Nexus Validation Test

Phase 3

Contents

1. Introduction

The Cisco Nexus line of data center product hardware and software must pass Cisco's comprehensive quality assurance process, which includes a multistage approach comprising extensive unit test, feature test, and system-level test. Each successive stage in the process adds increasingly higher levels of complexity in a multidimensional mix of features and topologies.

Nexus Validation Test (NVT) has been established as an additional quality assurance stage in order to leverage customer feedback and requirements into the product development cycle. NVT will validate and publish guidelines for deploying NX-OS switching and UCS solutions for data center networks.

This document describes the NVT Phase 3 network topologies, hardware and software configurations, test procedures and findings.

NVT Phase 3 testing is performed on the following networks:

- Data Center 1 (DC1): This network focuses on building and operating a data center with the Nexus 7000 Sup1 as the core routing and switching component. It also covers interoperability with the Nexus 5000, Nexus 3000, Nexus 2000, Catalyst 6500/4500 switches and UCS B Series Servers. This network uses virtual PortChannel (vPC) and FabricPath to deliver highly available unicast and multicast services.
- Data Center 2 (DC2): This network focuses on building and operating a data center with the Nexus 7000 and 7700 Sup2E as the core routing and switching component. It also covers interoperability with the Nexus 6000, Nexus 5000, Nexus 3548, Nexus 2000 and Catalyst 6500/4500 switches. This network uses virtual PortChannel (vPC) and FabricPath to deliver highly available unicast and multicast services.
- Data Center 3 (DC3): This network focuses on building and operating multi-tiered data center networks with different Nexus product line switches:
    - Network DC31: This network is focused on the Nexus 6000 to deliver highly available unicast and multicast services with multipath routing. The unicast coverage includes both IPv4 and IPv6. It also covers interoperability with the Nexus 7000, Nexus 3548 and Nexus 3000.
    - Network DC32: This network is focused on the Nexus 3548 to deliver highly available unicast and multicast services with multipath routing. It also covers interoperability with the Nexus 7000 and Nexus 3000.
    - Network DC33: This network is focused on the Nexus 3000 to deliver highly available unicast and multicast services with multipath routing. The unicast coverage includes both IPv4 and IPv6. It also covers interoperability with the Nexus 7000.
    - Network DC36: This network is focused on the Nexus 3000 to deliver highly available unicast services with multipath routing. The unicast coverage includes both IPv4 and IPv6. It also covers interoperability with the Nexus 7000.

This document is split into different sections. Within sections 2, 3 and 5, each data center is described independently. The sections are:

- Section 2 – This section describes hardware/software components, and physical topology design.
- Section 3 – This section describes logical network design and configuration.
- Section 4 – This section describes test methodology and automation strategy.
- Section 5 – This section describes caveats and recommended workarounds.
- Section 6 – This section shows NVT test results.

2. NVT Topology Design Overview

2.1 DC1

2.1.1 Network Logical Topology Design Overview

The topologies and test cases validate highly-available data center networks in order to provide unified fabric and computing services. This is achieved by using the Nexus 7010, Nexus 5548 and UCS B-series servers with features such as vPC and FabricPath.

2.1.1.1 Description of the Test Network

The data center site is built around the Nexus 7000 with Sup 1. This data center site is split into two halves:

- Nexus 7000 with back-to-back vPC to Nexus 5000 with Nexus 2000 FEX, Nexus 7000 with vPC to Nexus 5000 for access, Nexus 7000 with Nexus 2000 FEX.
- Nexus 7000 with FabricPath to Nexus 5000, Nexus 5000 FabricPath leaf with Nexus 2000 FEX, UCS 6200 Fabric Interconnect, UCS 5108 series chassis and M2/M3 series blade servers.

While the majority of test cases focus on integrated solutions using Nexus switching and UCS products, modular Catalyst switches are also included for interoperability between NX-OS and IOS.

Figure 1 DC1 Topology

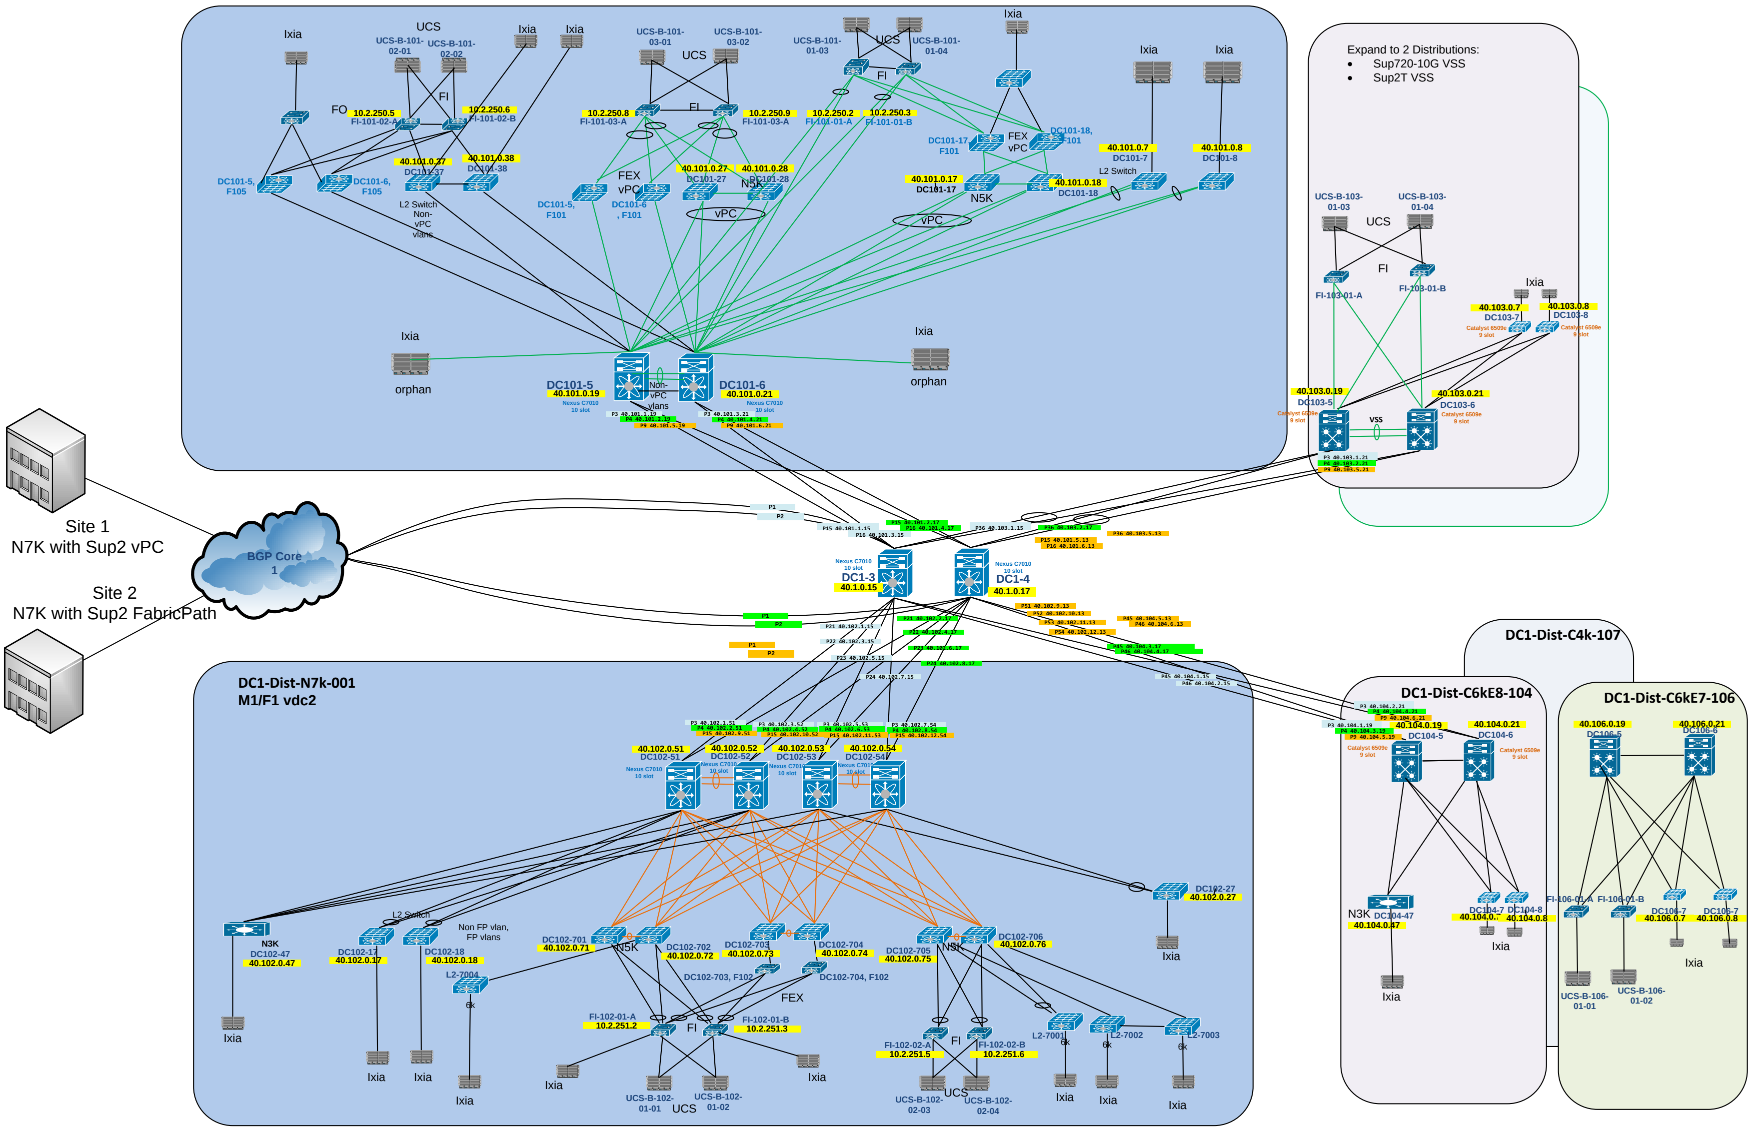

2.1.1.1.1 Core Routing

The core layer provides routing and high bandwidth connectivity between the aggregation-access blocks. The core layer of this data center is implemented using the Cisco Nexus 7000 Series Switch.

2.1.1.1.2 Aggregation-Access Blocks

The aggregation-access blocks provide connectivity and policy services for locally attached servers/hosts. These blocks are implemented as follows:

- Block 1 (DC101): Cisco Nexus 7000 Series Switch with virtual PortChannel (vPC).
- Block 2 (DC102): Cisco Nexus 7000 Series Switch with FabricPath (FP).
- Block 3-7: Blocks for Interoperability with Catalyst Platforms.

2.1.1.1.2.1 Block 1: Cisco Nexus 7000 Series Switch with virtual PortChannel (vPC)

Figure 2 Nexus 7000 vPC Topology

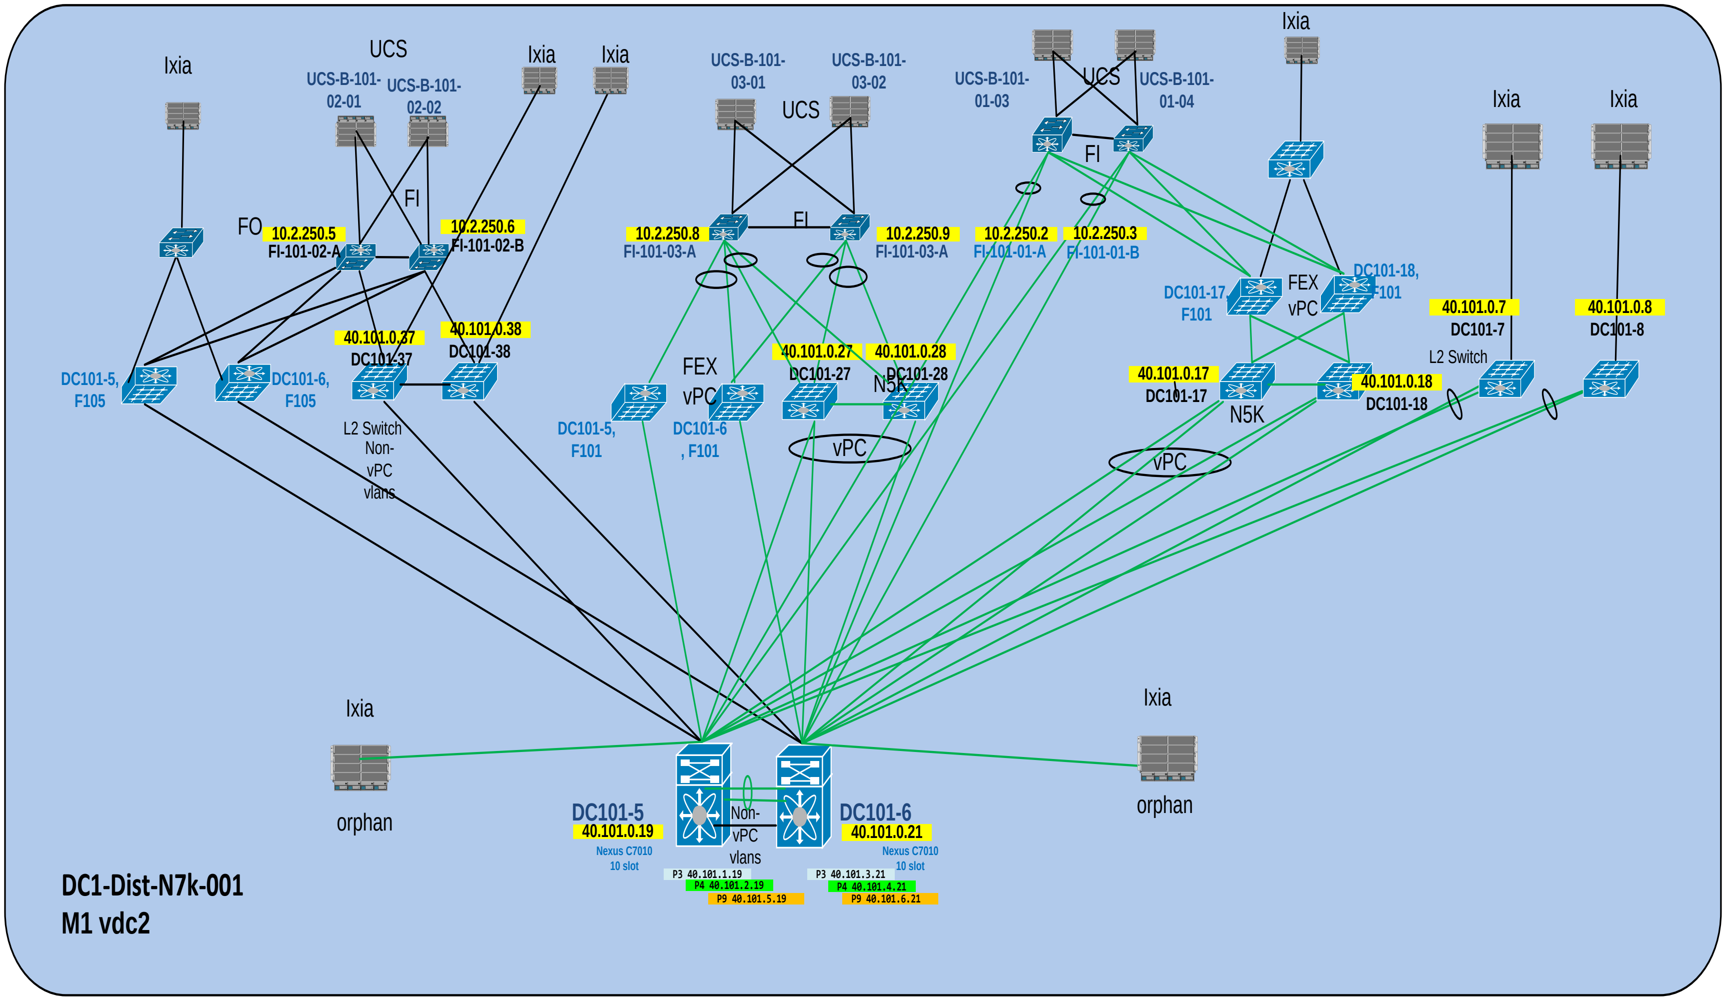

In this block the Nexus 7000 switches are used in vPC configuration on the aggregation level. The following types of Top of Rack devices are deployed:

- ToR FEX vPC: Fabric Extenders are directly attached to Nexus 7000 parent switches as well as the Nexus 5000 parent switches. The host ports are configured as vPC member ports.
- ToR Layer 2 Switch: Layer 2 switches are directly connected to the Nexus 7000 with vPC.
- ToR N5k vPC: A pair of Nexus 5000 switches is connected in a dual-sided vPC formation to the Nexus 7000 switches.

UCS B-series chassis are attached to UCS Fabric Interconnect (FI) clusters. The UCS FI clusters are directly connected to the Nexus 7000 switches as well as to the ToRs mentioned above, as shown in Figure 2.

2.1.1.1.2.2 Block 2: Cisco Nexus 7000 Series Switch with FabricPath (FP)

Figure 3 Nexus 7000 FabricPath Topology

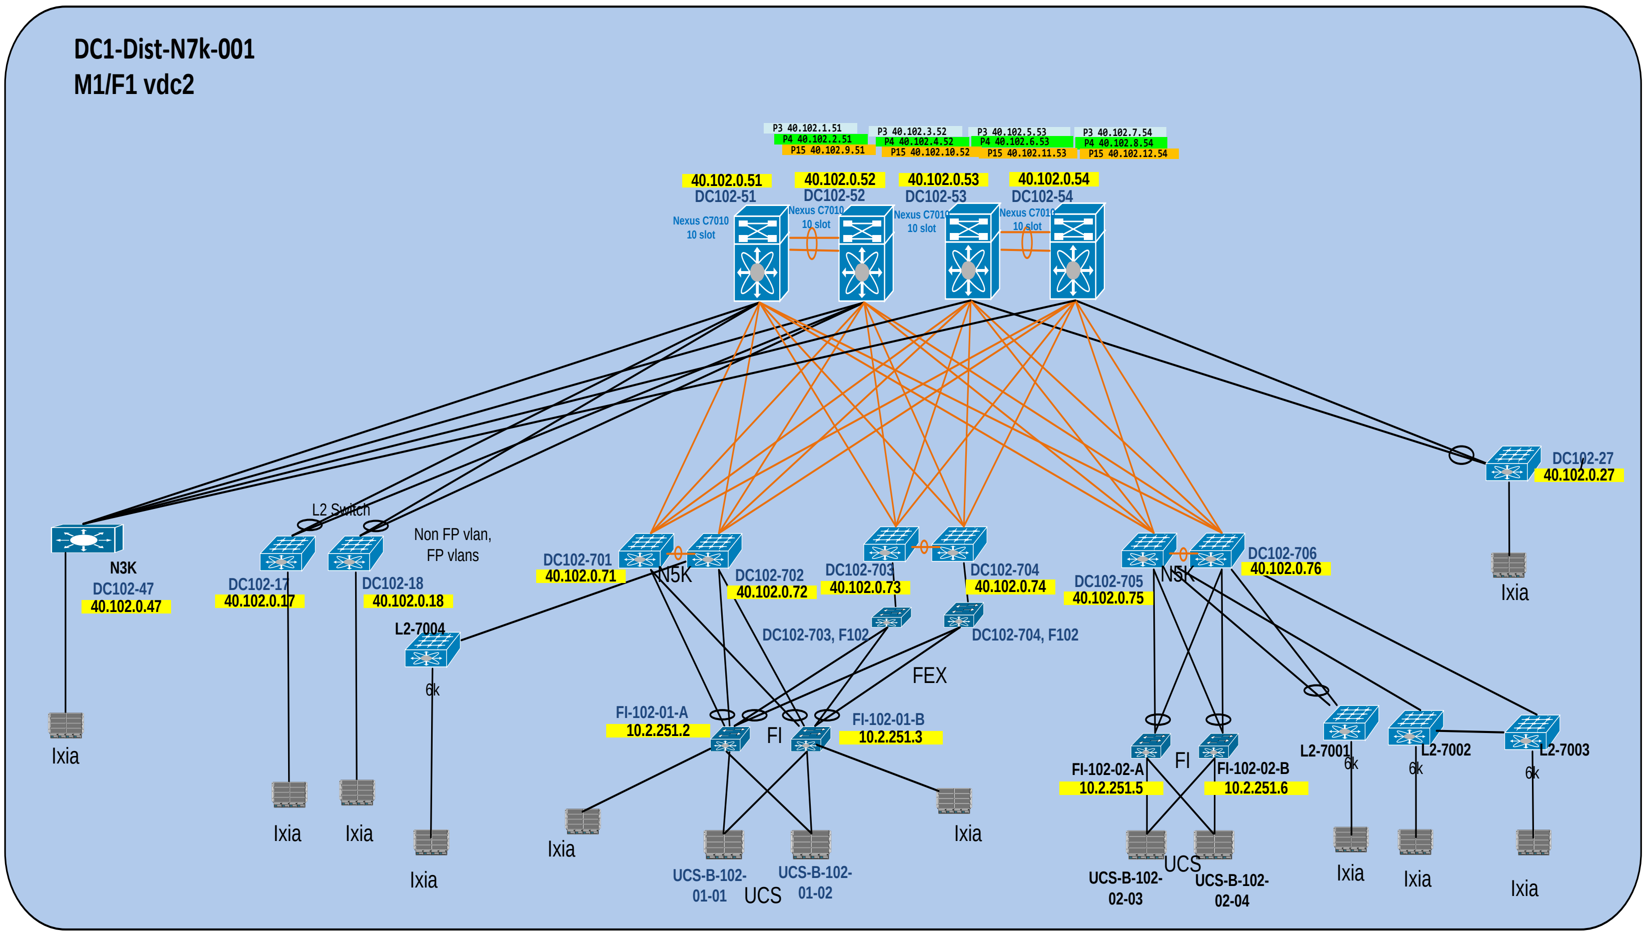

In this block the Nexus 7000 switches are used to form the spine layer for FabricPath. Nexus 5000 switches are deployed as the leaf layer. The following types of Top of Rack devices are deployed:

- ToR N5k FEX vPC+: Fabric Extenders are directly attached to Nexus 5000 parent switches on the FabricPath leaf. The host ports are configured as vPC+ member ports.
- ToR Layer 2 Switch: Layer 2 switches are directly connected to the Nexus 5000 switches on the FabricPath leaf.
- ToR Layer 2 Switch vPC+: Layer 2 switches are directly connected to the Nexus 7000 vPC+ on the FabricPath spine as well as the Nexus 5000 vPC+ on the FabricPath leaf.
- ToR N3k Layer 3: The Nexus 3000 is deployed as a Layer 3 access device. The Nexus 3000 are connected to the spine layer with routed links.

UCS B-series chassis are attached to UCS Fabric Interconnect (FI) clusters. The UCS FI clusters are directly connected to the Nexus 5000 leaf switches as well as some of the ToRs mentioned above, as shown in Figure 3.

2.1.1.1.2.3 Blocks for Interoperability with Catalyst Platforms

Blocks 3 to 7 are used to test interoperability of the Catalyst platform switches with the Nexus line of switches

- Block 3: Cisco Catalyst 6500 Series Switch Supervisor Engine 2T VSS
- Block 4: Cisco Catalyst 6500 Series Switch Supervisor Engine 2T
- Block 5: Cisco Catalyst 6500 Series Switch Supervisor Engine 720-10G VSS
- Block 6: Cisco Catalyst 6500 Series Switch Supervisor Engine 720
- Block 7: Cisco Catalyst 4500 Series Switch

UCS B-series chassis are attached to UCS Fabric Interconnect (FI) clusters. The UCS FI clusters are directly connected to Block 3 and Block 6.

