-
ping
set security policies from-zone servers to-zone untrust policy 1 then permit
```

Task 4.4: Data Center Connectivity to External Server

Ensure that the Servers DC2 and MAIN DC can ping the external server (8.8.8.8). The external server is preconfigured on vMX5.

- Traffic between the servers in DC2, MAIN DC and the external server should always transit the “dc-vpns” and “untrust” zones of the vSRX firewall.
- The external server should use destination IP address 1.1.1.2 if it sends ICMP echo replies to Servers in DC2 and MAIN DC.

vSRX:

```
set protocols bgp group external export default-2-bgp

set policy-options policy-statement default-2-bgp from protocol static
set policy-options policy-statement default-2-bgp from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default-2-bgp then accept

set routing-options static route 0.0.0.0/0 next-hop 1.1.1.1

set security policies from-zone dc-vpns to-zone untrust policy 1 match source-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match destination-address any
set security policies from-zone dc-vpns to-zone untrust policy 1 match application junos-icmp-
ping
set security policies from-zone dc-vpns to-zone untrust policy 1 then permit

set security nat source rule-set jncie from zone dc-vpns
set security nat source rule-set jncie from zone servers
set security nat source rule-set jncie to zone untrust
set security nat source rule-set jncie rule jncie-dc match source-address 0.0.0.0/0
set security nat source rule-set jncie rule jncie-dc match destination-address 0.0.0.0/0
set security nat source rule-set jncie rule jncie-dc match application any
set security nat source rule-set jncie rule jncie-dc then source-nat interface
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Task 4.5: EVPN VXLAN / EVPN MPLS Interworking

Prepare vMX4 for a future EVPN MPLS DCI connection:

- Create an EVPN MPLS instance called “evpn”.
- Interface xe-0/0/5.10 should be included in the EVPN MPLS instance.
- Ensure all servers in DC2 are able to ping TEST SERVER.

```
set chassis fpc 0 pic 0 tunnel-services
set interfaces xe-0/0/5 unit 10 family bridge interface-mode trunk
set interfaces xe-0/0/5 unit 10 family bridge vlan-id-list 10
set protocols mpls interface all
set routing-instances evpn instance-type virtual-switch
set routing-instances evpn interface lt-0/0/0.0
set routing-instances evpn interface xe-0/0/5.10
set routing-instances evpn route-distinguisher 192.168.1.44:2
set routing-instances evpn vrf-target target:1:2
set routing-instances evpn protocols evpn extended-vlan-list 10
set routing-instances evpn bridge-domains ie-10 vlan-id 10

set routing-instances vpn1 interface lt-0/0/0.1

set interfaces lt-0/0/0 unit 0 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 0 peer-unit 1
set interfaces lt-0/0/0 unit 0 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 0 family bridge vlan-id-list 10
set interfaces lt-0/0/0 unit 1 encapsulation ethernet-bridge
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family bridge interface-mode trunk
set interfaces lt-0/0/0 unit 1 family bridge vlan-id-list 10
```

Task 5: Security

Task 5.1: User Authentication and Authorization

Configure a user called “jncie-1” on vSRX with password “jncie123”. Helpful information should be shown when this user logs in to the device. User “jncie-1” should only be allowed to display the device configuration and view system and interfaces statistics.

Configure a user called “jncie-2” on vSRX with password “jncie123” and assign the operator role.

vSRX:

```
set system login class jncie login-tip
set system login class jncie permissions view
set system login class jncie permissions view-configuration
set system login user jncie-1 uid 2000
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set system login user jncie-1 class operator
set system login user jncie-1 authentication encrypted-password
"$1$zE4Ws2yJ$9uqNu2rzROynltuFASFBe0"
set system login user jncie-2 uid 2002
set system login user jncie-2 class operator
set system login user jncie-2 authentication encrypted-password
"$1$gygaN5Vq$zC.tqY4CfG7DmCajP8jBY"
```

Task 5.2: SRX Configuration

Configure the vSRX according to the following requirements:

See answers in DCI underlay and dc-vpn sections

Configure the vSRX according to the following requirements:

- Interface ge-0/0/1.2000 should be in the “dc-vpns” zone.
- Interface ge-0/0/1.3000 should be in the “untrust” zone.
- Configure a static default route towards the next-hop IP address (1.1.1.1) of AS 64513 in the untrust zone.
- Allow ICMP PING traffic towards vSRX interfaces, except for interfaces in the “untrust” zone.
- Allow ICMP traffic between all zones.
- Do not configure routing-instances on vSRX.
- Only allow traffic if required to complete other tasks in this superlab. Do not allow other traffic.