2.1.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- BGP: eBGP is configured between the core switches and the public cloud.
- OSPF: OSPF is the IGP running across the network. Each aggregation-access block is configured as a unique area with the core switches playing the role of the ABR.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. Each aggregation-access block is configured with the RP for the locally sourced groups.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs.
- vPC: vPC technology is deployed in the aggregation-access block DC1-Dist-N7k-101 as shown in Figure 1. In addition, dual-sided vPC is configured between the Nexus 7000 and Nexus 5000 switches

- FP: FabricPath is deployed in the aggregation block DC1-Dist-N7k-102. The spine layer is comprised of Nexus 7000 switches and the leaf switches are deployed using Nexus 5000 switches.
- VLAN trunking: VLAN trunking is used in the aggregation-access blocks to maintain segregation and security.
- STP: Rapid Spanning Tree Protocol is used to prevent Layer 2 loops in the aggregation-access blocks. The spanning tree root is placed on the aggregation level. Root Guard is configured on the aggregation level to enforce root placement. BPDU Filter, BPDU Guard and PortFast Edge are configured on the access ports towards hosts.
- HSRP: HSRP is used as the first hop gateway protocol for hosts.
- FEX: Multiple types of Fabric Extenders are deployed on Nexus 7000 and Nexus 5000 parent switches.
- IGMP: IGMP is used by hosts to join multicast groups of interest. IGMP snooping is enabled on all switches in the aggregation-access blocks to prevent flooding of multicast data traffic.
- LACP: LACP is used for link aggregation to form port-channels across the network.
- UDLD: UDLD aggressive mode is configured across the network to detect and prevent unidirectional links.
- DHCP relay: DHCP relay is enabled on the aggregation layer to provide IP address services to hypervisors and VMs running on UCS systems.
- End-Host Mode: All of the FI clusters are configured to run in End-Host Mode in order to prevent loops within the topology.
- VM-FEX: VM-FEX has been deployed to provide a direct connection for all of the virtual machines' network interfaces to the UCS Fabric Interconnect.

2.1.2 Hardware and Software Overview

DC 1:

Platform	Model No.	NVT 3.0
N7K	N7K-SUP1	6.2.6; 6.2.6a
N5K	N5K-C5548UP-SUP	5.2.1.N1.4
N3K	N3K-C3048TP-1GE-SUP	5.0.3.U5.1b
C6K	VS-SUP2T-10G	150-1.SY3
	VS-S720-10G	122-33.SXJ4
	WS-SUP720	122-33.SXJ4
	WS-SUP32-GE	122-33.SXJ
C4K	WS-X45-SUP7-E	03.03.02.SG.151-1.SG2
	WS-C4948	150-2.SG6-6.9
UCS	UCS-5108	N/A
	UCS-B200-M2	2.1(2a)B
	UCS-B22-M3	2.1(2a)B
	UCS-2208XP-IOM	2.1(2a)A
	UCS-6248UP-FI	2.1(2a)A
	UCS-6296UP-FI	2.1(2a)A
	UCS-M81KR-VIC	2.1(2a)B
	UCS-VIC-1280	2.1(2a)B

2.1.2.1 Nexus 7000 Line Cards and Fabric Extenders (FEX)

The following line cards are used on the Nexus 7000 devices:

- N7K-M108X2-12L
- N7K-M132XP-12L
- N7K-F132XP-15

The following types of FEX are utilized in the network:

- N2K-C2224TP-1GE
- N2K-C2248TP-E-1GE
- N2K-C2248TP-1GE
- N2K-C2232PP-10GE

2.1.2.2 Unified Computing System (UCS) Physical

2.1.2.2.1 Unified Computing System (UCS) Hardware

The hardware used in the NVT UCS setup contains the following:

- Cisco UCS 6248UP 48-Port Fabric Interconnect
- Cisco UCS 6296UP 96-Port Fabric Interconnect
- UCS 5108 Blade Server Chassis
- UCS 2208XP Fabric Extender (IOM)
- Cisco B200 M2 Blade Server
- Cisco B22 M3 Blade Server
- Cisco M81KR Virtual Interface Card
- Cisco Virtual Interface Card (VIC) 1280

2.1.2.2.2 Unified Computing System (UCS) Upstream Switch Connectivity

	Fabric Interconnect		Blade		Mezzanine		Chassis/ IOM
DC1	**Cisco UCS 6248UP**	**Cisco UCS 6296UP**	**Cisco B200 M2**	**Cisco B22 M3**	**Cisco UCS VIC 1280**	**Cisco UCS M81KR**	**UCS 5108/ UCS-IOM-2208XP**
N7k vpc (M1) (101-01) DC101-5/6		X		X	X		X
		X	X			X	X
N7k vpc (F1) (101-01) DC101-5/6		X		X	X		X
		X	X			X	X
N7k Fex (101-02) DC101-5/6,F105 (N2K-C2232PP-10GE)		X	X			X	X
N7k Fex vpc (101-03) DC101-5/6,F104 (N2K-C2232PP-10GE)	X		X			X	X
N5k vpc (101-03) DC101-27/28	X		X			X	X
N5k Fex vpc (101-01) DC101-17/18 (N2K-C2224TP-1GE)		X		X	X		X
		X	X			X	X
N5k FabricPath (102-01) DC102-701/702	X	X		X	X		X
	X	X	X			X	X
N5k FabricPath Fex (102-01) DC102-703-704 (N2K-C2232PP-10GE)		X	X			X	X
Cat6k Earl 8 VSS (103-01) DC103-VSS (WS-X6904-40G)	X			X	X		X

Cat6k Earl 7 standalone (106-01) DC106 (WS-X6708-10GE WS-X6704-10GE)	X			X	X		X
L2 Switch 4849 (101-02) DC101-37/38		X	X			X	X
L2 Switch 6509 (102-02) DC102-17/18	X			X		X	X

2.2 DC2

2.2.1 Network Logical Topology Design Overview

The topologies and test cases validate highly-available data center networks in order to provide unified fabric and computing services. This is achieved by using the Nexus 7000, Nexus 7700, Nexus 6000, Nexus 5000, Nexus 2000 and Nexus 3500 switches.

2.2.1.1 Description of the Test Network

Figure 4 illustrates the test network topology of DC2 data center, which is built around Nexus 7000 with Sup 2E. This data center site is split into two halves:

- Nexus 7000 with vPC to Nexus 5000 for access.
- Nexus 7000 with FabricPath to Nexus 5000, Nexus 6000 and Nexus 7700. Nexus 2000 is connected to Nexus 7000 FabricPath spine and to FabricPath leaf's: Nexus 5000 and Nexus 6000.

While the majority of test cases focus on integrated solutions using Nexus switching, modular Catalyst switches are also included for interoperability between NX-OS and IOS.

Figure 4 DC2 Topology

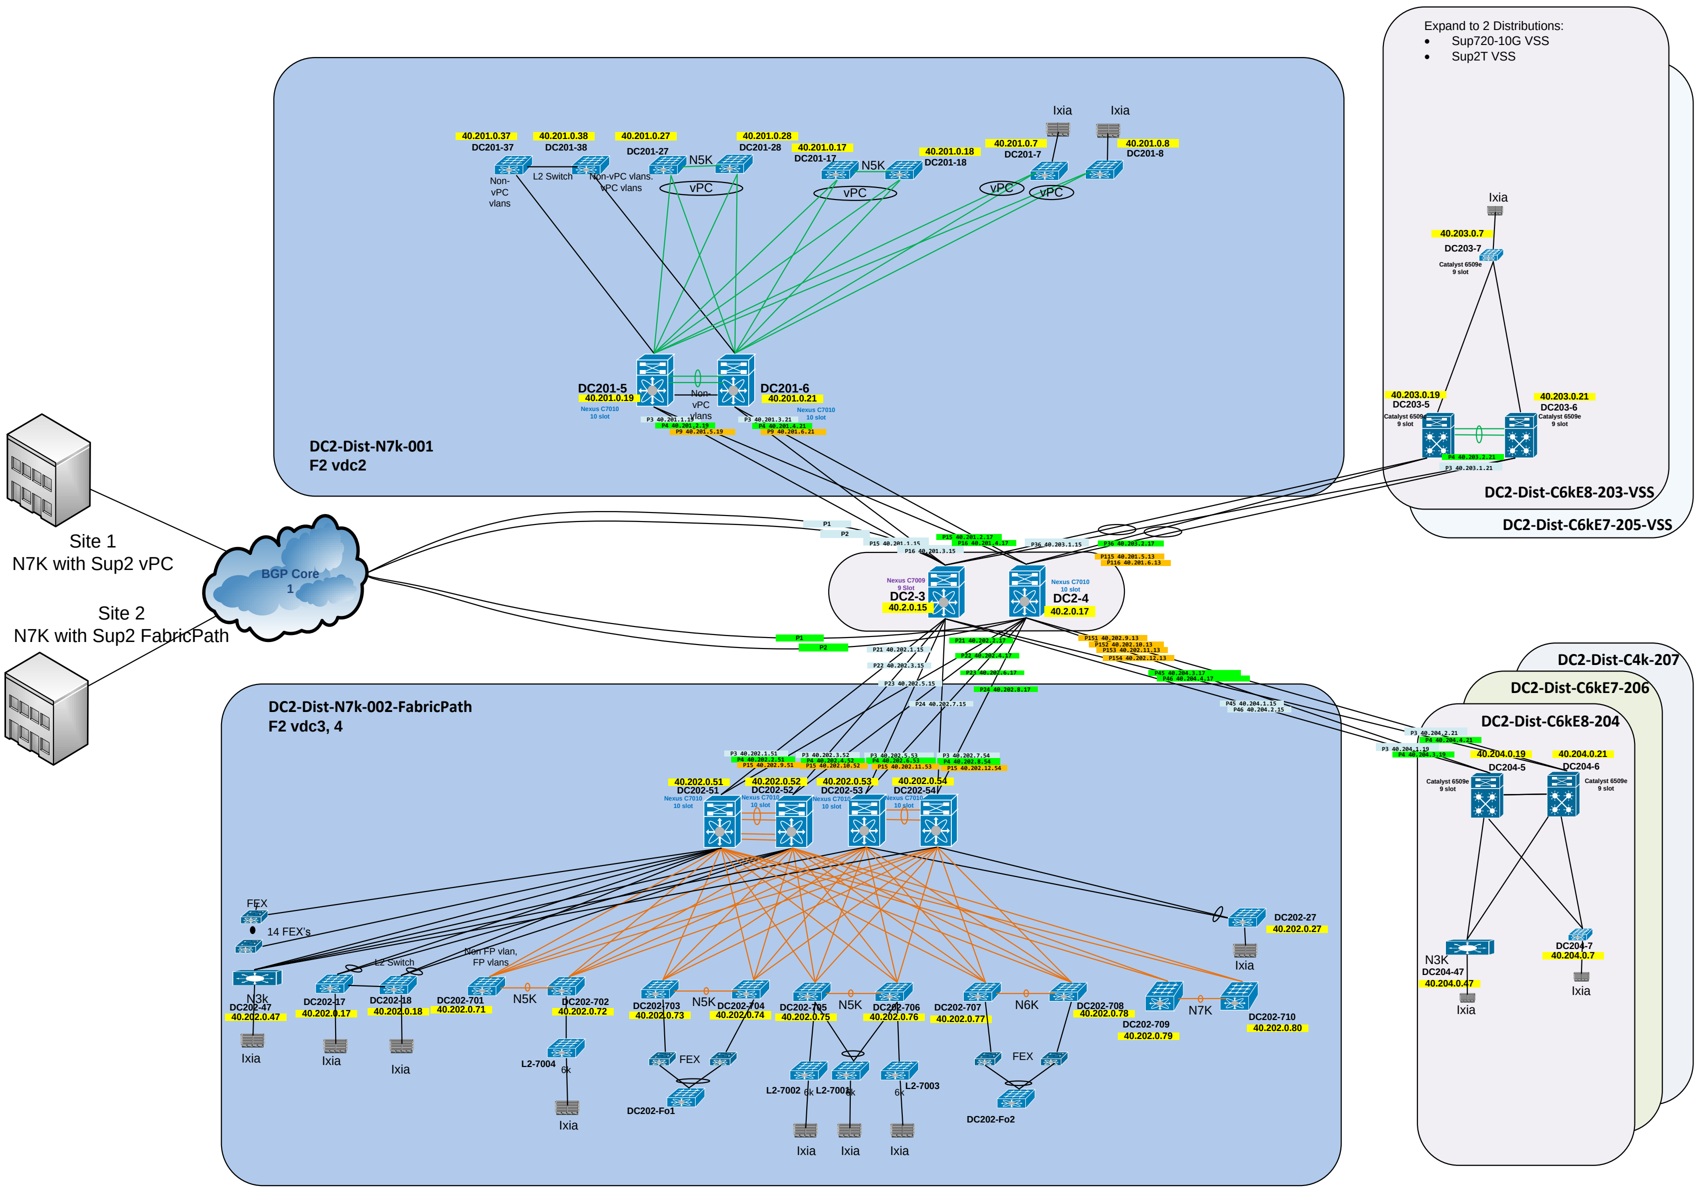

2.2.1.1.1 Core Routing

The core layer provides routing and high bandwidth connectivity between the aggregation-access blocks. The core layer of this data center is implemented using Nexus 7000 Series Switches.

2.2.1.1.2 Aggregation-Access Blocks

The aggregation-access blocks provide connectivity and policy services for locally attached servers/hosts. These blocks are implemented as follows:

- Block 1(DC201): Cisco Nexus 7000 Series Switch with virtual PortChannel (vPC).
- Block 2(DC202): Cisco Nexus 7000 Series Switch with FabricPath (FP).
- Blocks for Interoperability with Catalyst Platforms.

2.2.1.1.2.1 Block 1: Cisco Nexus 7000 Series Switch with virtual PortChannel (vPC)

Figure 5 Nexus 7000 vPC Topology

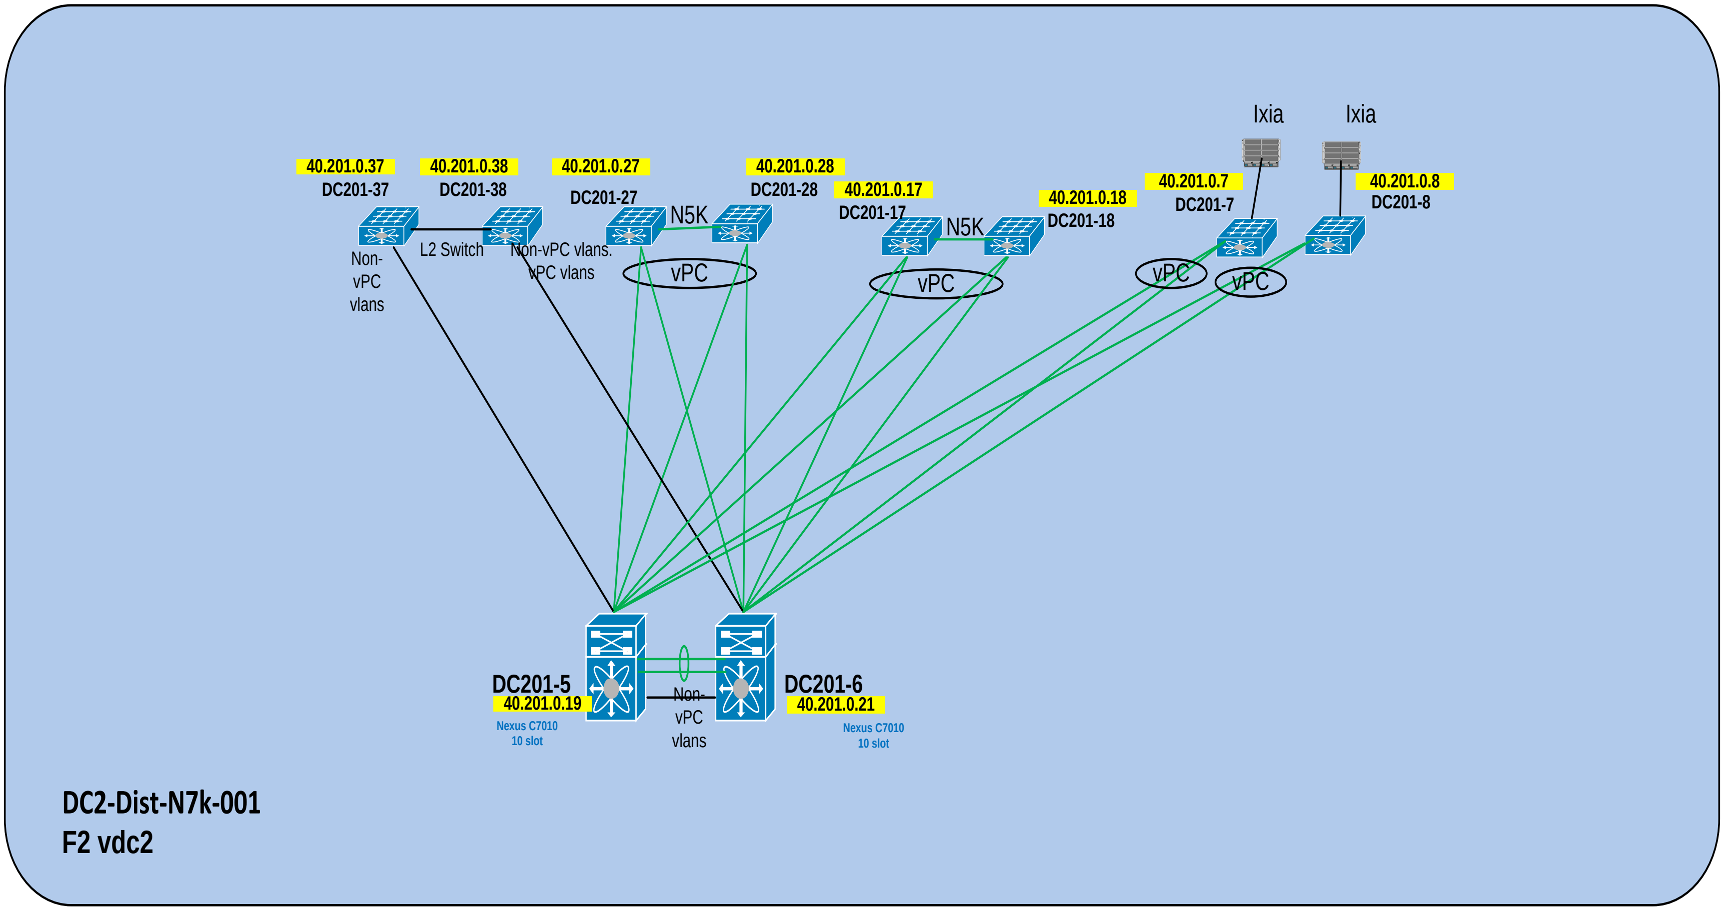

In this block the Nexus 7000 switches are used in vPC configuration on the aggregation level. The following types of Top of Rack devices are deployed:

- ToR Layer 2 Switch: Layer 2 switches are directly connected to the Nexus 7000 with vPC.
- ToR N5k vPC: A pair of Nexus 5000 switches is connected in a dual-sided vPC formation to the Nexus 7000 switches.

2.2.1.1.2.2 Block 2: Cisco Nexus 7000 Series Switch with FabricPath (FP)

Figure 6 Nexus 7000 FabricPath Topology

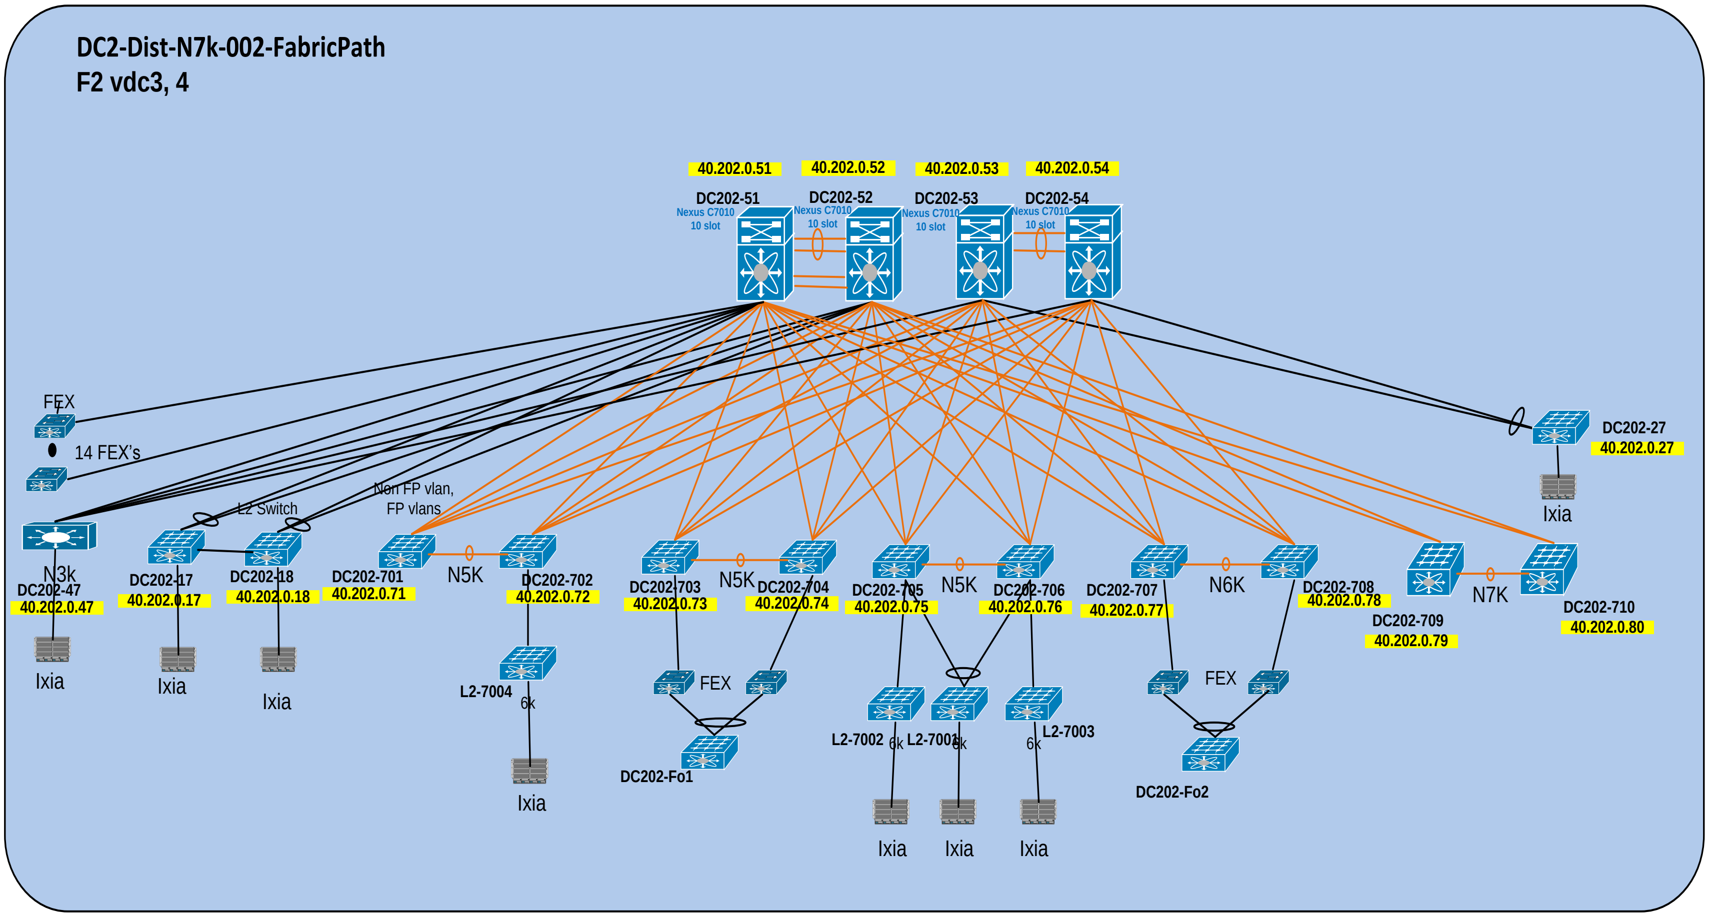

In this block the Nexus 7000 switches are used to form the spine layer for FabricPath. Nexus 5000, Nexus 6000 and Nexus 7700 switches are deployed at the leaf layer. The following types of Top of Rack devices are deployed:

- ToR N5k FEX vPC+: Fabric Extenders are directly attached to Nexus 5000 parent switches on the FabricPath leaf. The host ports are configured as vPC+ member ports.
- ToR N6k FEX vPC+: Fabric Extenders are directly attached to Nexus 6000 parent switches on the FabricPath leaf. The host ports are configured as vPC+ member ports.
- ToR N7k FEX vPC+: 14 Fabric Extenders are directly attached to one of the Nexus 7000 parent switches on the FabricPath spine.
- ToR Layer 2 Switch: Layer 2 switches are directly connected to the Nexus 5000 switches on the FabricPath leaf.
- ToR Layer 2 Switch vPC+: Layer 2 switches are directly connected to the Nexus 7000 vPC+ on the FabricPath spine as well as the Nexus 5000 vPC+ on the FabricPath leaf.
- ToR N3k Layer 3: The Nexus 3548 is deployed as a Layer 3 access device. The Nexus 3548 is connected to the spine layer with routed links.

2.2.1.1.2.3 Blocks for Interoperability with Catalyst Platforms

Blocks 3 to 7 are used to test interoperability of the Catalyst platform switches with the Nexus line of switches

- Block 3: Cisco Catalyst 6500 Series Switch Supervisor Engine 2T VSS
- Block 4: Cisco Catalyst 6500 Series Switch Supervisor Engine 2T
- Block 5: Cisco Catalyst 6500 Series Switch Supervisor Engine 720-10G VSS
- Block 6: Cisco Catalyst 6500 Series Switch Supervisor Engine 720
- Block 7: Cisco Catalyst 4500 Series Switch

2.2.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- BGP: eBGP is configured between the core switches and the public cloud.
- OSPF: OSPF is the IGP running across the network. Each aggregation-access block is configured as a unique area with the core switches playing the role of the ABR.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. Each aggregation-access block is configured with the RP for the locally sourced groups.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs.

- vPC: vPC technology is deployed in the aggregation-access block DC2-Dist-N7k-201. In addition, dual-sided vPC is configured between the Nexus 7000 and Nexus 5000 switches.
- FP: FabricPath is deployed in the aggregation blocks DC2-Dist-N7k-202. The spine layer is comprised of Nexus 7000 switches and the leaf switches are deployed using Nexus 5000, Nexus 6000 and Nexus 7700 switches.
- VLAN trunking: VLAN trunking is used in the aggregation-access blocks to maintain segregation and security.
- FP VLANs: On DC2-Dist-N7k-202, 2000 VLANs are deployed in mode FabricPath on all the spines and leaf's.
- STP: Rapid Spanning Tree Protocol is used to prevent Layer 2 loops in the aggregation-access block DC-Dist-N7K-201. MSTP is enabled on DC-Dist-N7K-202 for the same purpose wherever applicable. The spanning tree root is placed on the aggregation level. Root Guard is configured on the aggregation level to enforce root placement. BPDU Filter, BPDU Guard and PortFast Edge are configured on the access ports towards hosts.
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- HSRP: HSRP is used as the first hop gateway protocol for hosts.
- FEX: Multiple types of Fabric Extenders are deployed on Nexus 5000 parent switches.
- IGMP: IGMP is used by hosts to join multicast groups of interest. IGMP snooping is enabled on all switches in the aggregation-access blocks to prevent flooding of multicast data traffic.
- LACP: LACP is used for link aggregation to form port-channels across the network.
- UDLD: UDLD aggressive mode is configured across the network to detect and prevent unidirectional links.

2.2.2 Hardware and Software Overview

DC 2:

Platform	Model No.	NVT 3.0
N7000	N7K-SUP2E	6.2.6; 6.2.6a
N5000	N5K-C5548P -SUP	5.2.1.N1.4
	N5K-C5548UP-SUP	5.2.1.N1.3
N3548	N3K-C3548P-10G-SUP	5.0.3.A1.2
N6000	N6K-C6001-64P-SUP	6.0(2)N2(3)
N7700	N77-SUP2E	6.2.6; 6.2.6a
C6K	VS-SUP2T-10G	150-1.SY3
	VS-S720-10G	122-33.SXJ4
	WS-SUP720	122-33.SXJ4
C4K	WS-X45-SUP7-E	03.03.02.SG.151-1.SG2
	WS-C4948	150-2.SG6-6.9

2.2.2.1 Nexus 7000 and Nexus 7700 Line Cards and Fabric Extenders (FEX)

The following line cards are used on the Nexus 7000 devices:

- N7K-F248XP-25
- N7K-F248XP-25E
- N7K-M224XP-23L

The following line cards are used on the Nexus 7700 devices:

- N77-F248XP-23E

The following types of FEX are utilized in the network:

- N2K-C2224TP-1GE
- N2K-C2248TP-E-1GE

2.3 DC3 Core

2.3.1 Network Logical Topology Design Overview

The topologies and test cases validate highly-available data center networks. This is achieved by using three Nexus 7000 switches in the core network.

2.3.1.1 Description of the Test Network

Figure 7 illustrates the overall test network topology. The core layer consists of three Nexus 7000 switches connected in a meshed topology. The core layer provides routing and high bandwidth connectivity between the spine layers of different PODs.

Two of the Nexus 7000 core switches are connected to each of the spine switches for the following PODs:

- DC31:
  - Spine: Nexus 6004
  - Leaf: Nexus 6001, Nexus 3548, Nexus 3048, Nexus 7000 (used for peer-scale with VRF configuration)
- DC32:
  - Spine: Nexus 3548
  - Leaf: Nexus 3548, Nexus 3048, Nexus 7000 (used for peer-scale with VRF configuration)
- DC33:
  - Spine: Nexus 3048
  - Leaf: Nexus 3048, Catalyst 6500 (used for peer-scale with VRF configuration)
- DC36:
  - Spine: Nexus 3048, Nexus 3064
  - Leaf: Nexus 3048, Nexus 3064, Catalyst 6500 (used for peer-scale with VRF configuration)

Figure 7 DC3 Topology

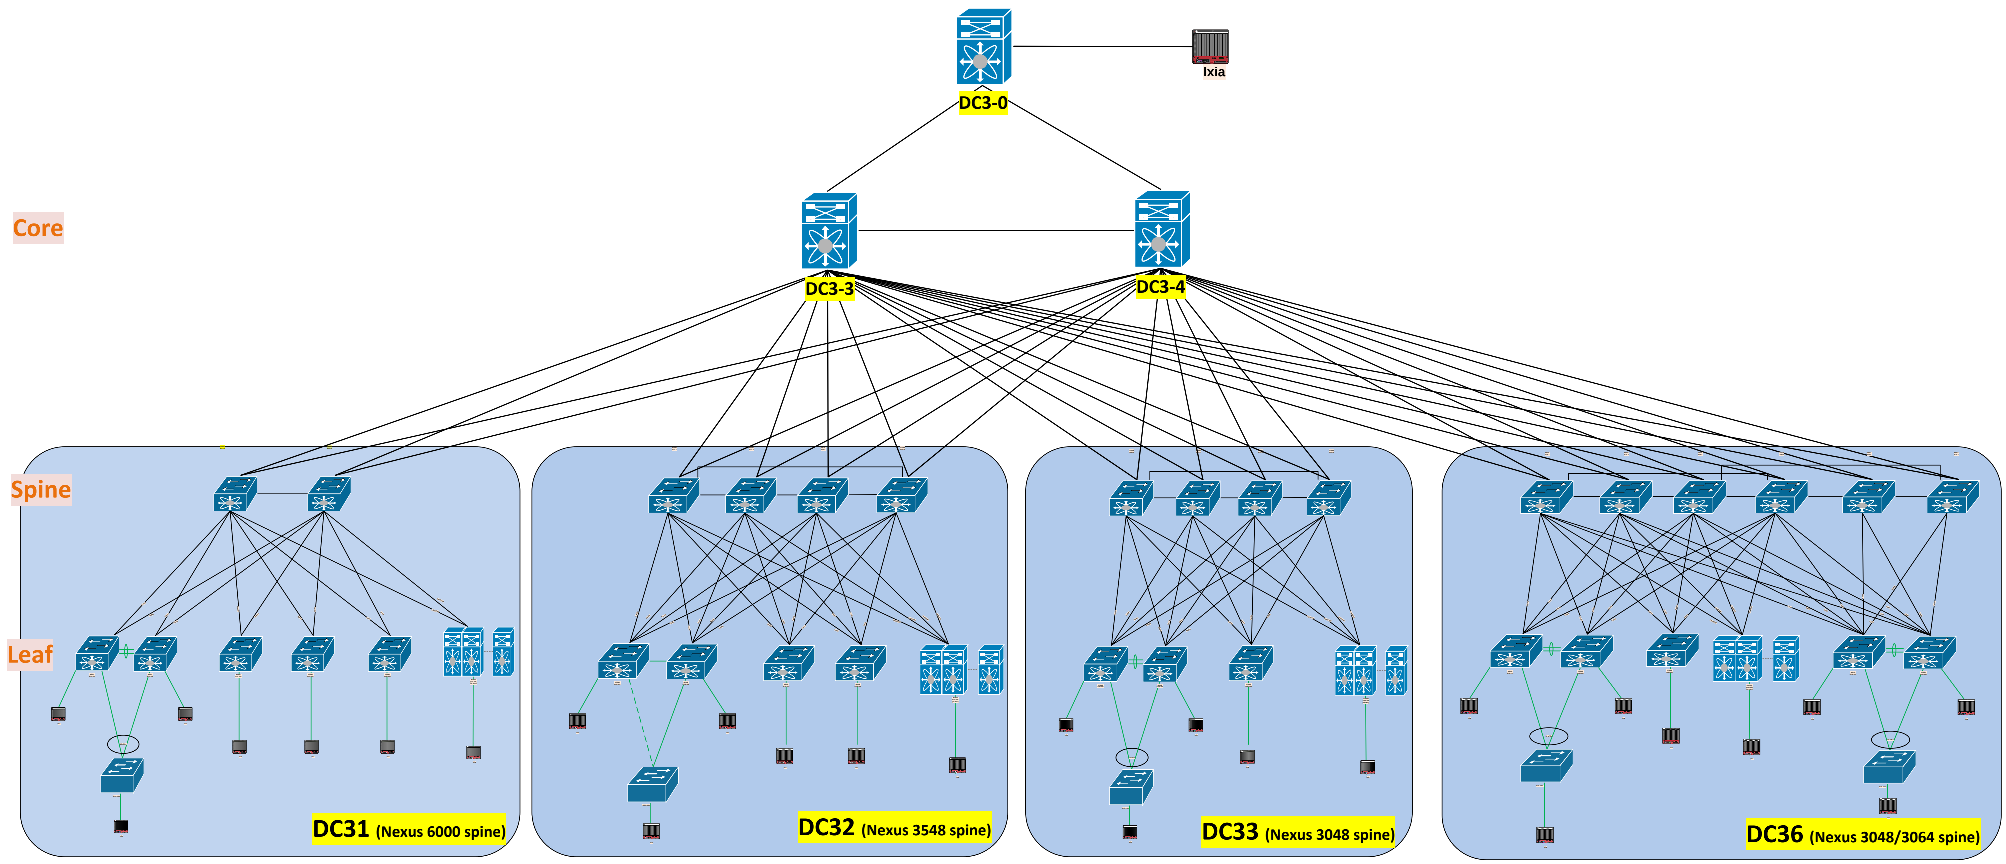

2.3.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- Jumbo MTU: Jumbo MTU is configured as 9216 across the network..
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- Dual Stack Interfaces: All Layer 3 interfaces including routed port, routed port-channel and SVI are configured as IPv4/IPv6 dual stack interfaces.
- BGP: An iBGP session is established between two core switches. Two eBGP sessions are configured between the 2 iBGP peers and the third core switch. Also, eBGP sessions are established between the 2 iBGP peers and each of the DC31, DC32, DC33, DC36 spine switches. IPv4/IPv6 address families are configured for all BGP peers with a maximum-path set to 32.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. The static RP is located at the two iBGP peers for all sourced groups.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs.
- VLAN trunking: VLAN trunking is used in one of the core switches to connect to the traffic simulator tool.
- UDLD: UDLD aggressive mode is configured to detect and prevent unidirectional links.
- LACP: LACP is used for link aggregation to form port-channels across the network.
- CDP/LLDP: CDP is used by default. LLDP is also used for link and neighbor discovery information.
- CoPP: CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

2.3.2 Hardware and Software Overview

Platform	Model No.	NVT 3.0
N7000	N7K-SUP2E	6.2(6)
N7000	N7K-SUP1	6.1(4)

The following line cards are used on the Nexus 7000 (N7K-SUP2E) devices:

- N7K-F248XP-25
- N7K-F312FQ-25

The following line cards are used on the Nexus 7000 (N7K-SUP1) device:

- N7K-F248XP-25

2.4 DC31

2.4.1 Network Logical Topology Design Overview

The topologies and test cases validate highly-available data center networks. This is achieved by using the Nexus 6004, Nexus 6001, Nexus 3048 and Nexus 3548 line of switches in a spine and leaf topology.

2.4.1.1 Description of the Test Network

Figure 8 illustrates the test network topology. The spine layer consists of 2 Nexus 6004 switches. The leaf layer is comprised of Nexus 6001, Nexus 3548, Nexus 3048 and Nexus 7000 switches with ECMP connections to each of the two spine switches.

During NVT Phase 3, the main focus has been the analysis and validation of ECMP deployments for both unicast IPv4 and IPv6 traffic as well as multicast multipath traffic on the above-mentioned Nexus platforms in a spine and leaf topology.

Figure 8 DC31 Topology

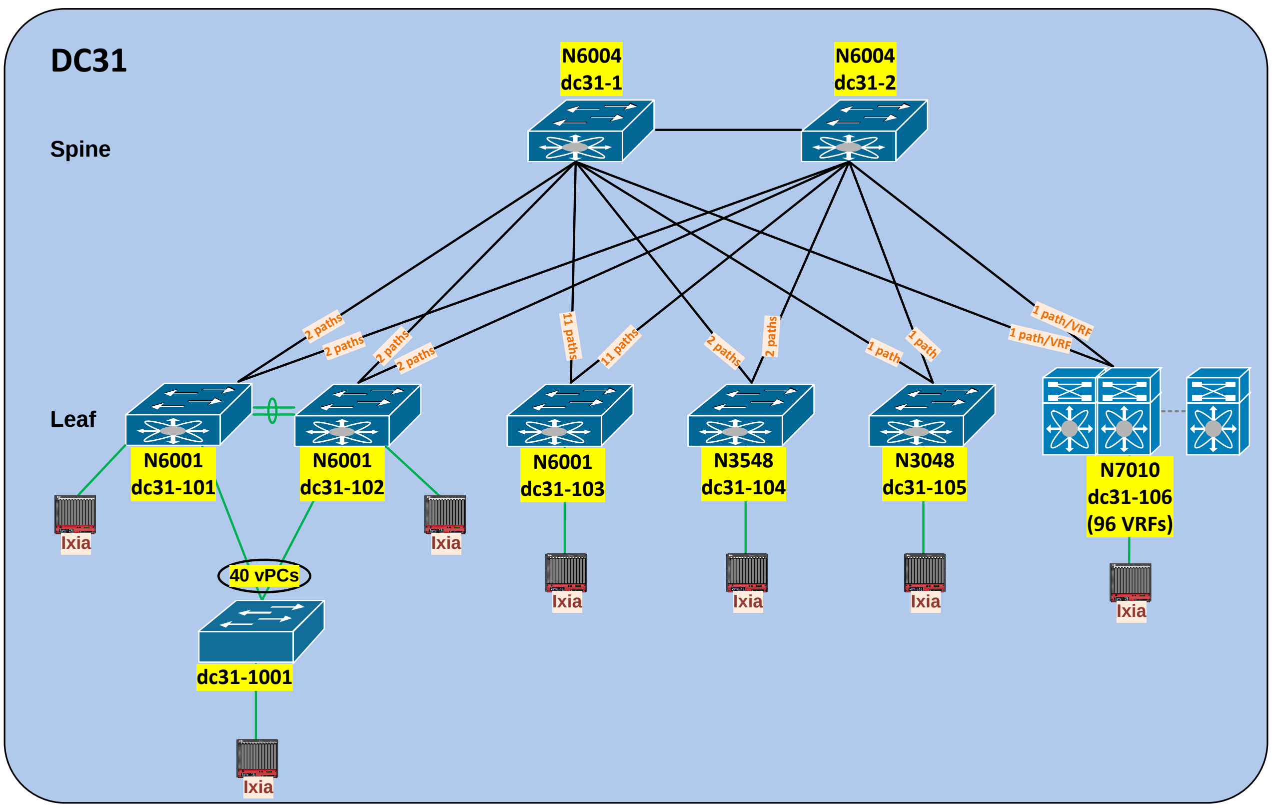

2.4.1.1.1 Spine Layer

The spine layer provides ECMP routing and high bandwidth connectivity between the Leaf/Access layer switches. The spine layer is implemented using the following platform type:

- Cisco Nexus 6004 Series Switch

2.4.1.1.2 Leaf/Access Layer

The Leaf/Access layer provides connectivity and policy services for loca lly attached hosts. These leaf switches are deployed as the following types of devices:

- Nexus 6001 switch with vPC: Two Nexus 6001 switches are configured as vPC peers. Two 40G ECMP interfaces are connected to each spine switch while 40 vPC port-channels with one member from each peer are connected to the Layer 2 access switch.
- Nexus 6001 switch: One Nexus 6001 switch is configured as a Layer 3 leaf and is connected to each spine switch with 11 port-channels (2 members/port-channel).
- Nexus 3548 switch: One Nexus 3548 switch is configured as a Layer 3 leaf and is connected to each spine switch with 2 port-channels (4 members/port-channel).
- Nexus 3048 switch: One Nexus 3048 switch is configured as a Layer 3 leaf and is connected to each spine switch with 1 port-channel (1 member/port-channel).
- Nexus 7000 switch: One Nexus 7000 switch is connected over 96 VRFs, with 1 routed interface to each spine switch for each VRF.

2.4.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- Jumbo MTU: Jumbo MTU is configured as 9000 across the network..
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- Dual Stack Interface: All Layer 3 interfaces including routed port, routed port-channel and SVI are configured as IPv4/IPv6 dual stack interfaces.
- BGP: eBGP is configured between the spine and the core, and between the spine and leaf. iBGP is configured between the spine switches. IPv4/IPv6 address families are configured for all BGP peers. Maximum-paths are configured for equal-cost multipath load balancing as 64 for both spine and leaf peers for IPv4/IPv6 address families..
- OSPF/OSPFv3: OSPF/OSPFv3 is used as the IGP to provide reachability for establishing iBGP peering at the spine layer.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. The RP is located at the spine layer.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs.

- vPC: vPC technology is deployed in the Leaf/Access layer.
- VLAN trunking: VLAN trunking is used in the Leaf/Access layer to maintain segregation and security.
- STP: Rapid Spanning Tree Protocol is used to prevent Layer 2 loops in the Leaf/Access layer. The spanning tree root is placed on the leaf layer. Root Guard is configured on the leaf layer to enforce root placement. BPDU Filter, BPDU Guard and PortFast Edge are configured on the access ports towards hosts.
- HSRP/HSRPv6: HSRP/HSRPv6 is used as the first hop gateway protocol for IPv4/IPv6 hosts.
- LACP: LACP is used for link spine and leaf layer to form port-channels across the network.
- CDP/LLDP: CDP is used by default. LLDP is also used for link and neighbor discovery information.
- IGMP: IGMP is used by hosts to join multicast groups of interest. IGMP snooping is enabled on all switches in the Leaf/Access layers to prevent flooding of multicast data traffic.
- CoPP: CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

2.4.2 Hardware and Software Overview

Platform	Model No.	NVT 3.0
N6004	N6K-C6004-96Q-SUP	6.0(2)N2(3)
N6001	N6K-C6001-64P-SUP	6.0(2)N2(3)
N3000	N3K-C3048TP-1GE-SUP	6.0(2)U1(3)
N3548	N3K-C3548P-10G-SUP	6.0(2)A1(1c)
N7000	N7K-SUP1	6.2.2

The following line cards are used on the Nexus 7000 devices:

- N7K-F248XP-25

2.5 DC32

2.5.1 Network Logical Topology Design Overview

The topologies and test cases validate highly-available data center networks. This is achieved by using Nexus 3548 and Nexus 3048 line of switches in a spine and leaf topology.

2.5.1.1 Description of the Test Network

Figure 9 illustrates the test network topology. The spine layer consists of 4 Nexus 3548 switches. The leaf switches are comprised of Nexus 3548, Nexus 3048, and Nexus 7000 platforms with ECMP connections to each spine.

During NVT Phase 3, the main focus has been the analysis and validation of ECMP deployments for unicast IPv4 traffic as well as multicast multipath traffic on the Nexus 3548, Nexus 3048, and Nexus 7000 platforms in a spine and leaf topology.