Task 5.3: RE Protection

Configure control-plane protection on vQFX2 with the following requirements:

- Allow incoming telnet and ssh sessions from all destinations.
- Allow BGP sessions to vMX3 and vMX4. Make sure that if new BGP neighbors are configured in the “protocols bgp” stanza, you do not need to change your RE protection configuration.
- Allow ICMP from all destinations.
- Allow specific traffic needed by other tasks in this superlab.
- Deny, count and log all other traffic. The control plane should not respond to denied traffic.

vQFX2:

```
set interfaces lo0 unit 0 family inet filter input re-protect

set policy-options prefix-list bgp apply-path "protocols bgp group <*> neighbor <*>"
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set firewall family inet filter re-protect term 1 from source-prefix-list bgp
set firewall family inet filter re-protect term 1 from protocol tcp
set firewall family inet filter re-protect term 1 from source-port bgp
set firewall family inet filter re-protect term 1 then accept
set firewall family inet filter re-protect term 2 from protocol icmp
set firewall family inet filter re-protect term 2 then accept
set firewall family inet filter re-protect term 3 from protocol tcp
set firewall family inet filter re-protect term 3 from destination-port ssh
set firewall family inet filter re-protect term 3 from destination-port telnet
set firewall family inet filter re-protect term 3 then accept
set firewall family inet filter re-protect term 4 from source-address 10.20.121.0/24
set firewall family inet filter re-protect term 4 from protocol ospf
set firewall family inet filter re-protect term 4 then accept
set firewall family inet filter re-protect term 5 from source-address 192.168.1.0/24
set firewall family inet filter re-protect term 5 from protocol udp
set firewall family inet filter re-protect term 5 from port 4789
set firewall family inet filter re-protect term 5 then count rejected
set firewall family inet filter re-protect term 5 then log
set firewall family inet filter re-protect term 5 then discard
```

Task 6: Class of Service

Task 6.1: Forwarding Classes

Configure the following forwarding classes and forwarding class sets on all vQFX devices

Forwarding Class	Queue ID	No-Loss
BE	0	Disabled
FCOE	3	Enabled
NOLOSS	4	Enabled
NC	7	Disabled

Forwarding Class Set	Forwarding Class
LOSS	BE
	NC
NOLOSS	FCOE
	NOLOSS

vQFX devices:

```
set class-of-service forwarding-classes class BE queue-num 0
set class-of-service forwarding-classes class FCOE queue-num 3
set class-of-service forwarding-classes class FCOE no-loss
set class-of-service forwarding-classes class NOLOSS queue-num 4
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

```
set class-of-service forwarding-classes class NOLOSS no-loss
set class-of-service forwarding-classes class NC queue-num 7

set class-of-service forwarding-class-sets LOSS class BE
set class-of-service forwarding-class-sets LOSS class NC
set class-of-service forwarding-class-sets NOLOSS class FCOE
set class-of-service forwarding-class-sets NOLOSS class NOLOSS
```

Task 6.2: Schedulers and Traffic Profiles

- Configure a scheduler mapping named “LOSS” that includes the BE and NC schedulers on all vQFX devices.
- Configure a scheduler mapping named “NOLOSS” that includes the FCOE and NOLOSS schedulers on all vQFX devices.
- Configure traffic profiles for each scheduler.

Scheduler name	Transmit rate	Buffer size	Priority
BE	55 %	remainder	low
FCOE	25 %	25 %	low
NOLOSS	19 %	19 %	low
NC	1 %	1 %	low

vQFX devices:

```
set class-of-service schedulers BE transmit-rate percent 55
set class-of-service schedulers BE buffer-size remainder
set class-of-service schedulers BE priority low

set class-of-service schedulers FCOE transmit-rate percent 25
set class-of-service schedulers FCOE buffer-size percent 29
set class-of-service schedulers FCOE priority low

set class-of-service schedulers NOLOSS transmit-rate percent 19
set class-of-service schedulers NOLOSS buffer-size percent 19
set class-of-service schedulers NOLOSS priority low

set class-of-service schedulers NC transmit-rate percent 1
set class-of-service schedulers NC buffer-size percent 1
set class-of-service schedulers NC priority low

set class-of-service scheduler-maps NOLOSS forwarding-class NOLOSS scheduler NOLOSS
set class-of-service scheduler-maps NOLOSS forwarding-class FCOE scheduler FCOE
set class-of-service scheduler-maps LOSS forwarding-class BE scheduler BE
set class-of-service scheduler-maps LOSS forwarding-class NC scheduler NC

set class-of-service traffic-control-profiles NOLOSS scheduler-map NOLOSS
set class-of-service traffic-control-profiles LOSS scheduler-map LOSS
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only

Note: vQFX does not support the usage of forwarding-class-sets and traffic profiles on xe-interfaces. For this reason, we do not ask you to activate them on the interfaces. Please make sure you understand how to apply them on physical QFX5100 devices interfaces.

See the following example:

```
set class-of-service interfaces xe-* forwarding-class-set LOSS output-traffic-control-profile LOSS
set class-of-service interfaces xe-* forwarding-class-set NOLOSS output-traffic-control-profile NOLOSS
```

Task 6.3: Policing

Configure vQFX2 so that NFS traffic from server 1 is policed to 2g and should use the NOLOSS forwarding class with low loss-priority.

vQFX2:

```
set firewall family ethernet-switching filter jnciedc term 1 from port nfsd
set firewall family ethernet-switching filter jnciedc term 1 then accept
set firewall family ethernet-switching filter jnciedc term 1 then forwarding-class BE
set firewall family ethernet-switching filter jnciedc term 1 then loss-priority low
set firewall family ethernet-switching filter jnciedc term 1 then policer 2g
set firewall family ethernet-switching filter jnciedc term 2 then accept

set firewall policer 2g if-exceeding bandwidth-limit 2g
set firewall policer 2g if-exceeding burst-size-limit 100m
set firewall policer 2g then discard

set interfaces xe-0/0/1 unit 0 family ethernet-switching filter input jnciedc
```


© 2020 Juniper Networks, Inc. All rights reserved.

Juniper Business Use Only