Figure 9 DC32 Topology

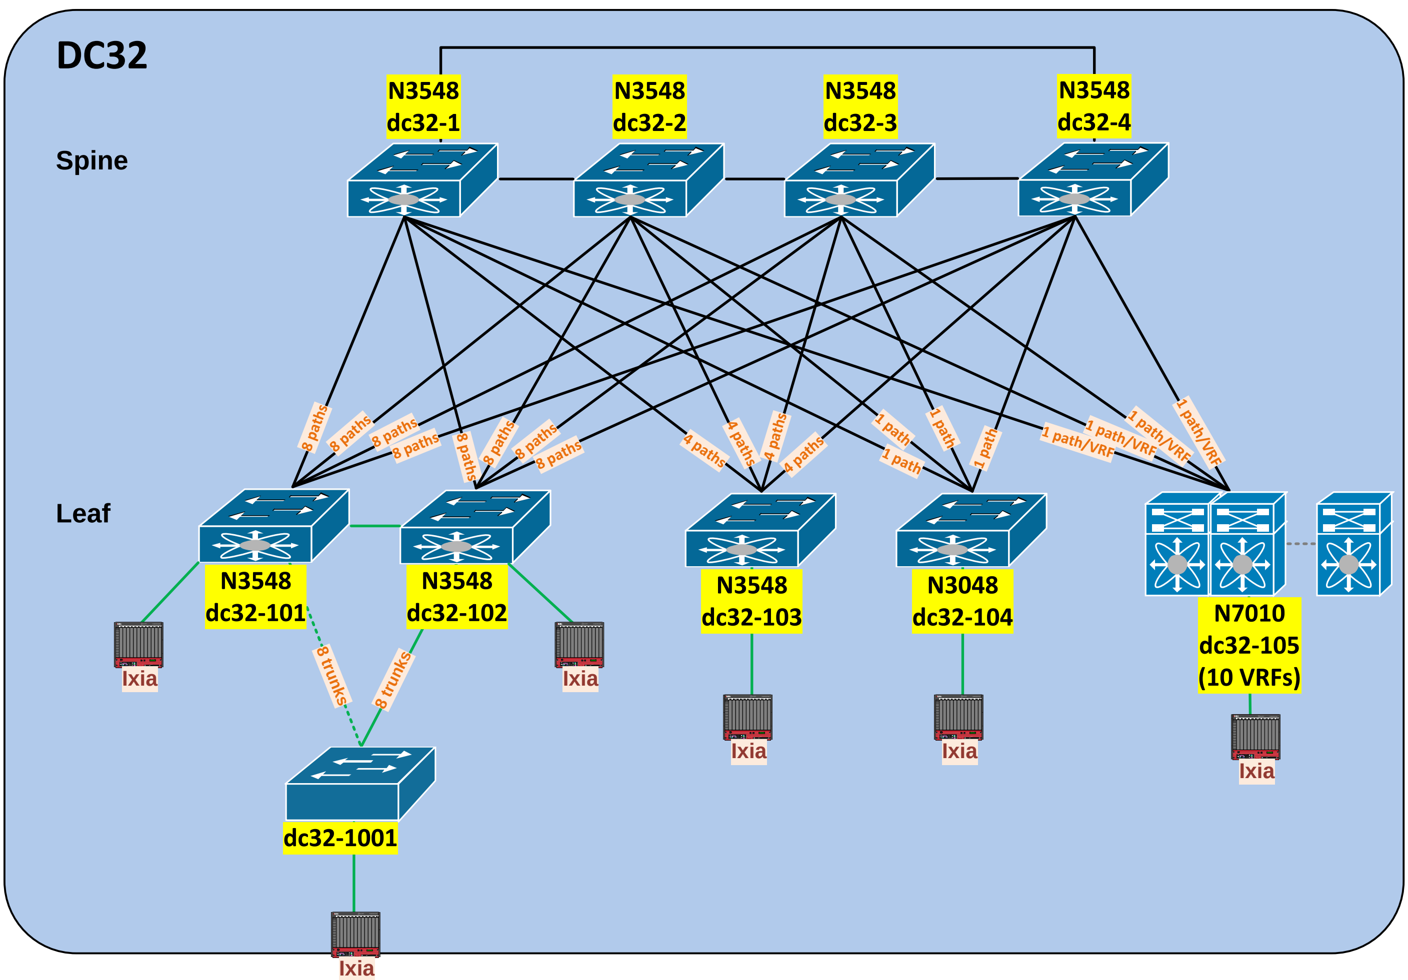

2.5.1.1.1 Spine Layer

The spine layer provides ECMP routing and high bandwidth connectivity between the Leaf/Access layer switches. The spine layer in the test network is implemented using the following platform type:

- Cisco Nexus 3548 Series Switch

2.5.1.1.2 Leaf/Access Layer

The Leaf/Access layer provides L3 and L2 connectivity and policy services for locally attached servers/hosts. These leaf switches are implemented as follows:

- Nexus 3548 with MSTP: Two standalone switches each with eight ECMP paths connected to each spine. While 8 port-channels with one member each are connected to the Layer 2 access switch.
- Nexus 3548: One Nexus 3548 switch is configured as a Layer 3 leaf and is connected to each spine switch with 4 port-channels (2 members/port-channel).
- Nexus 3048: One Nexus 3048 switch is configured as a Layer 3 leaf and is connected to two spines with individual routed interfaces and to two spines with a single port-channel (8 members/port-channel).
- Nexus 7000: One Nexus 7000 switch is configured as a Layer 3 leaf and is connected to each spine switch over 10 VRFs with 1 routed sub-interface to each switch for each VRF.

2.5.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- Jumbo MTU: Jumbo MTU is configured as 9216 across the network..
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- BGP: eBGP is configured between the spine and the core, and between the spine and leaf. iBGP is configured among spine switches. IPv4 address families are configured for all BGP peers. Maximum-paths are configured for equal-cost multipath load balancing as 32 for spine and leaf peers for IPv4.
- OSPF: OSPF is used as the IGP to provide reachability for establishing iBGP peering at the spine layer.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. The RP is located at the spine layer.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs located on spine layer.
- VLAN trunking: VLAN trunking is used in the Leaf/Access layers to maintain segregation and security.
- MSTP: Multiple Spanning Tree Protocol is used to prevent Layer 2 loops in the Leaf/Access layer. The spanning tree root is placed on the leaf layer.
- HSRP: HSRP is used as the first hop gateway protocol for host.

- IGMP: IGMP is used by hosts to join multicast groups of interest. IGMP snooping is enabled on all switches in the Leaf/Access layers to prevent flooding of multicast data traffic.
- LACP: LACP is used for link spine and leaf layers to form port-channels across the network.
- CDP/LLDP: CDP is used by default. LLDP is also used for link and neighbor discovery information.
- ECMP: Equal Cost Multipath is used to allow unicast routing over multiple equal cost paths for load sharing.
- Multicast Multipath: Multicast Multipath is used to allow multicast traffic to traverse multiple equal cost paths for load sharing.
- CoPP: CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

2.5.2 Hardware and Software Overview

Platform	Model No.	NVT 3.0
N3548	N3K-C3548P-10G-SUP	6.0(2)A1(1c)
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U1(3)
N7K	N7K-SUP1	6.2(2)

The following line cards are used on the Nexus 7000 devices:

- N7K-F248XP-25

2.6 DC33

2.6.1 Network Logical Topology Design Overview

The topology and test cases validate highly-available data center networks. This is achieved by using Nexus 3048 line of switches in a spine and leaf topology.

2.6.1.1 Description of the Test Network

The spine layer consists of 4 Nexus 3048 switches in a partial mesh topology. The leaf laye r is comprised of Nexus 3048, and Catalyst 6500 switches with ECMP Layer 3 connections to each of the four spine switches.

During NVT Phase 3, the main focus has been the analysis and validation of ECMP deployments for both unicast IPv4 and IPv6 traffic as well as multicast multipath traffic on the Nexus 3048 platform in a spine and leaf topology.

Figure 10 DC33 Topology

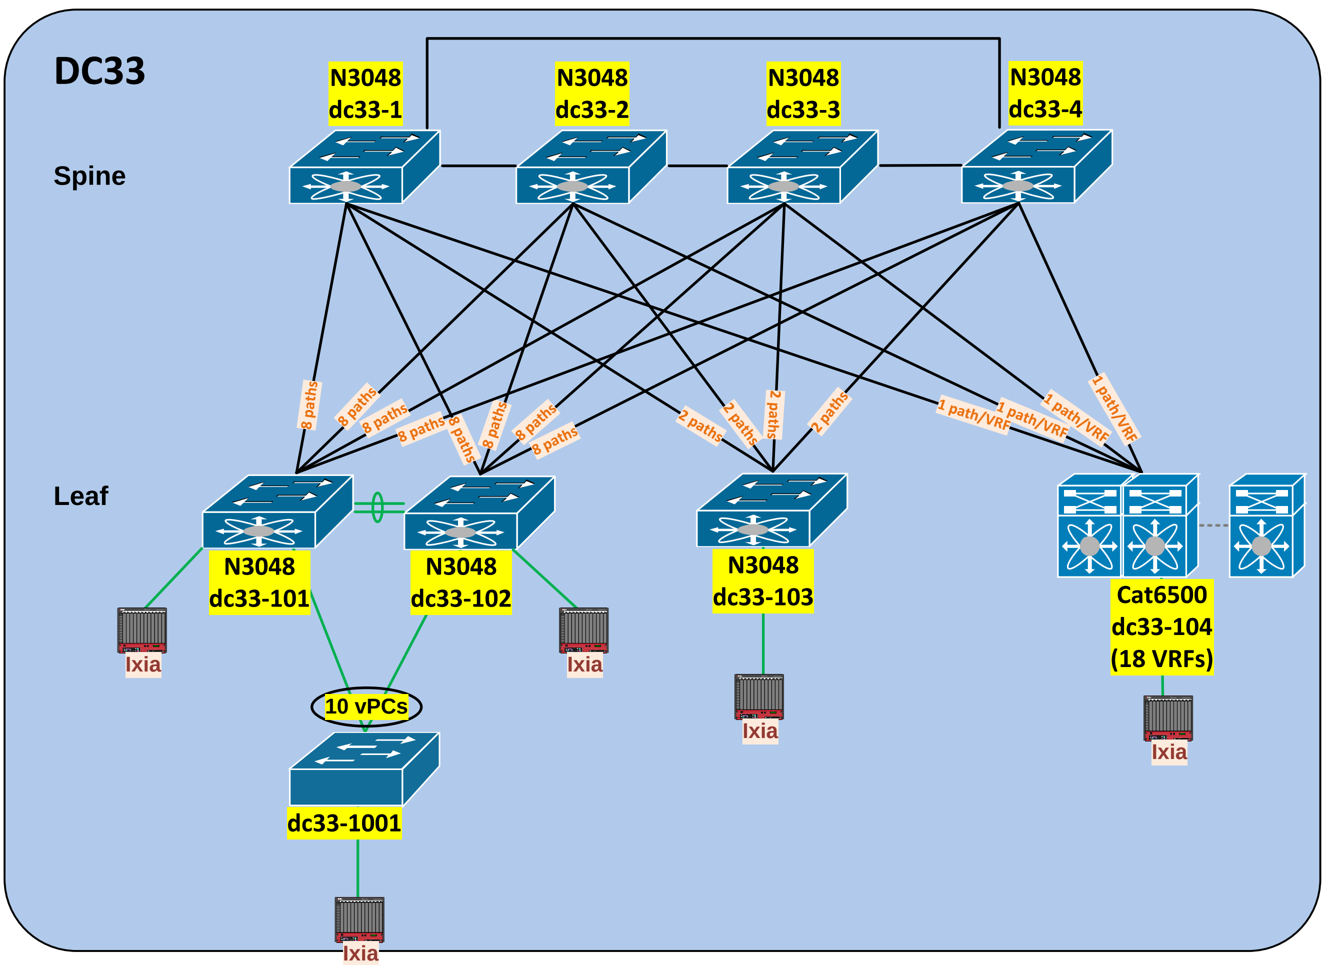

2.6.1.1.1 Spine Layer

The spine layer provides ECMP routing and high bandwidth connectivity between the Leaf/Access layer switches. The spine layer in the test network is implemented using the following platform type:

- Cisco Nexus 3048 Series Switch

2.6.1.1.2 Leaf/Access Layer

The Leaf/Access layers provide connectivity and policy services for locally attached hosts. These leaf switches are implemented as follows:

- Nexus 3048 switch with vPC: Two Nexus 3048 switches are configured as vPC peers. Eight ECMP interfaces are connected to each spine switch while 10 vPC port-channels with one member from each peer are connected to the Layer 2 access switch.
- Nexus 3048 switch: One Nexus 3048 switch is configured as a Layer 3 leaf and is connected to each spine switch with 2 port-channels (4 members/port-channel)..
- Catalyst 6500 switch: One Catalyst 6500 switch is configured as a Layer 3 leaf and is connected to each spine switch over 18 VRFs with 1 routed interface to each switch for each VRF.

2.6.1.2 Test Network Configuration

The following configurations are applied to the test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- Jumbo MTU: Jumbo MTU is configured as 9216 across the network.
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- Dual Stack Interface: All Layer 3 interfaces including routed port, routed port-channel and SVI are configured as IPv4/IPv6 dual stack interfaces.
- BGP: eBGP is configured between the spine and the core, and between the spine and leaf. iBGP is configured among spine switches. IPv4/IPv6 address families are configured for all BGP peers. Maximum-paths are configured for equal-cost multipath load balancing as 64 for both spine and leaf peers for IPv4/IPv6 address families.
- OSPF/OSPFv3: OSPF/OSPFv3 is used as the IGP to provide reachability for establishing iBGP peering at the spine layer.
- PIM-SM: PIM Sparse Mode/PIM Any Source Multicast is deployed across the network to support multicast. The RP is located at the spine layer.
- MSDP Anycast RP: MSDP is deployed to exchange source information between Anycast RPs.
- IGMP: IGMP is used by hosts to join multicast groups of interest. IGMP snooping is enabled on all switches in the Leaf/Access layers to prevent flooding of multicast data traffic.
- Static OIF: Static Outgoing Interfaces (OIFs) are used to statically bind the multicast groups to the outgoing interface (OIF), which is handled by the device hardware. The configuration is applied to all of the Nexus 3048 leafs.
- Static IGMP Snooping: Static IGMP Snooping is used to configure a Layer 2 port of a VLAN as a static member of the multicast groups on both vPC peers.
- vPC: vPC technology is deployed in the Leaf/Access layer.
- VLAN trunking: VLAN trunking is used in the Leaf/Access layer to maintain segregation and security.

- STP: Rapid Spanning Tree Protocol is used to prevent Layer 2 loops in the Leaf/Access layer. The spanning tree root is placed on the leaf layer. Root Guard is configured on the leaf layer to enforce root placement. BPDU Filter, BPDU Guard and PortFast Edge are configured on the access ports towards hosts.
- HSRP/HSRPv6: HSRP/HSRPv6 is used as the first hop gateway protocol for IPv4/IPv6 hosts.
- LACP: LACP is used for link spine and leaf layer to form port-channels across the network.
- CDP/LLDP: CDP is used by default. LLDP is also used for link and neighbor discovery information .
- ECMP: Equal Cost Multipath is used to allow unicast routing over multiple equal cost paths for load sharing.
- Multicast Multipath: Multicast Multipath is used to allow multicast traffic to traverse multiple equal cost paths for load sharing.
- CoPP: CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

2.6.2 Hardware and Software Overview

Platform	Model No.	NVT 3.0
Nexus 3048	N3K-C3048TP-1GE-SUP	6.0.2.U1.3
Catalyst 6500	WS-SUP720-BASE	151-1.SY1

The following line cards are used on the Catalyst 6500 device:

- WS-X6748-GE-TX
- WS-X6708-10GE

2.7 DC36

2.7.1 Network Logical Topology Design Overview

The topology and test cases validate highly-available data center networks to provide dual stack IPv4/IPv6 unicast ECMP. This is achieved by using a combination of Nexus 3048 and Nexus 3064 in a spine and leaf topology.

2.7.1.1 Description of the Test Network

Figure 11 illustrates the DC36 test network topology, consisting of the spine layer and leaf layer. The spine layer consists of four Nexus 3048 switches and two Nexus 3064 switches. The leaf layer is comprised of Nexus 3048, Nexus 3064 and Catalyst 6500 switches with ECMP connections to each of the six spine switches.

During NVT Phase 3, the main focus has been the analysis and validation of ECMP deployments for both unicast IPv4 and IPv6 traffic on the above-mentioned Nexus platforms in a spine and leaf topology.

Figure 11 DC36 Topology

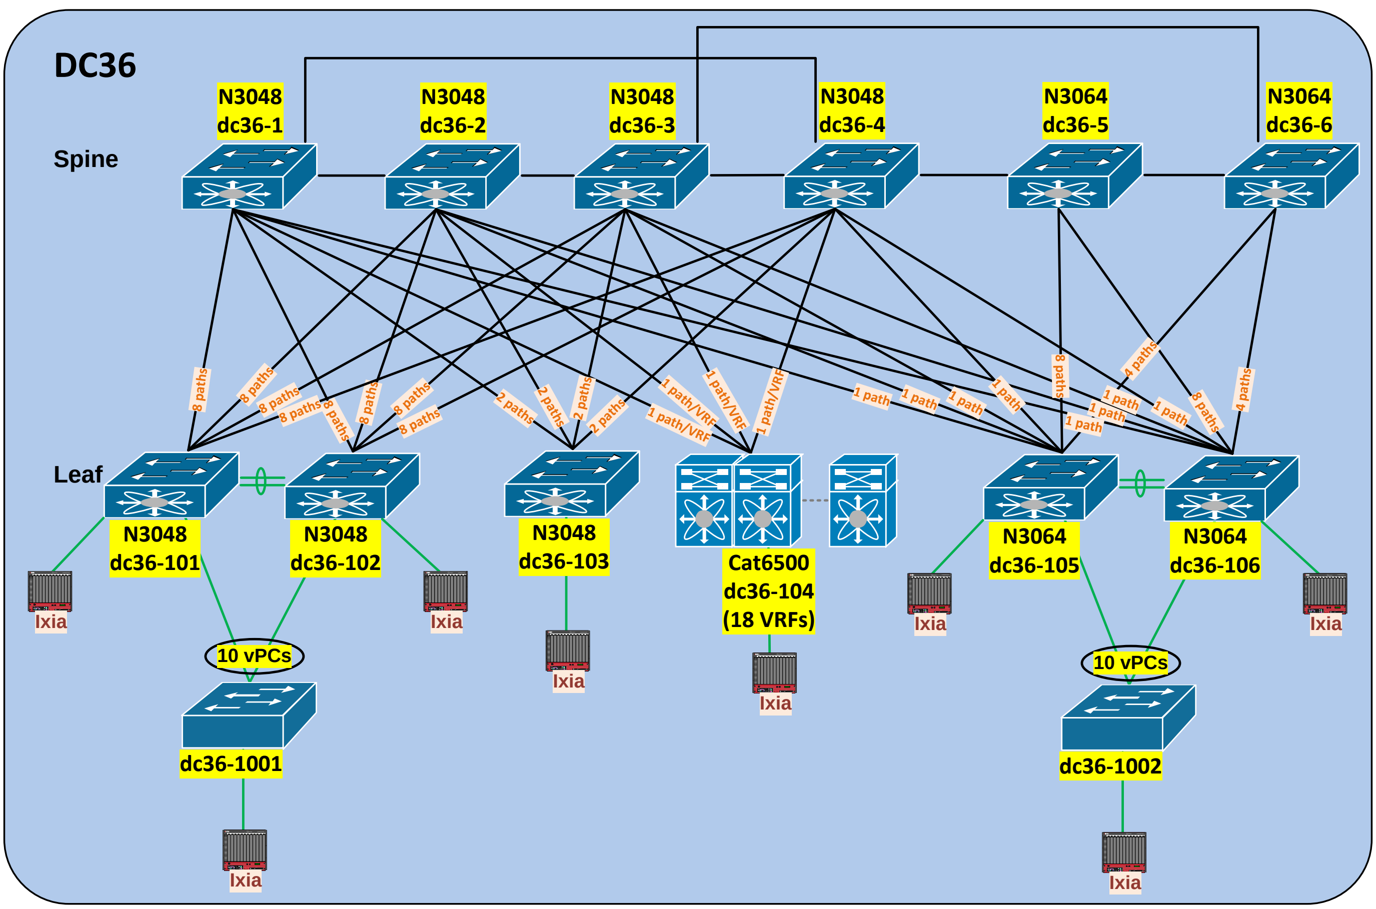

2.7.1.1.1 Spine Layer

The spine layer provides routing and high bandwidth IPv4/IPv6 ECMP connectivity between the spine and leaf layers. The spine layer in the network is implemented using the following two types of Nexus 3000 switches:

- Cisco Nexus 3048 Switch
- Cisco Nexus 3064 Switch

2.7.1.1.2 Leaf/Access Layer

The Leaf/Access layer provides connectivity and policy services for locally attached hosts. These leaf switches are deployed as the following types of devices:

- Nexus 3048 switch with vPC: Two Nexus 3048 switches are configured as vPC peers. Eight ECMP interfaces are connected to each spine switch while 10 vPC port-channels with one member from each peer are connected to the Layer 2 access switch.
- Nexus 3048 switch: One Nexus 3048 switch is configured as a Layer 3 leaf and is connected to each spine switch with 2 port-channels (4 members/port-channel).
- Catalyst 6500 switch: One Catalyst 6500 switch is configured as a Layer 3 leaf and is connected to each spine switch over 18 VRFs with 1 routed sub-interface to each switch for each VRF.
- Nexus 3064 switch with vPC: Two Nexus 3064 switches are configured as vPC peers. Each peer is connected by one individual routed interface to four spine switches, by Eight ECMP interfaces to one spine switch and by 4 ECMP port-channels to the remaining spine switch. 10 vPC port-channels with one member from each peer are connected to the Layer 2 access switch.

2.7.1.2 Test Network Configuration

The following configurations are applied to the DC36 test network:

- Common system control, management and accounting: Common system features like SSH, TACACS+, Syslog, SNMP, NTP, SPAN, DNS and Management VRF are configured.
- Jumbo MTU: Jumbo MTU is configured as 9216 across the network.
- SNMP: SNMP traps are enabled and SNMP scripts are used to collect system information and to monitor potential memory leaks.
- PFC: Priority flow control (refered to as Class-Based Flow Control) is configured to prevent frame loss due to congestion
- Dual Stack Interface: All Layer 3 interfaces including routed port, routed port-channel and SVI are configured as IPv4/IPv6 dual stack interfaces.
- BGP: eBGP is configured between the spine and the core, and between the spine and leaf. iBGP is configured among spine switches. IPv4/IPv6 address families are configured for all BGP peers. Maximum-paths are configured for equal-cost multipath load balancing as 64 for both spine and leaf peers for IPv4/IPv6 address families.
- OSPF/OSPFv3: OSPF/OSPFv3 is used as the IGP to provide reachability for establishing iBGP peering at the spine layer
- vPC: vPC technology is deployed between two leaf layer switches(dc36-101 and dc36-102, dc36-105 and dc36-106).

- VLAN trunking: VLAN trunking is used in the Leaf/Access layer to maintain segregation and security.
- STP: Rapid Spanning Tree Protocol is used to prevent Layer 2 loops in the Leaf/Access layer. The spanning tree root is placed on the leaf layer. Root Guard is configured on the leaf layer to enforce root placement. Portfast edge is configured on the access ports towards hosts.
- HSRP/HSRPv6: HSRP/HSRPv6 is used as the first hop gateway protocol for IPv4/IPv6 hosts.
- LACP: LACP is used for link aggregation to form port-channels across the network and LACP min-link is configured for all port-channels
- UDLD: UDLD aggressive mode is configured across the network to detect and prevent unidirectional links.
- CoPP: CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.
- CDP/LLDP: CDP is used by default. LLDP is also used for link and neighbor discovery information.
- ECMP: Equal Cost Multipath is used to allow unicast routing over multiple equal cost paths for load sharing.

2.7.2 Hardware and Software Overview

Platform	Model No.	NVT 3.0
N3048	N3K-C3048TP-1GE-SUP	6.0.2.U2.1
N3064	N3K-C3064PQ-10GE-SU	6.0.2.U2.1
N7000	N7K-SUP1	6.2.2
C6k	WS-SUP720-BASE	151-1.SY

The following line cards are used on the Nexus 7000 devices:

- N7K-F248XP-25

The following line cards are used on the Catalyst 6500 device:

- WS-X6748-GE-TX
- WS-X6708-10GE

3. NVT Network Implementation and Configuration

3.1 DC1 NVT Network Implementation and Configuration

3.1.1 DC1 Configuration of Platform Specific Features

3.1.1.1 Licensing

Feature-based licenses enable specific feature sets for the physical device. Any feature not included in a license package is bundled with the Cisco NX-OS software.

License Usage on Nexus 7000 in NVT:

```
N7K# show license usage
Feature                          Ins  Lic   Status Expiry Date Comments
                                      Count
--------------------------------------------------------------------------------
MPLS_PKG                          Yes   -    In use Never        -
STORAGE-ENT                       No    -    Unused              -
VDC_LICENSES                      No    0    Unused              -
ENTERPRISE_PKG                    No    -    Unused              -
FCOE-N7K-F132XP                   No    0    Unused              -
FCOE-N7K-F248XP                   No    0    Unused              -
ENHANCED_LAYER2_PKG               Yes   -    Unused Never        -
SCALABLE_SERVICES_PKG             Yes   -    In use Never        -
TRANSPORT_SERVICES_PKG            Yes   -    In use Never        -
LAN_ADVANCED_SERVICES_PKG         Yes   -    In use Never        -
LAN_ENTERPRISE_SERVICES_PKG       Yes   -    In use Never        -
--------------------------------------------------------------------------------
```

License Usage on Nexus 5000 in NVT:

```
dc102-701# show license usage
Feature                          Ins  Lic   Status Expiry Date Comments
                                      Count
--------------------------------------------------------------------------------
FCOE_NPV_PKG                      No    -    Unused              -
FM_SERVER_PKG                     No    -    Unused              -
ENTERPRISE_PKG                    No    -    Unused              -
FC_FEATURES_PKG                   No    -    Unused              -
VMFEX_FEATURE_PKG                 Yes   -    Unused Never        -
ENHANCED_LAYER2_PKG               Yes   -    In use Never        -
LAN_BASE_SERVICES_PKG             Yes   -    In use Never        -
LAN_ENTERPRISE_SERVICES_PKG       No    -    Unused              -
--------------------------------------------------------------------------------
```

License Usage on Nexus 3000 in NVT:

```
dc102-47# show license usage
Feature                          Ins  Lic   Status Expiry Date Comments
                                      Count
--------------------------------------------------------------------------------
LAN_BASE_SERVICES_PKG             Yes   -    In use Never        -
ALGO_BOOST_SERVICES_PKG           No    -    Unused              -
LAN_ENTERPRISE_SERVICES_PKG       No    -    Unused              -
--------------------------------------------------------------------------------
```

3.1.1.2 Out-of-Band Management Network

NVT makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic. Specifically, NVT makes use of the mgmt0 ports on the Nexus devices on a separate management VRF.

Configuration:

```
interface mgmt0
  vrf member management
  ip address 10.1.101.21/16
```

3.1.1.3 Common Configurations

3.1.1.3.1 SSH and TACACS+

SSH is enabled by NVT to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration:

```
feature tacacs+


ip tacacs source-interface mgmt 0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
    use-vrf management

DC5-DC101-5# sh ssh server
ssh version 2 is enabled

DC5-DC101-5# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
interop  pts/0        Feb 10 11:37   .          3995 (taro.interop.cisco.com) session=ssh *
```

3.1.1.3.2 CDP and LLDP

CDP is pervasively used on the NVT testbed for inter-device discovery. LLDP is used where CDP is not supported on host interfaces on Nexus 2000.

```
DC5-DC101-5# sh run cdp all

version 6.2(6)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
cdp format device-id system-name

interface mgmt0
  cdp enable

interface Ethernet1/1
  cdp enable

<TRUNCATED>

interface Ethernet1/52
  cdp enable

DC101-5# sh cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform      Port ID
mgmt-sw1.interop.cisco.com
                    mgmt0          157    R S I     WS-C6509-E    Gig1/1
DC101-6.interop.cisco.com(TBM12450204)
```

```
                    Eth1/1          153     R S s      N7K-C7010      Eth1/1
DC1-3.interop.cisco.com(JAF1529DGCA)
                    Eth1/2          163     R S s      N7K-C7010      Eth2/1
```

```
DC5-DC101-5# sh run lldp all

!Command: show running-config lldp all
!Time: Thu Feb 20 20:30:28 2014

version 6.2(6)
feature lldp

lldp holdtime 120
lldp reinit 2
lldp timer 30
lldp tlv-select port-description
lldp tlv-select system-name
lldp tlv-select system-description
lldp tlv-select system-capabilities
lldp tlv-select management-address
lldp tlv-select dcbxp
lldp tlv-select port-vlan

interface mgmt0
  lldp transmit
  lldp receive

interface Ethernet1/1
  lldp transmit
  lldp receive

<TRUNCATED>

interface Ethernet1/52
  lldp transmit
  lldp receive

DC101-5# sh lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID            Local Intf      Hold-time  Capability  Port ID
Fo1.interop.cisco.comEth101/1/1      120        BR          Gi1/1
Fo1.interop.cisco.comEth101/1/2      120        BR          Gi1/2
Fo1.interop.cisco.comEth101/1/3      120        BR          Gi1/3
Fo1.interop.cisco.comEth101/1/4      120        BR          Gi1/4
Fo1.interop.cisco.comEth101/1/5      120        BR          Gi1/5
Fo1.interop.cisco.comEth101/1/6      120        BR          Gi1/6
```

3.1.1.3.3 Syslog

Syslog is used to record all network events on the DC1 test bed. Whenever possible, NVT uses a separate management VRF for syslog.

Configuration:

```
logging server syslog.interop.cisco.com 5 use-vrf management facility local6

DC1-3 sh logging server
Logging server:               enabled
{syslog.interop.cisco.com}
        server severity:      notifications
        server facility:      local6
        server VRF:           management
```

3.1.1.3.4 SNMP

SNMP is used for system monitoring in NVT. Scripts are used to poll the systems asynchronously during the course of all NVT test execution.

Configuration:

```
DC1-3# show running-config snmp


!Command: show running-config snmp
!Time: Tue Mar 11 21:35:15 2014

version 6.2(6)
power redundancy-mode combined force

snmp-server user admin network-admin auth md5 0xb22e88f075fb25fd56268bcf4628d1a7 priv
0xb22e88f075fb25fd56268bcf4628d1a7 localizedkey
snmp-server user snmpv3 network-admin auth md5 0x46176d732506e914a5ddbf47c4fea173 priv
0x46176d732506e914a5ddbf47c4fea173 localizedkey
snmp-server user ciscoMd5 network-operator auth md5 0x7cc743011a2d8b997d8f99081db6b873 localizedkey
snmp-server user ciscoSha network-operator auth sha 0x545809b573f5dfbab909345cd16ea8543a8d5caa
localizedkey
snmp-server user ciscoMd5Aes network-operator auth md5 0x7cc743011a2d8b997d8f99081db6b873 priv aes-128
0x7cc743011a2d8b997d8f99081db6b873 localizedkey
snmp-server user ciscoMd5Des network-operator auth md5 0x7cc743011a2d8b997d8f99081db6b873 priv
0x7cc743011a2d8b997d8f99081db6b873 localizedkey
snmp-server user ciscoShaAes network-operator auth sha 0x545809b573f5dfbab909345cd16ea8543a8d5caa priv
aes-128 0x545809b573f5dfbab909345cd16ea8543a8d5caa localizedkey
snmp-server user ciscoShaDes network-operator auth sha 0x545809b573f5dfbab909345cd16ea8543a8d5caa priv
0x545809b573f5dfbab909345cd16ea8543a8d5caa localizedkey
snmp-server host 172.28.92.81 traps version 2c public udp-port 2162
snmp-server host 172.28.84.38 traps version 1 public
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-operator
snmp-server community private group network-admin
snmp-server community cisco group network-admin
snmp-server community interop group network-operator
DC1-3# sh snmp trap
--------------------------------------------------------------------------------
Trap type                         Description                   Enabled
--------------------------------------------------------------------------------
entity              : entity_mib_change                           Yes
entity              : entity_module_status_change                 Yes
entity              : entity_power_status_change                  Yes
entity              : entity_module_inserted                      Yes
entity              : entity_module_removed                       Yes
entity              : entity_unrecognised_module                  Yes
entity              : entity_fan_status_change                    Yes
entity              : entity_power_out_change                     Yes
link                : linkDown                                    Yes
link                : linkUp                                      Yes
link                : extended-linkDown                           Yes
link                : extended-linkUp                             Yes
link                : cieLinkDown                                 Yes
link                : cieLinkUp                                   Yes
link                : connUnitPortStatusChange                    Yes
link                : delayed-link-state-change                   Yes
callhome            : event-notify                                No
```

```
callhome             : smtp-send-fail                              No
cfs                  : state-change-notif                          No
cfs                  : merge-failure                               No
rf                   : redundancy_framework                        Yes
aaa                  : server-state-change                         No
license              : notify-license-expiry                       Yes
license              : notify-no-license-for-feature               Yes
license              : notify-licensefile-missing                  Yes
license              : notify-license-expiry-warning               Yes
upgrade              : UpgradeOpNotifyOnCompletion                 Yes
upgrade              : UpgradeJobStatusNotify                      Yes
feature-control      : FeatureOpStatusChange                       No
sysmgr               : cseFailSwCoreNotifyExtended                 No
rmon                 : risingAlarm                                 Yes
rmon                 : fallingAlarm                                Yes
rmon                 : hcRisingAlarm                               Yes
rmon                 : hcFallingAlarm                              Yes
```

3.1.1.3.5 NTP

NTP is used to synchronize the clocks on all NVT devices to provide consistent timestamps on all network logs and events.

Configuration:

```
DC1-3# show running-config ntp

ntp distribute
ntp server 172.28.92.1
ntp commit


DC1-3# show ntp status
Distribution : Enabled
Last operational state: No session

DC1-3# show ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                8   64      377   0.00104 default
```

3.1.1.3.6 SPAN

SPAN has been enabled on NVT switches to provide packet captures to assist in network debugging.

Configuration:

```
monitor session 1
  source interface port-channel36 both
  destination interface Ethernet2/15
  destination interface Ethernet2/32
  no shut

DC1-3# sh monitor session 1
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Po36
    tx            : Po36
    both          : Po36
```

```
source VLANs      :
    rx            :
    tx            :
    both          :
source exception  :
filter VLANs      : filter not specified
destination ports : Eth2/15          Eth2/32

Feature       Enabled   Value   Modules Supported       Modules Not-Supported
-----------------------------------------------------------------------------
MTU-Trunc     No
rate-limit-rx No
rate-limit-tx No
Sampling      No
MCBE          No
L3-TX         -         -       1  2  5  7              -
RB span       No


Legend:
  MCBE  = Multicast Best Effort
  L3-TX = L3 Multicast Egress SPAN
  ExSP-X = Exception Span for type X (L3, FP, or misc)
```

3.1.1.3.7 DNS

DNS has been enabled to provide name lookup in NVT network.

Configuration:

```
ip domain-lookup
ip domain-name interop.cisco.com
ip domain-list cisco.com
ip domain-list interop.cisco.com
ip name-server 172.28.92.9 172.28.92.10

DC1-3# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=1.631 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=1.754 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=1.578 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=1.409 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=1.374 ms

--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.374/1.549/1.754 ms
```

3.1.1.3.8 NDE

NetFlow data export is used to identify packet flows for both ingress and egress IP packets and provide statistics based on these packet flows.

Configuration:

```
DC1-4# show running-config netflow

!Command: show running-config netflow
!Time: Fri Mar  7 12:29:30 2014

version 6.2(6)
feature netflow

flow exporter export-out
```

```
  destination 172.28.92.112
  transport udp 9991
  source loopback0
  version 9
flow exporter export-out1
  transport udp 9995
  version 5
flow record my-flow-record
  description custom-flow-record
  match ipv4 source address
  match ipv4 destination address
  match transport destination-port
  collect counter bytes
  collect counter packets
flow monitor my-flow-monitor
  record my-flow-record
  exporter export-out

interface port-channel1
  ip flow monitor my-flow-monitor input

interface port-channel2
  ip flow monitor my-flow-monitor input

DC5-DC101-5# sh flow monitor my-flow-monitor
Flow Monitor my-flow-monitor:
    Use count: 2
    Flow Record: netflow-original
    Flow Exporter: export-out

DC5-DC101-5# sh flow record my-flow-record
Flow record my-flow-record:
    Description: custom-flow-record
    No. of users: 0
    Template ID: 0
    Fields:
        match ipv4 source address
        match ipv4 destination address
        match transport destination-port
        match interface input
        match interface output
        match flow direction
        collect counter bytes
        collect counter packets

DC5-DC101-5# sh flow exporter export-out
Flow exporter export-out:
    Destination: 172.28.92.112
    VRF: default (1)
    Destination UDP Port 9991
    Source Interface port-channel4 (40.101.3.19)
    Export Version 9
    Exporter Statistics
        Number of Flow Records Exported 57205
        Number of Templates Exported 18
        Number of Export Packets Sent 4394
        Number of Export Bytes Sent 3081412
        Number of Destination Unreachable Events 0
        Number of No Buffer Events 0
        Number of Packets Dropped (No Route to Host) 11
        Number of Packets Dropped (other) 0
        Number of Packets Dropped (LC to RP Error) 0
        Number of Packets Dropped (Output Drops) 1
        Time statistics were last cleared: Never
```

3.1.1.3.9 UDLD

UDLD is used to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a device detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

Configuration:

```
DC1-3# show running-config udld

!Command: show running-config udld
!Time: Fri Mar  7 12:30:50 2014

version 6.2(6)
feature udld

udld aggressive
udld message-time 90

DC1-3# sh udld neighbors
Port            Device Name    Device ID   Port ID         Neighbor State
-------------------------------------------------------------------------
Ethernet1/2     017DF55300     1           Te1/2           bidirectional
Ethernet1/3     TBM12450199    1           Ethernet4/25    bidirectional
Ethernet1/6     017DF55300     1           Te1/4           bidirectional
Ethernet2/1     TBM12450199    1           Ethernet1/2     bidirectional
Ethernet2/2     TBM12450199    1           Ethernet4/1     bidirectional
Ethernet2/3     TBM12450204    1           Ethernet1/3     bidirectional
Ethernet2/4     TBM12450204    1           Ethernet4/1     bidirectional
Ethernet2/6     025B4CF66C0    1           Te2/1/5         bidirectional
Ethernet2/7     025B4CF4B40    1           Te1/1           bidirectional
Ethernet2/8     025B4CF4A0     1           Te1/1           bidirectional
```

3.1.1.3.10 DHCP Relay

DHCP relay is enabled on the aggregation layer to provide IP address services to hypervisors and VMs running on UCS systems.

Configuration:

```
DC5-DC101-5# show running-config dhcp

!Command: show running-config dhcp
!Time: Fri Mar  7 12:32:09 2014

version 6.2(6)
feature dhcp

service dhcp
ip dhcp relay

interface Vlan11
  ip dhcp relay address 94.253.253.2
  ip dhcp relay address 94.1.1.2

DC101-5# sh ip dhcp relay
DHCP relay service is enabled
Insertion of option 82 is disabled
Insertion of VPN suboptions is disabled
Insertion of cisco suboptions is disabled
Global smart-relay is disabled

Smart-relay is enabled on the following interfaces:
-----------------------------------------------------

Subnet-broadcast is enabled on the following interfaces:
```

```
-----------------------------------------------------

Helper addresses are configured on the following interfaces:
 Interface        Relay Address     VRF Name
 -------------    -------------     --------
 Vlan10            94.1.1.2

 Vlan10            94.253.253.2

 Vlan11            94.253.253.2
```

3.1.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

When the switch comes up for the first time, there are multiple CoPP configuration templates that are presented: *strict, moderate, lenient* and *dense*. NVT has chosen the *lenient* template.

Default Lenient CoPP on Nexus 7000 for Software Release 6.2.x as Used in DC1:

```
copp profile lenient

DC5# sh policy-map type control-plane copp-system-p-policy-lenient

  policy-map type control-plane copp-system-p-policy-lenient
    class copp-system-p-class-critical
      set cos 7
      police cir 36000 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-important
      set cos 6
      police cir 1400 kbps bc 1500 ms
        conform transmit violate drop
    class copp-system-p-class-multicast-router
      set cos 6
      police cir 2600 kbps bc 1000 ms
        conform transmit violate drop
    class copp-system-p-class-management
      set cos 2
      police cir 10000 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-multicast-host
      set cos 1
      police cir 1000 kbps bc 1000 ms
        conform transmit violate drop
    class copp-system-p-class-normal
      set cos 1
      police cir 680 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-ndp
      set cos 6
      police cir 680 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-normal-dhcp
      set cos 1
      police cir 1500 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-normal-dhcp-relay-response
      set cos 1
      police cir 1800 kbps bc 750 ms
        conform transmit violate drop
    class copp-system-p-class-redirect
      set cos 1
```

```
    police cir 280 kbps bc 375 ms
      conform transmit violate drop
  class copp-system-p-class-exception
    set cos 1
    police cir 360 kbps bc 375 ms
      conform transmit violate drop
  class copp-system-p-class-monitoring
    set cos 1
    police cir 130 kbps bc 1500 ms
      conform transmit violate drop
  class copp-system-p-class-l2-unpoliced
    police cir 8 gbps bc 5 mbytes
      conform transmit violate transmit
  class copp-system-p-class-undesirable
    set cos 0
    police cir 32 kbps bc 375 ms
      conform drop violate drop
  class copp-system-p-class-fcoe
    set cos 6
    police cir 1060 kbps bc 1500 ms
      conform transmit violate drop
  class copp-system-p-class-l2-default
    police cir 100 kbps bc 375 ms
      conform transmit violate drop
  class class-default
    set cos 0
    police cir 100 kbps bc 250 ms
      conform transmit violate drop
```

Default CoPP on Nexus 5000 as Used in NVT:

```
Dc102-706# show policy-map type control-plane name copp-system-policy-default

policy-map type control-plane copp-system-policy-default
  class copp-system-class-igmp
    police cir 1024 kbps bc 65535 bytes
  class copp-system-class-pim-hello
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-bridging
    police cir 20000 kbps bc 4800000 bytes
  class copp-system-class-arp
    police cir 1024 kbps bc 3600000 bytes
  class copp-system-class-dhcp
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-mgmt
    police cir 12000 kbps bc 4800000 bytes
  class copp-system-class-lacp
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-lldp
    police cir 2048 kbps bc 4800000 bytes
  class copp-system-class-udld
    police cir 2048 kbps bc 4800000 bytes
  class copp-system-class-isis
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-msdp
    police cir 9600 kbps bc 4800000 bytes
  class copp-system-class-cdp
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-fip
    police cir 1024 kbps bc 4800000 bytes
  class copp-system-class-bgp
    police cir 9600 kbps bc 4800000 bytes
  class copp-system-class-eigrp
    police cir 9600 kbps bc 4800000 bytes
  class copp-system-class-exception
    police cir 64 kbps bc 4800000 bytes
  class copp-system-class-glean
```

```
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-hsrp-vrrp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-icmp-echo
      police cir 64 kbps bc 3600000 bytes
    class copp-system-class-ospf
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-pim-register
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-rip
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-l3dest-miss
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-mcast-miss
      police cir 256 kbps bc 3200000 bytes
    class copp-system-class-excp-ip-frag
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-same-if
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-ttl
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-default
      police cir 512 kbps bc 6400000 bytes
```

Default CoPP on Nexus 3000 as Used in DC1:

```
dc102-47# sh policy-map type control-plane expand name copp-system-policy

  policy-map type control-plane copp-system-policy
    class copp-s-selfIp
      police pps 500
    class copp-s-default
      police pps 400
    class copp-s-l2switched
      police pps 200
    class copp-s-ping
      police pps 100
    class copp-s-l3destmiss
      police pps 100
    class copp-s-glean
      police pps 500
    class copp-s-l3mtufail
      police pps 100
    class copp-s-ttl1
      police pps 100
    class copp-s-ipmcmiss
      police pps 400
    class copp-s-l3slowpath
      police pps 100
    class copp-s-dhcpreq
      police pps 300
    class copp-s-dhcpresp
      police pps 300
    class copp-s-dai
      police pps 300
    class copp-s-igmp
      police pps 400
    class copp-s-routingProto2
      police pps 1300
    class copp-s-v6routingProto2
      police pps 1300
    class copp-s-eigrp
      police pps 200
    class copp-s-pimreg
      police pps 200
```

```
    class copp-s-pimautorp
      police pps 200
    class copp-s-routingProto1
      police pps 1000
    class copp-s-arp
      police pps 200
    class copp-s-ptp
      police pps 1000
    class copp-s-bfd
      police pps 350
    class copp-s-bpdu
      police pps 12000
    class copp-icmp
      police pps 200
    class copp-telnet
      police pps 500
    class copp-ssh
      police pps 500
    class copp-snmp
      police pps 500
    class copp-ntp
      police pps 100
    class copp-tacacsradius
      police pps 400
    class copp-stftp
      police pps 400
```

3.1.1.5 Rate Limiters

Rate limiters are an additional set of features on Nexus 7000 to prevent undesirable packets from overwhelming the CPU on the supervisor module.

Default Values:

```
DC1-3# show hardware rate-limiter

Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters


Module: 1
  R-L Class           Config           Allowed         Dropped           Total
 +------------------+--------+---------------+---------------+-----------------+
  L3 mtu                 500               0               0                 0
  L3 ttl                 500               3               0                 3
  L3 control           10000               0               0                 0
  L3 glean               100             143             133               276
  L3 mcast dirconn    Disable
  L3 mcast loc-grp      3000               0               0                 0
  L3 mcast rpf-leak      500               0               0                 0
  L2 storm-ctrl       Disable
  access-list-log        100               0               0                 0
  copy                 30000          147942               0            147942
  receive              30000          341520               0            341520
  L2 port-sec            500               0               0                 0
  L2 mcast-snoop       10000               0               0                 0
  L2 vpc-low            4000               0               0                 0
  L2 l2pt                500               0               0                 0
  f1 rl-1               4500                               0
  f1 rl-2               1000                               0
  f1 rl-3               1000                               0
  f1 rl-4                100                               0
```

```
  f1 rl-5                 1500                             0
  L2 vpc-peer-gw          5000               0             0               0
  L2 lisp-map-cache       5000               0             0               0
  L2 dpss                  100               0             0               0
  L3 glean-fast            100               0             0               0
  L2 otv                   100               0             0               0
  L2 netflow               500               0             0               0

  Port group with configuration same as default configuration
       Eth1/1-4       Eth1/5-8

Module: 2
  R-L Class           Config           Allowed         Dropped           Total
 +------------------+--------+---------------+---------------+-----------------+
  L3 mtu                   500               0             0               0
```

3.1.1.6 VDCs and Resource Allocation

VDCs on the Nexus 7000 are used in the NVT testbed to partition a single physical device into multiple logical devices that provide fault isolation, management isolation, address allocation isolation, service differentiation domains, and adaptive resource management.

```
DC6# show vdc

Switchwide mode is m1 f1 m1xl f2 m2xl f2e

vdc_id  vdc_name                        state               mac                 type        lc
------  --------                        -----               ----------          ---------   ------
1       DC6                             active              00:23:ac:64:bb:c1   Ethernet    m1 f1 m1xl m2xl
2       DC101-6                         active              00:23:ac:64:bb:c2   Ethernet    m1 f1 m1xl m2xl
3       DC102-52                        active              00:23:ac:64:bb:c3   Ethernet    m1 f1 m1xl m2xl
4       DC102-54                        active              00:23:ac:64:bb:c4   Ethernet    m1 f1 m1xl m2xl
```

Resource allocation for VDC's is done from the main VDC based on the requirements. The configuration used in the NVT testbed is as shown below.

The Following Command Can Be Used to Help Estimate the VDC Resource Allocation:

```
N7k# show routing memory estimate routes 68000 nex 2
Shared memory estimates:
  Current max     16 MB;  13743 routes with 16 nhs
          in-use   7 MB;  23290 routes with  2 nhs (average)
  Configured max  16 MB;  13743 routes with 16 nhs
  Estimate        17 MB;  68000 routes with  2 nhs
```

Configuration:

```
vdc DC6 id 1
  limit-resource module-type m1 f1 m1xl m2xl
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  allocate interface Ethernet8/9,Ethernet8/11,Ethernet8/13,Ethernet8/15
  allocate interface Ethernet9/1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
```

```
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc DC101-6 id 2
  limit-resource module-type m1 f1 m1xl m2xl
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  allocate interface Ethernet1/1-7
  allocate interface Ethernet7/7-12
  allocate interface Ethernet8/1-8,Ethernet8/10,Ethernet8/12,Ethernet8/14,Ethernet8/16-32
  allocate interface Ethernet9/41-44
  allocate interface Ethernet10/1-32
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc DC102-52 id 3
  limit-resource module-type m1 f1 m1xl m2xl
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  allocate interface Ethernet4/1-16
  allocate interface Ethernet7/1-6,Ethernet7/13-16
  allocate interface Ethernet9/25-40,Ethernet9/45-48
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc DC102-54 id 4
  limit-resource module-type m1 f1 m1xl m2xl
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  allocate interface Ethernet1/8
  allocate interface Ethernet4/17-32
  allocate interface Ethernet7/17-32
  allocate interface Ethernet9/2-24
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
```

3.1.2 Image Upgrade and Downgrade

NVT makes use of ISSU/D to upgrade/downgrade software images whenever possible.

On the Nexus 7000, to check if the process will be disruptive or not, perform: *show install all impact system <system_image_name> kickstart <kickstart_image_name>:*

```
DC1-3# show install all impact system bootflash:n7000-s1-dk9.6.1.4.bin kickstart  n7000-s1-
kickstart.6.1.4.bin
Installer will perform impact only check. Please wait.

Verifying image bootflash:/n7000-s1-kickstart.6.1.4.bin for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n7000-s1-dk9.6.1.4.bin for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[####################] 100% -- SUCCESS

Extracting "lc1n7k" version from image bootflash:/n7000-s1-dk9.6.1.4.bin.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n7000-s1-dk9.6.1.4.bin.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n7000-s1-dk9.6.1.4.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/n7000-s1-kickstart.6.1.4.bin.
[####################] 100% -- SUCCESS


"Running-config contains configuration that is incompatible with the new image (strict incompatibility).
 Please run 'show incompatibility-all system <image>' command to find out which feature needs to be
disabled.".
Pre-upgrade check failed. Return code 0x40930029 (Current running-config is not supported by new image).
```

Running the command show incompatibility-all system <image-name> will show the incompatible configuration and the necessary steps needed achieve non-disruptive upgrade/downgrade:

```
DC1-3# show incompatibility-all system bootflash:n7000-s1-dk9.6.1.4.bin

Checking incompatible configuration(s) for vdc 'DC1-3':
-------------------------------------------------------
The following configurations on active are incompatible with  the system image
1) Service : confcheck , Capability : CAP_FEATURE_ISSD_PRE621_DENIED
Description : ISSD from current image is not supported.
Capability requirement : STRICT
Enable/Disable command : There is no workaround. If ISSD is required, please
configure the boot variables and reload the switch(disruptive).

2) Service : ipqosmgr , Capability : CAP_FEATURE_IPQOS_DCE_TEMPLATE_8E_4Q4Q
Description : The DCE-QoS template 8e-4q4q exists.
Capability requirement : STRICT
Enable/Disable command : Detach template of type 8e-4q4q from all the interfaces and system qos. Remove
DCE-QoS template 8e-4q4q using the command " clear qos policies 8e-4q4q" from default-vdc at the exec mode


Checking dynamic incompatibilities for vdc 'DC1-3':
---------------------------------------------------
No incompatible configurations

Checking incompatible configuration(s) for vdc 'DC1-3':
------------------------------------------------------
The following configurations on active are incompatible with  the system image
1) Service : confcheck , Capability : CAP_FEATURE_ISSD_PRE621_DENIED
Description : ISSD from current image is not supported.
Capability requirement : STRICT
```

```
Enable/Disable command : There is no workaround. If ISSD is required, please
configure the boot variables and reload the switch(disruptive).

2) Service : ipqosmgr , Capability : CAP_FEATURE_IPQOS_DCE_TEMPLATE_8E_4Q4Q
Description : The DCE-QoS template 8e-4q4q exists.
Capability requirement : STRICT
Enable/Disable command : Detach template of type 8e-4q4q from all the interfaces and system qos. Remove
DCE-QoS template 8e-4q4q using the command " clear qos policies 8e-4q4q" from default-vdc at the exec mode


Checking dynamic incompatibilities for vdc 'DC1-3':
--------------------------------------------------
No incompatible configurations

Checking incompatible configuration(s) for vdc 'DC1-3':
------------------------------------------------------
The following configurations on active are incompatible with  the system image
1) Service : confcheck , Capability : CAP_FEATURE_ISSD_PRE621_DENIED
Description : ISSD from current image is not supported.
Capability requirement : STRICT
Enable/Disable command : There is no workaround. If ISSD is required, please
configure the boot variables and reload the switch(disruptive).

2) Service : ipqosmgr , Capability : CAP_FEATURE_IPQOS_DCE_TEMPLATE_8E_4Q4Q
Description : The DCE-QoS template 8e-4q4q exists.
Capability requirement : STRICT
Enable/Disable command : Detach template of type 8e-4q4q from all the interfaces and system qos. Remove
DCE-QoS template 8e-4q4q using the command " clear qos policies 8e-4q4q" from default-vdc at the exec mode


Checking dynamic incompatibilities for vdc 'DC1-3':
--------------------------------------------------
No incompatible configurations
```

The following caveats apply to ISSU/D:

1. When performing a software release upgrade or downgrade without ISSU in a system with FEX, the host interface configurations on the FEX will be lost after the reload to activate the new image. An extra step is required to reapply the configuration after the FEX module is fully online (*CSCuh58086*). A future FEX pre-provisioning feature will take care of this issue (*CSCuh57942*).
2. When performing ISSU process with OTV configuration, the following error was encountered: Conversion function failed for service "otv" (error-id 0xFFFFFFFF)
   With OTV configured, ISSU will be disruptive and requires shutting down the overlay interface. An enhancement request has been filed to place a configuration compatibility check and throw a message to disallow the procedure until the overlay interface is shutdown (*CSCug73006*).

3.1.3 Routing Design Overview

3.1.3.1 Unicast

3.1.3.1.1 BGP Routing Design

From edge/core switches to public cloud, NVT has enabled eBGP to establish peering between data center autonomous systems and public cloud autonomous systems to exchange routing updates. BGP policy has been applied to the eBGP peering configuration to control route updates between peers.

NVT has configured route maps to filter the redistribution of OSPF routes from DC1 into BGP. The filters are configured based on IP prefix matching.

NSF is a high availability feature on modular switches running NX-OS or IOS with a redundant supervisor. On the Nexus 7000, data packets are forwarded by the hardware forwarding engines on the linecards. These engines are programmed with information learned from the routing control plane running on the supervisors. If the active supervisor were to fail, the forwarding tables on the linecards are preserved. All interface states are also preserved while the standby supervisor takes over active control of the system. This high availability system prevents any drop in traffic during the failure of the active control plane.

BGP graceful restart is a BGP feature that prevents disruption to the control and data plane. It allows for the graceful recovery of BGP sessions after a peer has failed. When combined with the NSF feature, any GR capable peers connected to a switch going through supervisor switchover will continue to forward traffic seamlessly.

Nonstop Forwarding (NSF) and graceful restart (GR) for BGP are enabled by default on NX-OS. SSO/NSF and graceful restart must be explicitly enabled for the system and for BGP, respectively, for Catalyst 6500 and 4500 running IOS.

NVT BGP Configuration:

```
DC1-4# show runn bgp

!Command: show running-config bgp
!Time: Fri Mar  7 12:37:10 2014

version 6.2(6)
feature bgp

router bgp 100
  router-id 40.1.0.15
  graceful-restart stalepath-time 120
  log-neighbor-changes
  address-family ipv4 unicast
    redistribute direct route-map CONN
    redistribute ospf 1 route-map CONN
    maximum-paths 8
    maximum-paths ibgp 8
  neighbor 40.90.1.11 remote-as 100090
    address-family ipv4 unicast
      prefix-list NO_SELF in
  neighbor 40.90.3.13 remote-as 100090
    address-family ipv4 unicast
      prefix-list NO_SELF in
```

3.1.3.1.2 OSPF Routing Design

OSPF has been chosen as the IGP routing protocol for NVT DC1. OSPF has been deployed from Core to Aggregation to L3 Access in NVT data center.

NVT DC1 core switches are configured as backbone Area 0. Each aggregation-access block is configured as a different non-backbone area. The multi-area design reduces computational work for OSPF routers during a topology change.

NVT OSPF Configuration:

```
DC1-4# show running-config ospf

!Command: show running-config ospf
!Time: Fri Mar  7 12:37:44 2014

version 6.2(6)
 feature ospf
router ospf 1
  router-id 40.1.0.15
  redistribute bgp 100 route-map BGPCORE-TO-DC1
  log-adjacency-changes
  timers throttle spf 100 200 500
  timers throttle lsa 50 100 300
  auto-cost reference-bandwidth 1000000
  default-metric 1

interface loopback0
  ip router ospf 2 area 0.0.0.0

interface loopback1
  ip router ospf 2 area 0.0.0.0

interface port-channel15
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 3 a667d47acc18ea6b
  ip router ospf 1 area 0.0.0.101
```

3.1.3.1.2.1 OSPF Router-ID

Each switch in the OSPF routing domain is identified by a Router ID. NVT has configured a loopback interface IP address as OSPF Router-ID for each switch in DC1 to identify each OSPF instance. If there is no OSPF Router-ID, NX-OS will choose the available loopback IP address as OSPF Router-ID and if there is no loopback address available, NX-OS will choose the highest interface IP address as OSPF Router-ID. If the interface IP address is used as the OSPF Router-ID, it will cause routing re-convergence when that interface goes down.

Router-ID is configured per OSPF process instance. NVT testing only creates one instance per VDC.

To Verify the OSPF Router ID:

```
DC1-3# show ip ospf

 Routing Process 1 with ID 40.1.0.15 VRF default
 Routing Process Instance Number 1

DC1-3# sh ip ospf neighbors
 OSPF Process ID 1 VRF default
 Total number of neighbors: 14
 Neighbor ID     Pri State            Up Time  Address         Interface
 40.101.0.19       1 FULL/DR          16:33:16 40.101.1.19     Po15
 40.101.0.21       1 FULL/BDR         01:09:37 40.101.2.21     Po16
 40.102.0.51       1 FULL/DR          1d03h    40.102.1.51     Po21
 40.102.0.52       1 FULL/DR          1d03h    40.102.2.52     Po22
 40.102.0.53       1 FULL/DR          1d03h    40.102.5.53     Po23
 40.102.0.54       1 FULL/DR          1d03h    40.102.6.54     Po24
```

3.1.3.1.2.2 OSPF Reference Bandwidth

The default OSPF Auto-Cost reference bandwidth for calculating OSPF metric is 40Gbps for NX-OS and 100Mbps for IOS. The reference bandwidth should be configured to be the same across the entire network; NVT has configured 100Gbps as the reference bandwidth.

To Verify OSPF Reference Bandwidth:

```
DC1-3# show ip ospf

 Routing Process 1 with ID 40.1.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 This router is an area border and autonomous system boundary.
 Redistributing External Routes from
   bgp-100
 Administrative distance 110
 Reference Bandwidth is 1000000 Mbps
```

3.1.3.1.2.3 OSPF Network Type

NVT has configured point-to-point OSPF Network Type on all interfaces between the core and aggregation switches. It removes the OSPF designated router and backup designated router (DR/BDR) election and reduces the OSPF neighbor adjacency negotiation process.

To Verify OSPF Point-to-Point OSPF Network:

```
DC1-4# show ip ospf interface po15
 port-channel15 is up, line protocol is up
    IP address 40.101.3.17/24, Process ID 1 VRF default, area 0.0.0.101
    Enabled by interface configuration
    State BDR, Network type BROADCAST, cost 50
    Index 6, Transmit delay 1 sec, Router Priority 1
    Designated Router ID: 40.101.0.19, address: 40.101.3.19
    Backup Designated Router ID: 40.1.0.17, address: 40.101.3.17
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:07
    No authentication
    Number of opaque link LSAs: 0, checksum sum 0

DC1-4# sh ip ospf neighbors
 OSPF Process ID 1 VRF default
 Total number of neighbors: 11
 Neighbor ID     Pri State            Up Time  Address         Interface
 40.101.0.19       1 FULL/DR          2d03h    40.101.3.19     Po15
 40.101.0.21       1 FULL/DR          2d03h    40.101.4.21     Po16
```

3.1.3.1.2.4 OSPF Authentication

Cisco NX-OS supports two authentication methods, simple password authentication and MD5 authentication digest. Authentication can be configured for an OSPFv2 area or per interface.

NVT has configured MD5 authentication for each interface.

To Verify OSPF Authentication:

```
DC1-3# show ip ospf interface p15
 port-channel15 is up, line protocol is up
```

```
    IP address 40.101.1.15/24, Process ID 1 VRF default, area 0.0.0.101
    Enabled by interface configuration
    State DR, Network type BROADCAST, cost 50
    Index 6, Transmit delay 1 sec, Router Priority 1
    Designated Router ID: 40.1.0.15, address: 40.101.1.15
    Backup Designated Router ID: 40.101.0.19, address: 40.101.1.19
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:09
    No authentication
    Number of opaque link LSAs: 0, checksum sum 0
    Message-digest authentication, using key id 1
    Number of opaque link LSAs: 0, checksum sum 0
```

3.1.3.1.2.5 Route Redistribution

Route redistribution is configured on the Core/Edge switches for DC1 to learn routes from BGP. Route maps are used to control which external routes are redistributed. NVT has configured IP prefix-list to filter IP addresses.

3.1.3.1.2.6 OSPF High Availability and Graceful Restart

Cisco provides multilevel high-availability architecture for OSPF: Non Stop Routing (NSR) and Graceful Restart (GR) with NSF.

With NSR, OSPF preserves the running state of the protocol data and sessions in persistent memory. If the OSPF application fails or needs to be restarted for any reason, it will restart from the preserved state to ensure that there is no disruption seen by any of its OSPF peers. The internal applications that manage the routing table and hardware forwarding tables will also not experience any failure, allowing for non-disruptive OSPF process restarts.

OSPF GR and NSF allow for non-disruptive failure of the supervisor on Cisco modular switches. On the Nexus 7000, the hardware routing engines are programed per linecard. On active supervisor failure, the forwarding tables on the linecards are preserved while the standby supervisor takes over active control of the system. There is no disruption to packet forwarding during this process. GR prevents OSPF peers from restarting during a supervisor failure; thus, preserving their packet forwarding states. The combination of OSPF GR and SSO/NSF allows the entire network to continue operating seamlessly during a supervisor failure.

OSPF NSR and graceful restart are enabled by default on NX-OS. SSO/NSF and graceful restart must be explicitly enabled for the system and for OSPF, respectively, for Catalyst 6500 and 4500 running IOS.

To Verify OSPF Graceful Restart:

```
DC1-3# show ip ospf

 Routing Process 1 with ID 40.1.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
```

3.1.3.1.2.7 Passive Interfaces

All servers/hosts facing SVIs (Switched Virtual Interfaces) are configured as OSPF passive interfaces. This is to ensure that server farm subnets are advertised into OSPF, while preventing the formation of unnecessary OSPF adjacencies through the access layer.

To Verify OSPF Passive Interface:

```
DC101-5# sh ip ospf interface vlan 12
 Vlan12 is up, line protocol is up
    IP address 101.12.0.19/16, Process ID 2 VRF default, area 0.0.0.101
    Enabled by interface configuration
    State DR, Network type BROADCAST, cost 1000
    Index 9, Passive interface
```

3.1.3.1.2.8 OSPF Timers and Optimization

NVT has kept the OSPF hello/hold timers at their default values. This allows other resilience features such as SSO/NSF to provide high availability. BFD should be used for networks where fast peer failure detection is desired. NVT has left all OSPF hello/hold timers as default for DC1.

To Verify OSPF Timers and Optimization:

```
DC1-3# show ip ospf

 Routing Process 1 with ID 40.1.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 This router is an area border and autonomous system boundary.
 Redistributing External Routes from
   bgp-100
 Administrative distance 110
 Reference Bandwidth is 1000000 Mbps
 SPF throttling delay time of 100.000 msecs,
   SPF throttling hold time of 200.000 msecs,
   SPF throttling maximum wait time of 500.000 msecs
 LSA throttling start time of 50.000 msecs,
   LSA throttling hold interval of 100.000 msecs,
   LSA throttling maximum wait time of 300.000 msecs
 Minimum LSA arrival 1000.000 msec
 LSA group pacing timer 10 secs
 Maximum paths to destination 8
 Number of external LSAs 77, checksum sum 0x2c86b2
  Number of opaque AS LSAs 0, checksum sum 0
 Number of areas is 8, 8 normal, 0 stub, 0 nssa
 Number of active areas is 8, 8 normal, 0 stub, 0 nssa
 Install discard route for summarized external routes.
 Install discard route for summarized internal routes.
   Area BACKBONE(0.0.0.0) (Inactive)
        Area has existed for 2w1d
        Interfaces in this area: 6 Active interfaces: 6
        Passive interfaces: 0  Loopback interfaces: 3
        No authentication available
        SPF calculation has run 3483 times
         Last SPF ran for 0.000644s
        Area ranges are
        Number of LSAs: 442, checksum sum 0xccce3d
   Area (0.0.0.101)
        Area has existed for 2w1d
        Interfaces in this area: 3 Active interfaces: 3
```

```
        Passive interfaces: 0  Loopback interfaces: 0
        No authentication available
        SPF calculation has run 3483 times

DC1-3# show ip ospf interface port-channel 15
 port-channel15 is up, line protocol is up
    IP address 40.101.1.15/24, Process ID 1 VRF default, area 0.0.0.101
    Enabled by interface configuration
    State DR, Network type BROADCAST, cost 50
    Index 6, Transmit delay 1 sec, Router Priority 1
    Designated Router ID: 40.1.0.15, address: 40.101.1.15
    Backup Designated Router ID: 40.101.0.19, address: 40.101.1.19
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:05
    No authentication
    Number of opaque link LSAs: 0, checksum sum 0
```

3.1.3.2 Unicast Forwarding Verification

On NX-OS platforms, routing is performed using hardware forwarding engines. The following sequence of commands illustrates verification of the programming of a host on a directly connected subnet on the Nexus 7000.

This Switch is the Authoritative Router for a Directly Connected Subnet on VLAN 11: 10.11.0.0/16:

```
DC101-6# show running-config interface vlan 11

!Command: show running-config interface Vlan11
!Time: Thu Feb 13 19:42:58 2014

version 6.2(6)

interface Vlan11
  no ip redirects
  ip address 101.11.0.21/16
  ip address 101.111.0.21/16 secondary
  ipv6 address 2001:1:101:11::21/64
  ip router ospf 1 area 0.0.0.101
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 101.11.0.1


hsrp 2
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 101.111.0.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 2001:1:101:11::1
  ip dhcp relay address 94.253.253.2
  ip dhcp relay address 94.1.1.2
  no shutdown
```

The Host 101.11.7.1 has been Learned via ARP on this Subnet:

```
DC101-6# show ip arp 101.11.7.1
```

```
Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
101.11.7.1      00:04:45  0065.0b07.0100  Vlan11
```

On NX-OS, "show ip route" will also Show Directly Connected Hosts as /32 Routes:

```
DC101-6# sh ip route 101.11.7.1

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

101.11.7.1/32, ubest/mbest: 1/0, attached
    *via 101.11.7.1, Vlan11, [250/0], 00:02:43, am
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
DC101-6# sh ip adjacency 101.11.7.1

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
101.11.7.1      0065.0b07.0100  50   arp        Vlan11
```

Find the PO Interface on which this MAC Address is Learnt:

```
DC101-6# sh mac address-table address 0065.0b07.0100
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0065.0b07.0100    dynamic   0          F    F  Po7
```

Display PO7 Member Interface with Module Information:

```
DC101-6# sh port-channel summary | in Po7
7     Po7(SU)     Eth      LACP      Eth8/1(P)
```

Display Adjacency Index for this Route in Hardware Table:

```
DC101-6# sh system internal forwarding ip route 101.11.7.1 module 8
Routes for table default/base

----+---------------------+----------+----------+------+-----------
Dev | Prefix              | PfxIndex | AdjIndex | LIFB | LIF
----+---------------------+----------+----------+------+-----------
 1    101.11.7.1/32         0x4202     0x4300f       0   0x7b
```

Display DMAC Entry Programmed in Adjacency Table:

```
DC101-6# sh system internal forwarding adjacency module 8 entry 0x4300f  detail
 Device: 1   Index: 0x4300f   DMAC: 0065.0b07.0100 SMAC: 0023.ac64.bbc2
            LIF: 0x7b (Vlan11) DI: 0x0     ccc: 4   L2_FWD: NO  RDT: NO
            packets: 0   bytes: 549755813888zone enforce: 0
```

Display Allocated Bridge Domain Matches in the Hardware Table:

```
DC101-6# sh vlan internal bd-info vlan-to-bd 11

VDC Id  Vlan Id  BD Id
------  -------  -------
2         11       123
```

Display LTL Entry for this MAC Address Associated with the Bridge Domain:

```
DC101-6# sh hardware mac address-table 8 vlan 11
FE | Valid| PI|  BD  |       MAC       |  Index| Stat| SW   | Modi| Age| Tmr| GM| Sec| TR| NT| RM| RMA| Cap| Fld|Always
   |      |   |      |                 |       |  ic |      | fied|Byte| Sel|   | ure| AP| FY|   |    |TURE|    | Learn
---+------+---+------+-----------------+-------+-----+------+-----+----+----+---+----+---+---+---+----+----+----+------
0    1      1   123   0065.0b07.0100  0x00a2b   0   0x003    0     247   1    0    0   0   0   0    0    0    0    0
```

Display DMAC Sent to LTL Index for PO7:

```
DC101-6# sh system internal pixm info ltl 0x00a2b

PC_TYPE    PORT    LTL      RES_ID       LTL_FLAG      CB_FLAG      MEMB_CNT
---------------------------------------------------------------------------
Normal    Po7     0x0a2b    0x16000006   0x00000000    0x00000002   1
```

3.1.3.3 Multicast Routing Design

Multicast routing has been enabled across the entire NVT network on DC1. On NX-OS, multicast routing is enabled by default, while it needs to be explicitly enabled on IOS.

NVT Multicast Configuration:

```
feature pim
ip pim rp-address 40.1.50.1 group-list 230.2.0.0/16
ip pim rp-address 40.1.50.1 group-list 239.1.1.1/32
ip pim send-rp-announce loopback1 group-list 230.201.0.0/16
ip pim send-rp-discovery loopback1
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt

interface loopback1
  ip address 40.101.51.1/32
  ip router ospf 2 area 0.0.0.201
  ip pim sparse-mode
```

```
feature msdp
ip msdp originator-id loopback0
ip msdp peer 40.101.0.19 connect-source loopback0

interface loopback0
  ip address 40.101.0.21/32
```

```
  ip router ospf 1 area 0.0.0.101
  ip pim sparse-mode
```

3.1.3.3.1 PIM-ASM Rendezvous Point

The NVT topology relies heavily on vPC and as such PIM Sparse Mode has been configured as the protocol of choice for multicast routing. NX-OS does not support PIM SSM and PIM Bidir operating over vPC.

3.1.3.3.1.1 Auto-RP

The NVT testbed is designed to have an RP for each POD in DC1 data centers to support the groups sourced from that particular POD. Each RP is configured on the aggregation switches for a given POD. NVT makes use of Auto-RP to automate distribution of RP information in the network.

To Verify PIM RP:

```
DC101-6# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.107.51.1, uptime: 22:40:53, expires: 00:02:38
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
RP: 40.1.50.1, (0), uptime: 22:50:21, expires: 00:02:38 (A),
  priority: 0, RP-source: 40.107.51.1 (A), (local), group ranges:
      239.1.1.1/32   230.2.0.0/16
RP: 40.101.51.1*, (0), uptime: 22:48:58, expires: 00:02:38,
  priority: 0, RP-source: 40.107.51.1 (A), group ranges:
      230.201.0.0/16

DC101-6# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.2.0.0/16       ASM       40.1.50.1        -
230.201.0.0/16     ASM       40.101.51.1      -
239.1.1.1/32       ASM       40.1.50.1        -
```

3.1.3.3.1.1.1 Auto-RP Forward Listen

NVT has enabled the Auto-RP listening and forwarding feature so that the Auto-RP mechanism can dynamically inform routers in the PIM domain of the group-to-RP mapping since PIM dense mode is not supported on NX-OS. By default, listening or forwarding of Auto-RP messages is not enabled on NX-OS.

3.1.3.3.1.2 Static RP

The NVT network is configured with a backup RP on the core routers for all groups in the network. This RP is statically configured on all routers in the network. Auto-RP takes precedence over static RP.

To Verify PIM RP:

```
DC101-6# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.107.51.1, uptime: 00:53:17, expires: 00:02:13
BSR RP Candidate policy: None
```

```
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.1.50.1, (0), uptime: 01:12:54, expires: 00:02:13 (A),
  priority: 0, RP-source: 40.107.51.1 (A), (local), group ranges:
      230.1.0.0/16
RP: 40.101.51.1*, (0), uptime: 01:09:19, expires: 00:02:13,
  priority: 0, RP-source: 40.107.51.1 (A), group ranges:
      230.101.0.0/16
RP: 40.102.51.1, (0), uptime: 01:09:19, expires: 00:0

DC101-6# show ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.1.0.0/16       ASM       40.1.50.1        -
230.101.0.0/16     ASM       40.101.51.1      -
230.102.0.0/16     ASM       40.102.51.1      -
```

3.1.3.3.1.3 Anycast RP with MSDP

NVT has configured Anycast RP with MSDP within each POD at the aggregation layer. NVT has also configured Anycast RP with MSDP among the core switches.

NVT Anycast RP and MSDP Configuration:

```
N7K aggregation 1:

!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list
230.101.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  ip address 40.101.51.1/32
  ip router ospf 1 area 0.0.0.101
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.101.0.21 connect-source loopback0
interface loopback0
  ip address 40.101.0.19/32
  ip router ospf 2 area 0.0.0.101
  ip pim sparse-mode
```

```
N7K aggregation 2:

!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list
230.101.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  ip address 40.101.51.1/32
  ip router ospf 2 area 0.0.0.101
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.101.0.19 connect-source loopback0
interface loopback0
  ip address 40.101.0.21/32
  ip router ospf 2 area 0.0.0.101
  ip pim sparse-mode
```

To Verify MSDP Peer and SA_Cache:

```
DC101-5# sh ip msdp sa-cache
MSDP SA Route Cache for VRF "default" - 100 entries
Source          Group           RP              ASN         Uptime
201.11.7.1      230.201.0.1     40.101.0.21     0           16:23:37
201.11.7.2      230.201.0.1     40.101.0.21     0           16:12:19
201.11.7.3      230.201.0.1     40.101.0.21     0           16:23:37
201.11.7.4      230.201.0.1     40.101.0.21     0           16:12:19
201.11.7.5      230.201.0.1     40.101.0.21     0           16:23:37
201.11.7.6      230.201.0.1     40.101.0.21     0           16:12:19

DC101-5# sh ip msdp sum
MSDP Peer Status Summary for VRF "default"
Local ASN: 0, originator-id: 40.101.0.19

Number of configured peers:  1
```

```
Number of established peers: 1
Number of shutdown peers:    0

Peer            Peer        Connection      Uptime/   Last msg  (S,G)s
Address         ASN         State           Downtime  Received  Received
40.101.0.21   0           Established     17:34:46  00:00:35  100
```

3.1.3.3.2 PIM SPT-Threshold

NVT has enabled *ip pim spt-threshold infinity* on the last hop non-vPC PIM routers to decrease the multicast entries hardware usage across the network. Nexus 7000 vPC does not support PIM spt-threshold configuration.

3.1.3.3.3 Multicast Multipath

Cisco NX-OS Multicast Multipath is enabled by default and the load sharing selection algorithm is based on the source and group addresses. On Cisco IOS, Multicast Multipath is disabled by default. When multipath is enabled on Cisco IOS, the default load sharing selection algorithm is source-based. The algorithm on IOS can be configured to match the behavior on NX-OS with the command "*ip multicast multipath s-g-hash basic".*

NVT has enabled multicast multipath across the whole network on all applicable platforms.

3.1.3.4 Multicast Forwarding Verification

The following sequence of commands illustrates the verification of the Cisco NX-OS multicast L2 and L3 forwarding.

Displays a Specific Multicast Route 230.101.0.1 with Incoming Interface Information:

```
DC6-DC101-6# show ip mroute 230.102.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.102.0.1/32), uptime: 00:21:33, igmp ip pim
  Incoming interface: port-channel4, RPF nbr: 40.101.4.17
  Outgoing interface list: (count: 20)
    Vlan2010, uptime: 00:21:28, igmp
    Vlan2004, uptime: 00:21:28, igmp
    Vlan17, uptime: 00:21:28, igmp
    Vlan16, uptime: 00:21:28, igmp
    Vlan15, uptime: 00:21:28, igmp
    Vlan14, uptime: 00:21:28, igmp
    Vlan2009, uptime: 00:21:33, igmp
    Vlan2008, uptime: 00:21:33, igmp
    Vlan2007, uptime: 00:21:33, igmp
    Vlan2006, uptime: 00:21:33, igmp
    Vlan2005, uptime: 00:21:33, igmp
    Vlan2003, uptime: 00:21:33, igmp
    Vlan2002, uptime: 00:21:33, igmp
    Vlan2001, uptime: 00:21:33, igmp
    Vlan20, uptime: 00:21:33, igmp
    Vlan19, uptime: 00:21:33, igmp
    Vlan18, uptime: 00:21:33, igmp
    Vlan13, uptime: 00:21:33, igmp
    Vlan12, uptime: 00:21:33, igmp
```

```
    Vlan11, uptime: 00:21:33, igmp

(102.11.17.1/32, 230.102.0.1/32), uptime: 00:08:48, ip mrib pim
  Incoming interface: port-channel4, RPF nbr: 40.101.4.17
  Outgoing interface list: (count: 20)
    Vlan2010, uptime: 00:08:48, mrib
    Vlan2009, uptime: 00:08:48, mrib
    Vlan2008, uptime: 00:08:48, mrib
    Vlan2007, uptime: 00:08:48, mrib
    Vlan2006, uptime: 00:08:48, mrib
    Vlan2005, uptime: 00:08:48, mrib
    Vlan2004, uptime: 00:08:48, mrib
    Vlan2003, uptime: 00:08:48, mrib
    Vlan2002, uptime: 00:08:48, mrib
    Vlan2001, uptime: 00:08:48, mrib
    Vlan20, uptime: 00:08:48, mrib
    Vlan19, uptime: 00:08:48, mrib
    Vlan18, uptime: 00:08:48, mrib
    Vlan17, uptime: 00:08:48, mrib
    Vlan16, uptime: 00:08:48, mrib
    Vlan15, uptime: 00:08:48, mrib
    Vlan14, uptime: 00:08:48, mrib
    Vlan13, uptime: 00:08:48, mrib
    Vlan12, uptime: 00:08:48, mrib
    Vlan11, uptime: 00:08:48, mrib
```

Display DR Information for Interface Vlan11:

```
DC101-6# sh ip pim interface brief
PIM Interface Status for VRF "default"
Interface            IP Address      PIM DR Address  Neighbor  Border
                                                     Count     Interface
Vlan11               101.11.0.21     101.11.0.21     1         no
Vlan2001             101.201.0.21    101.201.0.21    1         no
port-channel3        40.101.2.21     40.101.2.21     1         no
port-channel4        40.101.4.21     40.101.4.21     1         no
port-channel9        40.101.6.21     0.0.0.0         0         no
loopback0            40.101.0.21     40.101.0.21     0         no
loopback1            40.101.51.1     40.101.51.1     0         no
```

Displays Mroute RPF Interface and Forwarding Counters in L3 Hardware Table:

```
DC6-DC101-6# sh forwarding multicast route group 230.102.0.1 source 102.11.17.1

slot  1
=======


  (102.11.17.1/32, 230.102.0.1/32), RPF Interface: port-channel4, flags:
    Received Packets: 13820 Bytes: 1326720
    Number of Outgoing Interfaces: 20
    Outgoing Interface List Index: 6
      Vlan11 Outgoing Packets:35186683 Bytes:3377921568
      Vlan12 Outgoing Packets:26230679 Bytes:2518145184
      Vlan13 Outgoing Packets:26230679 Bytes:2518145184
      Vlan14 Outgoing Packets:26230679 Bytes:2518145184
      Vlan15 Outgoing Packets:26230679 Bytes:2518145184
      Vlan16 Outgoing Packets:26230679 Bytes:2518145184
      Vlan17 Outgoing Packets:26230679 Bytes:2518145184
      Vlan18 Outgoing Packets:26230679 Bytes:2518145184
      Vlan19 Outgoing Packets:26230679 Bytes:2518145184
      Vlan20 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2001 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2002 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2003 Outgoing Packets:39346009 Bytes:3777216864
      Vlan2004 Outgoing Packets:26230679 Bytes:2518145184
```

```
      Vlan2005 Outgoing Packets:39346028 Bytes:3777218688
      Vlan2006 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2007 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2008 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2009 Outgoing Packets:26230679 Bytes:2518145184
      Vlan2010 Outgoing Packets:26230679 Bytes:2518145184
```

Displays the Multicast Routing Table with Packet Counts and Bit Rates for All Sources:

```
DC6-DC101-6# sh ip mroute 230.102.0.1 summary
IP Multicast Routing Table for VRF "default"

Total number of routes: 1018
Total number of (*,G) routes: 17
Total number of (S,G) routes: 1000
Total number of (*,G-prefix) routes: 1
Group count: 17, rough average sources per group: 58.8

Group: 230.102.0.1/32, Source count: 400
Source          packets      bytes           aps   pps       bit-rate     oifs
(*,G)           65428        5363312         81    0         0.000   bps  20
102.11.17.1     14727        1207606         81    20        13.186  kbps 20
102.11.17.2     14629        1199578         82    20        13.186  kbps 20
102.11.17.3     14689        1204482         81    20        13.186  kbps 20
```

Display IGMP Snooping Groups Information:

```
DC101-6# sh ip igmp snooping groups 230.102.0.1 vlan 11
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
11    230.102.0.1        v2   D     Po7 Po8
```

Displays Detected Multicast Routers for VLAN:

```
DC101-6# sh ip igmp snooping mrouter vlan 11

Type: S - Static, D - Dynamic, V - vPC Peer Link
      I - Internal, F - FabricPath core port
      U - User Configured
Vlan  Router-port   Type      Uptime      Expires
11    Vlan11        I         02:04:10    never
11    Po5           SVD       01:38:00    00:04:54
```

Displays IGMP Snooping Querier Information for VLAN:

```
DC101-6# sh ip igmp snooping querier vlan 11
Vlan  IP Address       Version    Expires     Port
11    101.11.0.19      v2         00:03:51    port-channel5
```

Display L2 MFDM Software Entries for Group/VLAN 11:

```
DC6-DC101-6# sh forwarding distribution ip igmp snooping vlan 11 group 230.102.0.1
Vlan: 11, Group: 230.102.0.1, Source: 0.0.0.0
  Outgoing Interface List Index: 76
  Reference Count: 12
  Platform Index: 0x7fc7
  Number of Outgoing Interfaces: 4
    port-channel5
    port-channel7
    port-channel8
```

```
    Replicator1/2/5

Vlan: 11, Aggregated Group: 230.102.0.1, Source: 0.0.0.0
  Outgoing Interface List Index: 82
  Reference Count: 120
  Platform Index: 0x7fc1
  Number of Outgoing Interfaces: 3
    port-channel5
    port-channel7
    port-channel8
```

Display L2 Hardware Entry for Group/VLAN:

```
DC6-DC101-6# sh system internal ip igmp snooping vlan 11 group 230.102.0.1 module 8

VDC: 2
Lookup Mode : IP

Vlan   Group           Source           Epoch    RID   DTL      hwptr     Ref#    GS Entry#
11    230.102.0.1                       1        76    0x7fc7   0x4a3f    1       0
```

Display DTL Sent to LTL Index for PO7:

```
DC6-DC101-6# sh system internal pixm info ltl 0x7fc7
MCAST LTLs allocated for VDC:2
============================================
LTL    IFIDX/RID   LTL_FLAG CB_FLAG
0x7fc7 0x0000004c 0x00     0x0002

mi  | v4_fpoe | v5_fpoe | clp_v4_l2 | clp_v5_l2 | clp20_v4_l3 | clp_cr_v4_l3 | flag | proxy_if_index
0xd | 0xb     | 0x0     | 0x7       | 0x0       | 0x0         | 0x47         | 0x0  | repl1/2/5

Member info
------------------
IFIDX           LTL
---------------------------------
Po8                0x0a2d
Po7                0x0a2b
Po5                0x0a29
```

3.1.4 Layer-2/ Layer-3 Aggregation/Access Layer Network Design Overview

3.1.4.1 vPC

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco NX-OS switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

Figure 12 Creating a Single Logical Node through vPC (virtual PortChannel) Technology

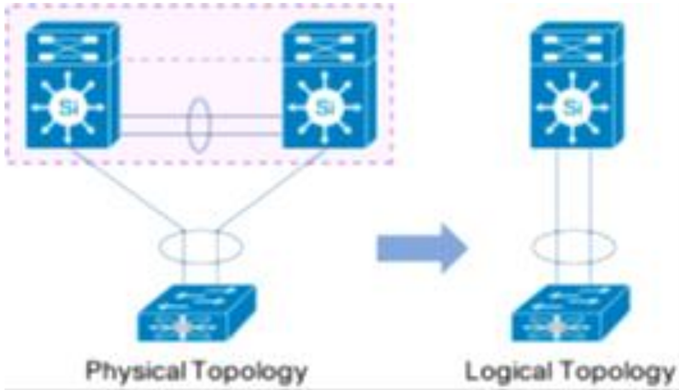

vPC Peers Configuration:

```
N7K 1:
feature vpc

! vpc domain config
vpc domain 95
  peer-switch
  role priority 200
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  track 10
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel6N7K-2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1-100,2001-
2010,3001-3010,3951-3960
  spanning-tree port type network
  vpc peer-link
! vpc peer-link member config
interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1-100,2001-
2010,3001-3010,3951-3960
  channel-group 6 mode active
  no shutdown

! vpc peer-keepalive config
interface Ethernet1/1
  vrf member vpc-keepalive
  ip address 1.1.1.1/24
  no shutdown

! vpc member port-channel config
interface port-channel7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  vpc 7
! vpc member port config
interface Ethernet8/1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
```

```
N7K 1:
feature vpc

! vpc domain config
vpc domain 95
  peer-switch
  role priority 200
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  track 10
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel5
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1-100,2001-
2010,3001-3010,3951-3960
  spanning-tree port type network
  vpc peer-link
! vpc peer-link member config
interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1-100,2001-
2010,3001-3010,3951-3960
  channel-group 5 mode active
  no shutdown

! vpc peer-keepalive config
interface Ethernet1/1
  vrf member vpc-keepalive
  ip address 1.1.1.2/24
  no shutdown

! vpc member port-channel config
interface port-channel7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  vpc 7
! vpc member port config
interface Ethernet8/1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
```

2010,3001-3010
  channel-group 7 mode active
  no shutdown
 	
!vpc object tracking
!! uplinks
track 1 interface port-channel3 line-protocol
track 2 interface port-channel4 line-protocol
!!vpc peer-link
track 3 interface port-channel6 line-protocol
track 10 list boolean or
  object 1
  object 2
  object 3
 	
! PIM prebuild SPT(only for non F2 mode)
ip pim pre-build-spt
	2010,3001-3010
	  channel-group 7 mode active
	  no shutdown
	
	!vpc object tracking
	!! uplinks
	track 1 interface port-channel3 line-protocol
	track 2 interface port-channel4 line-protocol
	!!vpc peer-link
	track 3 interface port-channel5 line-protocol
	track 10 list boolean or
	  object 1
	  object 2
	  object 3
	
	! PIM prebuild SPT(only for non F2 mode)
	ip pim pre-build-spt

Display vPC Status:

```
N7K-2# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 95
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 108
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po5    up     1-100,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------
id   Port      Status Consistency Reason                 Active vlans
--   ----      ------ ----------- ------                 ------------
7    Po7       up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
8    Po8       up     success     success                   1,11-20,200
                                                            1-2010,3001
```

3.1.4.1.1 LACP

NVT makes use of LACP mode active for all link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
N7K-2# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
```

```
       M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
3     Po3(RU)     Eth      LACP      Eth1/3(P)    Eth1/5(P)
4     Po4(RU)     Eth      LACP      Eth1/2(P)    Eth1/6(P)
5     Po5(SU)     Eth      LACP      Eth1/4(P)    Eth1/7(P)
7     Po7(SU)     Eth      LACP      Eth8/1(P)
8     Po8(SU)     Eth      LACP      Eth8/2(P)

DC6-DC101-6# show lacp interface ethernet 8/1
Interface Ethernet8/1 is up
  Channel group is 7 port channel is Po7
  PDUs sent: 2381
  PDUs rcvd: 2577
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(7f9b, 0-23-4-ee-be-5f, 8007, 0, 0), (8000, 0-1b-90-25-44-0, 6, 0, 0)] ]
Operational as aggregated link since Tue Aug 13 12:15:43 2013

Local Port: Eth8/1   MAC Address= 0-23-ac-64-bb-c2
  System Identifier=0x8000,  Port Identifier=0x8000,0x801
  Operational key=32775
  LACP_Activity=passive
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=60
Actor Oper State=60
Neighbor: 0x103
  MAC Address= 0-1b-90-25-44-0
  System Identifier=0x8000,  Port Identifier=0x8000,0x103
  Operational key=6
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=61
Partner Oper State=61
Aggregate or Individual(True=1)= 1
```

3.1.4.1.2 VLAN Trunking

NVT makes use of VLAN trunking in the aggregation-access blocks to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Display VLAN Information for Nexus 7000:

```
N7K-2# show vlan internal usage

VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035,4048-4059     Online Diagnostic
4036-4039,4060-4087     ERSPAN
4042                    Satellite
4040                    Fabric scale
3968-4095               Current
```

```
N7K-2# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po5, Po7, Po8, Po17, Po27, Po71
                                                Po72, Po73, Po74, Po77, Po78
                                                Po201, Po221, Po401, Po421
                                                Po441, Po501, Po521, Eth1/4
                                                Eth1/7, Eth8/1, Eth8/2, Eth8/16
                                                Eth8/18, Eth8/29, Eth8/30
                                                Eth9/42, Eth10/31, Eth102/1/1
                                                Eth102/1/21, Eth102/1/41
                                                Eth104/1/25, Eth104/1/26
                                                Eth104/1/27, Eth104/1/28
                                                Eth104/1/29, Eth104/1/30
                                                Eth104/1/31, Eth104/1/32

VLAN Type         Vlan-mode
---- -----        ----------
11   enet         CE

Remote SPAN VLAN
----------------
Disabled

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------
```

Display VLAN Information for Nexus 5000:

```
dc102-701# show vlan internal usage

VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035               Online Diagnostic
4036-4039               ERSPAN
4042                    Satellite
3968-4047,4094          Current
dc102-701#
```

Display VLAN Information for Nexus 3000:

```
dc102-47# show vlan internal usage

VLAN         DESCRIPTION
---------    -------------------------------------------------------
3968-4031    Multicast
4032         Online diagnostics vlan1
4033         Online diagnostics vlan2
4034         Online diagnostics vlan3
4035         Online diagnostics vlan4
4036-4047    Reserved
4094         Reserved
```

3.1.4.1.3 Spanning Tree

vPC technology helps build a loop free topology by leveraging port-channels from access devices to the vPC domain. A port-channel is seen as a logical link from the spanning tree's standpoint, so a vPC domain with vPC-attached access devices forms a star topology at Layer 2 (there are no STP blocked ports in this type of topology). In this case, STP is used as a fail-safe mechanism to protect against any network loops.

NVT makes use of Rapid-PVST which is the default spanning tree protocol on NX-OS. For networks with larger logical port counts, MST is recommended.

Display Spanning Tree Information:

```
N7K-2# show spanning-tree vlan 11

VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    24587
             Address     0023.04ee.be5f
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24587  (priority 24576 sys-id-ext 11)
             Address     0023.04ee.be5f
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po5              Desg FWD 1000      128.4100 (vPC peer-link) Network P2p
Po7              Desg FWD 200       128.4102 (vPC) P2p
Po8              Desg FWD 200       128.4103 (vPC) P2p
Po17             Desg FWD 200       128.4112 (vPC) P2p
Po71             Desg FWD 200       128.4166 (vPC) Edge P2p
Po77             Desg FWD 200       128.4172 (vPC) Edge P2p
Po78             Desg FWD 200       128.4173 (vPC) Edge P2p
Eth102/1/1       Desg FWD 20000     128.4197 Edge P2p
Eth102/1/21      Desg FWD 20000     128.4197 Edge P2p
Eth102/1/41      Desg FWD 20000     128.4197 Edge P2p

N7K-2# show spanning-tree summary totals
Switch is in rapid-pvst mode
Root bridge for: VLAN0001-VLAN0006, VLAN0009-VLAN0100, VLAN2001-VLAN2010
  VLAN3001-VLAN3010, VLAN3951-VLAN3960
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is long
vPC peer switch                          is enabled (operational)
STP-Lite                                 is enabled

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
    130 ans                   0         0        0        488        488
```

3.1.4.1.4 vPC Peer Switch Feature

The vPC Peer Switch feature allows a pair of vPC peer devices to appear as a single Spanning Tree Protocol root in the Layer 2 topology (they have the same bridge ID). vPC peer switch must be configured on both vPC peer devices to become operational.

This feature simplifies Spanning Tree Protocol configuration by configuring vPC VLANs on both peer devices with the same Spanning Tree Protocol priority. A vPC Peer Switch eliminates the need to map the Spanning Tree Protocol root to the vPC primary peer device.

Figure 13 vPC Peer Switch

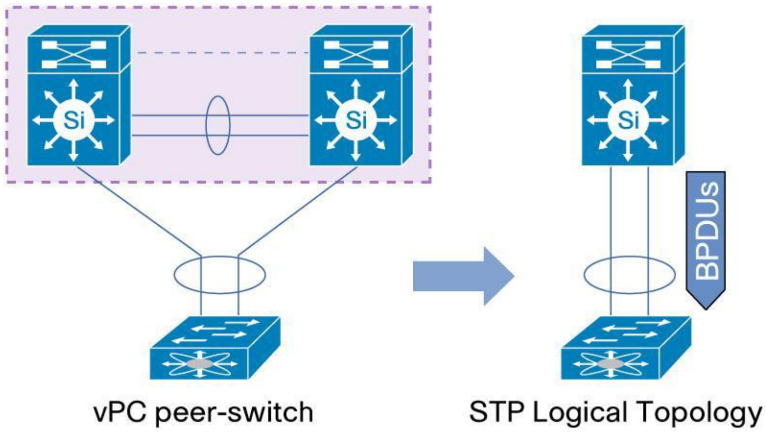

3.1.4.1.5 Configuration Parameters Consistency

After the vPC feature is enabled and the vPC peer-link on both peer devices is configured, Cisco Fabric Services messages provide a copy of the local vPC peer device configuration to the remote vPC peer device. The systems then determine whether any of the crucial configuration parameters differ on the two devices.

When a Type 1 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 1 consistency check failure, all vPC member ports are set to down state.
- For a vPC interface configuration Type 1 consistency check failure, the misconfigured vPC is set to down state.

When a Type 2 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 2 consistency check failure, all vPC member ports remain in up state and vPC systems trigger protective actions.
- For a vPC interface configuration Type 2 consistency check failure, the misconfigured vPC remains in up state. However, depending on the discrepancy type, vPC systems will trigger protective actions. The most typical misconfiguration deals with the allowed VLANs in the vPC interface trunking configuration. In this case, vPC systems will disable the vPC interface VLANs that do not match on both sides.

Display vPC Consistency Parameters:

```
N7K-2# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
```

```
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled              Disabled
STP Bridge Assurance        1     Enabled               Enabled
STP Port Type, Edge         1     Normal, Disabled,     Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled              Disabled
STP MST Simulate PVST       1     Enabled               Enabled
Interface-vlan admin up     2     1,10-20,2001-2010,3951 1,10-20,2001-2010,3951
                                  -3960                 -3960
Interface-vlan routing      2     1,10-20,2001-2010,3951 1,10-20,30,75,2001-201
capability                        -3960                 0,3951-3960
VTP domain                  2     interop               interop
VTP version                 2     1                     1
VTP mode                    2     Server                Server
VTP password                2
VTP pruning status          2     Disabled              Disabled
Allowed VLANs               -     1-100,2001-2010,3001-3 1-100,2001-2010,3001-3
                                  010,3951-3960         010,3951-3960
Local suspended VLANs       -     -                     -

N7K-2# show vpc consistency-parameters interface port-channel 7

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value           Peer Value
-------------               ----  --------------------- ----------------------
lag-id                      1     [(7f9b,               [(7f9b,
                                  0-23-4-ee-be-5f, 8007, 0-23-4-ee-be-5f, 8007,
                                   0, 0), (8000,         0, 0), (8000,
                                  0-1b-90-25-44-0, 6, 0, 0-1b-90-25-44-0, 6, 0,
                                   0)]                   0)]
mode                        1     passive               passive
STP Port Type               1     Default               Default
STP Port Guard              1     Default               Default
STP MST Simulate PVST       1     Default               Default
Native Vlan                 1     1                     1
Port Mode                   1     trunk                 trunk
MTU                         1     1500                  1500
Duplex                      1     full                  full
Speed                       1     10 Gb/s               10 Gb/s
Admin port mode             1     trunk                 trunk
Interface type              1     port-channel          port-channel
LACP Mode                   1     on                    on
vPC card type               1     Earl8                 Earl8
Allowed VLANs               -     1,10-20,2001-2010,3001 1,11-20,2001-2010,3001
                                  -3010                 -3010
Local suspended VLANs       -     10                    -
```

3.1.4.1.6 vPC in mixed chassis mode (M1/F1 ports in same system or VDC)

Mixed chassis mode is a system where both M1 ports and F1 ports are used simultaneously.

M1 Series line cards provide scalable Layer 2 and Layer 3 capabilities. F1 Series line cards provide high-density cost-effective Layer 2 10-Gigabit Ethernet connectivity. Interoperability between M1 and F1 ports is provided by L3 internal proxy routing where M1 ports are used for L3 proxy when traffic entering a F1 port needs to be routed (L3 traffic for inter VLAN routing or traffic going outside of data center). M1 line cards typically host the interface VLAN (i.e SVI - Switch Virtual Interface) on behalf of F1 line cards.

A vPC system in mixed chassis mode with peer-link on F1 ports presents the following characteristics:

- The total number of MAC addresses supported is 16K (capacity of one forwarding engine [i.e switch on chip] on the F1 series line card)
- M1 ports are used only for L3 uplinks
- F1 ports are used for vPC member ports (can use M1 ports as well if needed)
- Must use the *peer-gateway exclude-vlan <VLAN list>* knob to exclude VLANs that belong to backup routing path. This will avoid the transit traffic between vPC peer devices using the vPC peer-link from being punted to CPU, allowing direct HW switching (this command only applies to a vPC system in mixed chassis mode with vPC peer-link on F1)

NVT makes use of M1 ports for the vPC peer-link and F1 ports for the vPC member ports. A vPC system in mixed chassis mode with the peer-link on M1 ports presents the following characteristics:

- The total number of MAC addresses supported is 128K (capacity of forwarding engine on the M1 series line card)
- M1 ports are used for L3 uplinks and vPC peer-link
- F1 ports are used for vPC member ports (can use M1 ports as well if needed). Non-overlapping assignment of vlans on the F1 card SoC's ensures the best use of mac address table space.
- There is no need to use the *peer-gateway exclude-vlan <VLAN list>* knob

Display Layer 3 proxy details:

```
N7K-2# show hardware proxy layer-3 detail

Global Information:
       Layer-2 only Modules: Count: 1, Slot: 7
       Layer-3 Modules supporting proxy layer-3: Count: 4, Slot: 1,8-10

       Replication Rebalance Mode:            Manual
       Number of proxy layer-3 forwarders:    22
       Number of proxy layer-3 replicators:   14

Forwarder Interfaces                     Status     Reason
-------------------------------------------------------------------------------
Eth1/1                                   up         SUCCESS
Eth1/2                                   up         SUCCESS
Eth1/3                                   up         SUCCESS
Eth1/4                                   up         SUCCESS
Eth1/5                                   up         SUCCESS
Eth1/6                                   up         SUCCESS
Eth1/7                                   up         SUCCESS
Eth8/1, Eth8/3, Eth8/5, Eth8/7           up         SUCCESS
Eth8/2, Eth8/4, Eth8/6, Eth8/8           up         SUCCESS
Eth8/10, Eth8/12, Eth8/14, Eth8/16       up         SUCCESS
Eth8/17, Eth8/19, Eth8/21, Eth8/23       up         SUCCESS
Eth8/18, Eth8/20, Eth8/22, Eth8/24       up         SUCCESS
Eth8/25, Eth8/27, Eth8/29, Eth8/31       up         SUCCESS
Eth8/26, Eth8/28, Eth8/30, Eth8/32       up         SUCCESS
Eth9/41-44                               up         SUCCESS
Eth10/1, Eth10/3, Eth10/5, Eth10/7       up         SUCCESS
Eth10/2, Eth10/4, Eth10/6, Eth10/8       up         SUCCESS
Eth10/9, Eth10/11, Eth10/13, Eth10/15    up         SUCCESS
Eth10/10, Eth10/12, Eth10/14, Eth10/16   up         SUCCESS
Eth10/17, Eth10/19, Eth10/21, Eth10/23   up         SUCCESS
Eth10/18, Eth10/20, Eth10/22, Eth10/24   up         SUCCESS
Eth10/25, Eth10/27, Eth10/29, Eth10/31   up         SUCCESS

RE = Replication Engine
Replicator Interfaces (RE instance)      #Interface-Vlan   Interface-Vlan
-------------------------------------------------------------------------------
Eth1/1-2 (0)                             4                 1,10-12
Eth1/3-4 (1)                             3                 13-15
Eth1/5-6 (2)                             3                 16-18
```

```
Eth1/7-8 (3)                              3              19-20,2001
Eth8/1, Eth8/3, Eth8/5, Eth8/7, Eth8/9,   3              2002-2004
Eth8/11, Eth8/13, Eth8/15 (3)
Eth8/2, Eth8/4, Eth8/6, Eth8/8, Eth8/10,  3              2005-2007
 Eth8/12, Eth8/14, Eth8/16 (0)
Eth8/17, Eth8/19, Eth8/21, Eth8/23,       3              2008-2010
Eth8/25, Eth8/27, Eth8/29, Eth8/31 (2)
Eth8/18, Eth8/20, Eth8/22, Eth8/24,       3              2950-2952
Eth8/26, Eth8/28, Eth8/30, Eth8/32 (1)
Eth9/1-24 (0)                             3              2953-2955
Eth9/25-48 (1)                            3              2956-2958
Eth10/1, Eth10/3, Eth10/5, Eth10/7,       3              2959-2960,3951
Eth10/9, Eth10/11, Eth10/13, Eth10/15
(3)
Eth10/2, Eth10/4, Eth10/6, Eth10/8,       3              3952-3954
Eth10/10, Eth10/12, Eth10/14, Eth10/16
(0)
Eth10/17, Eth10/19, Eth10/21, Eth10/23,   3              3955-3957
Eth10/25, Eth10/27, Eth10/29, Eth10/31
(2)
Eth10/18, Eth10/20, Eth10/22, Eth10/24,   3              3958-3960
Eth10/26, Eth10/28, Eth10/30, Eth10/32
(1)
```

3.1.4.1.7 vPC Role Priority

There are two defined vPC roles: primary and secondary. The vPC role defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDUs) and responds to Address Resolution Protocol (ARP).

In case of a tie (same role priority value defined on both peer devices), the lowest system MAC will dictate the primary peer device.

Display vPC Role, System-MAC, System-Priority:

```
N7K-2# show vpc role

vPC Role status
----------------------------------------------------
vPC role                        : primary
Dual Active Detection Status    : 0
vPC system-mac                  : 00:23:04:ee:be:5f
vPC system-priority             : 32667
vPC local system-mac            : 00:23:ac:64:bb:c2
vPC local role-priority         : 110
```

3.1.4.1.8 vPC Peer-Link

The vPC peer-link is a standard 802.1Q trunk that performs the following actions:

- Carry vPC and non-vPC VLANs.
- Carry Cisco Fabric Services (CFS) messages that are tagged with CoS=4 for reliable communication CoS=4 for reliable communication.
- Carry flooded traffic between the vPC peer devices.
- Carry STP BPDUs, HSRP hello messages, and IGMP updates.

When the vPC peer-link fails and the vPC peer-keepalive link is still up, the vPC secondary peer device performs the following operations:

- Suspends its vPC member ports

- Shuts down the SVI associated to the vPC VLAN

Display vPC Peer-link Information:

```
DC101-5# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 95
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 108
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po6    up     1-100,300,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- ------                     ------------
7    Po7    up     success     success                    1,11-20,300
                                                          ,2001-2010,
8    Po8    up     success     success                    1,11-20,300
                                                          ,2001-2010,
                                                          3001-3010
17   Po17   up     success     success                    1,11-20,200
                                                          1-2010,3001
                                                          -3010
27   Po27   up     success     success                    1,10-20,200
                                                          1-2010,3001
                                                          -3010
```

3.1.4.1.9 vPC Peer-Keepalive Link

The vPC peer-keepalive link is a Layer 3 link that joins one vPC peer device to the other vPC peer device and carries a periodic heartbeat between those devices. It is used at the boot up of the vPC systems to guarantee that both peer devices are up before forming the vPC domain. It is also used when the vPC peer-link fails, in which case, the vPC peer-keepalive link is leveraged to detect split brain scenario (both vPC peer devices are active-active).

Default Values for VPC Peer-Keepalive Links:

Timer	Default value
Keepalive interval	1 seconds
Keepalive hold timeout (on vPC peer-link loss)	3 seconds
Keepalive timeout	5 seconds

When building a vPC peer-keepalive link, use the following in descending order of preference:

1. Dedicated link(s) (1-Gigabit Ethernet port is enough) configured as L3. A port-channel with 2 X 1G port is preferred.

2. Mgmt0 interface (along with management traffic).
3. As a last resort, route the peer-keepalive link over the Layer 3 infrastructure.

NVT makes use of the 1st option.

Display vPC Peer-Keepalive Information:

```
DC101-5# sh vpc peer-keepalive

vPC keep-alive status             : peer is alive
--Peer is alive for             : (8755) seconds, (95) msec
--Send status                   : Success
--Last send at                  : 2014.02.18 00:56:35 559 ms
--Sent on interface             : Eth1/1
--Receive status                : Success
--Last receive at               : 2014.02.18 00:56:35 651 ms
--Received on interface         : Eth1/1
--Last update from peer         : (0) seconds, (504) msec

vPC Keep-alive parameters
--Destination                   : 1.1.1.2
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc-keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

3.1.4.1.10 vPC Member Link

As suggested by the name, a vPC member port is a port-channel member of a vPC. A port-channel defined as a vPC member port always contains the keywords *vpc <vpc id>*.

A vPC only supports Layer 2 port-channels. The port-channel can be configured in access or trunk switchport mode. Any VLAN allowed on the vPC member port is by definition called a vPC VLAN. Whenever a vPC VLAN is defined on a vPC member port, it must also be defined on the vPC peer-link. Not defining a vPC VLAN on the vPC peer-link will cause the VLAN to be suspended.

The configuration of the vPC member port must match on both the vPC peer devices. If there is an inconsistency, a VLAN or the entire port channel may be suspended (depending on Type-1 or Type-2 consistency check for the vPC member port). For instance, a MTU mismatch will suspend the vPC member port.

Display vPC Member Port-channel Information:

```
N7K-2# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 95
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 108
Track object                      : 10
Peer Gateway                      : Disabled
```

```
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po5    up     1-100,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------
id   Port      Status Consistency Reason                 Active vlans
--   ----      ------ ----------- ------                 -----------
7    Po7       up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
8    Po8       up     success     success                   1,11-20,200
                                                            1-2010,3001
N7K-2# show vpc consistency-parameters interface port-channel 7

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-5f, 8007, 0-23-4-ee-be-5f, 8007,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-1b-90-25-44-0, 6, 0, 0-1b-90-25-44-0, 6, 0,
                                   0)]                    0)]
mode                        1     passive                passive
STP Port Type               1     Default                Default
STP Port Guard              1     Default                Default
STP MST Simulate PVST       1     Default                Default
Native Vlan                 1     1                      1
Port Mode                   1     trunk                  trunk
MTU                         1     1500                   1500
Duplex                      1     full                   full
Speed                       1     10 Gb/s                10 Gb/s
Admin port mode             1     trunk                  trunk
Interface type              1     port-channel           port-channel
LACP Mode                   1     on                     on
vPC card type               1     Earl8                  Earl8
Allowed VLANs               -     1,10-20,2001-2010,3001 1,11-20,2001-2010,3001
                                  -3010                  -3010
Local suspended VLANs       -     10                     -
```

3.1.4.1.11 vPC ARP Synchronization

The vPC ARP Synchronization feature improves the convergence time for Layer 3 flows (North to South traffic). When the vPC peer-link fails and subsequently recovers, vPC ARP Synchronization performs an ARP bulk synchronization over Cisco Fabric Services (CFS) from the vPC primary peer device to the vPC secondary peer device.

Displays vPC ARP Synchronization Information:

```
N7K-2# show ip arp sync-entries

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Address          Age       MAC Address     Interface
101.39.59.101    00:01:22  0050.5601.0009  Vlan3959
101.39.59.102    00:01:22  0050.5601.0109  Vlan3959
101.39.59.103    00:01:22  0050.5601.0209  Vlan3959
```

```
101.39.59.104    00:01:22  0050.5601.0309  Vlan3959
```

3.1.4.1.12 vPC Delay Restore

After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from the access to the core until the Layer 3 connectivity is reestablished.

The vPC Delay Restore feature delays the vPCs leg bringup on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on the vPC leg. The result provides a graceful restoration along with zero packet loss during the recovery phase (traffic still gets diverted to the alive vPC peer device).

This feature is enabled by default with a vPC restoration default timer of 30 seconds, which NVT maintains in the testbed.

3.1.4.1.13 vPC Object-Tracking

A vPC deployment with a single Cisco Nexus 7000 Series M132XP-12 module or M108XP-12 module, where the L3 core uplinks and vPC peer-link interfaces are localized on the same module, is vulnerable to access layer isolation if the 10-Gbps module fails on the primary vPC (vPC member ports are defined on both 1-Gbps line cards and on 10-Gbps line card).

In this scenario, the vPC Object Tracking feature shuts down vPC member ports on the peer device where M1 10-Gbps is damaged (irrespective of vPC role primary or secondary). This triggered action allows traffic flows (southbound and northbound) to go through the other peer device where the M1 10-Gbps line card is up.

Figure 14 vPC Object Tracking Feature – Behavior when vPC Peer-link Fails

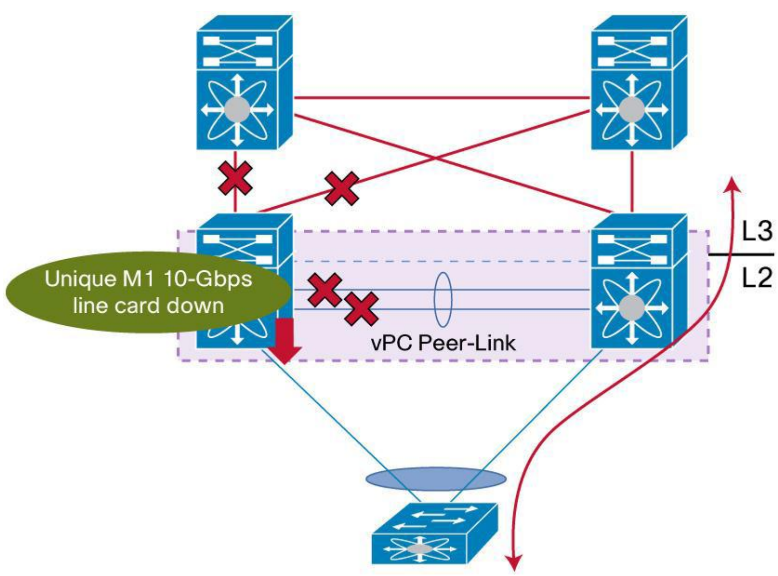

The vPC Object Tracking feature suspends the vPCs on the impaired device so that traffic can be diverted over the remaining vPC peer.

To use vPC object tracking, track both Peer-link interfaces and L3 core interfaces as a list of Boolean objects. Note that the Boolean AND operation is not supported with vPC object tracking. The vPC object tracking configuration must be applied on both vPC peer devices.

Sample Configuration:

```
! Track the vpc peer link
track 1 interface port-channel5 line-protocol
! Track the uplinks to the core
track 2 interface port-channel3 line-protocol
track 3 interface port-channel4 line-protocol
! Combine all tracked objects into one.
! "OR" means if ALL objects are down, this object will go down
! ==> lost all connectivity to the L3 core and the peer link
track 10 list boolean OR
  object 1
  object 2
  object 3
! If object 10 goes down on the primary vPC peer,
! system will switch over to other vPC peer and disable all local vPCs
vpc domain 95
  track 10
```

Display Tracked Object Status:

```
N7K-2# show track 10
Track 10
  List  Boolean or
  Boolean or is UP
  4 changes, last change 1d00h
  Track List Members:
    object 3 UP
    object 2 UP
    object 1 UP
  Tracked by:
    vPCM Domain 95
```

3.1.4.1.14 vPC Auto-Recovery

vPC auto-recovery feature was designed to address 2 enhancements to vPC.

- To provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure (vPC auto-recovery feature).
- To handle a specific case where both vPC peer devices reload but only one comes back to life (vPC auto-recovery reload-delay feature).

The switch which unsuspends its vPC role with vPC auto-recovery continues to remain primary even after peer-link is on. The other peer takes the role of secondary and suspends its own vPC until a consistency check is complete. Therefore, to avoid this situation from occurring erroneously, auto-recovery reload-delay-timer should be configured to be long enough for the system to fully complete its bootup sequence.

Helpful Commands for vPC Object Tracking:

Show vpc brief	Displays Auto-recovery status

Configuration Check:

```
N7K-2# show vpc brief
```

```
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 95
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 108
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
    1      Po5    up     1-100,2001-2010,3001-3010,3951-3960
```

3.1.4.1.15 HSRP Active/Active with vPC

HSRP in the context of vPC has been improved from a functional and implementation standpoint to take full benefits of the L2 dual-active peer devices nature offered by vPC technology. HSRP operates in active-active mode from a data plane standpoint, as opposed to classical active/standby implementation with a STP based network. No additional configuration is required. As soon as a vPC domain is configured and interface VLAN with an associated HSRP group is activated, HSRP will behave by default in active/active mode (on the data plane side).
From a control plane standpoint, active-standby mode still applies for HSRP in context of vPC; the active HSRP instance responds to ARP request. ARP response will contain the HSRP vMAC which is the same on both vPC peer devices. The standby HSRP vPC peer device just relays the ARP request to active HSRP/VRRP peer device through the vPC peer-link.

Sample Configuration:

```
! N7K-1:
interface Vlan11
  no ip redirects
  ip address 101.11.0.21/16
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 101.11.0.1
  no shutdown

! N7K-2:
interface Vlan11
  no ip redirects
  ip address 101.11.0.19/16
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 101.11.0.1
  no shutdown
```

Helpful Commands for HSRP Active/Active with vPC:

Show hsrp brief	Displays hsrp status
Show mac address-table vlan <vlan id>	Displays mac addresses including HSRP vMAC; check for G-flag on vMAC for active/active HSRP

Configuration Check:

```
N7K-2# show hsrp brief
*:IPv6 group   #:group belongs to a bundle
                     P indicates configured to preempt.
                     |
 Interface   Grp  Prio P State    Active addr      Standby addr     Group addr
  Vlan1       1    200  P Active   local            101.0.1.19       101.0.1.1       (conf)
  Vlan10      1    200  P Active   local            101.10.0.19      101.10.0.1      (conf)
  Vlan10      2    200  P Active   local            101.10.0.19      101.110.0.1     (conf)
  Vlan11      1    200  P Active   local            101.11.0.19      101.11.0.1      (conf)
N7K-2# show mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G 11       0000.0c9f.f001    static       -       F    F  sup-eth1(R)
G 11       0000.0c9f.f002    static       -       F    F  sup-eth1(R)
```

3.1.4.1.16 PIM Pre-build-spt with vPC

PIM Pre-build SPT on non-forwarder attracts multicast traffic by triggering upstream PIM J/Ps (Join/Prune) without setting any interface in the OIF (Outgoing Interface) list. Multicast traffic is then always pulled to the non-active forwarder and finally dropped due to no OIFs.

The immediate effect of enabling PIM Pre-build SPT is to improve the convergence time upon active forwarder failure (1 to 3 seconds of convergence time). The other vPC peer device (which is the non-active forwarder) does not need to create any new upstream multicast state and can quickly transition to the active forwarder role by properly programming the OIF (Outgoing Interface) list.
The impact of enabling PIM prebuild SPT is the consumption of bandwidth and replication capacity on the primary and secondary data path (i.e. on vPC primary and secondary peer devices) in steady state .

In the vPC implementation in F2-mode, because of a hardware limitation, the PIM dual DR mode is disabled. In this case (with F2 mode), even if the **ip pim pre-build-spt** command is configured, there is no value added because the corresponding (S,G) route is not created in the first place.

As shown below, on the non-forwarder/secondary, the (S,G) is created with no OIFs.

On Non-Forwarder:

```
DC5-DC101-5# sh ip mrout 230.102.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.102.0.1/32), uptime: 00:24:01, igmp ip pim
  Incoming interface: port-channel3, RPF nbr: 40.101.1.15
  Outgoing interface list: (count: 20)
    Vlan20, uptime: 00:22:39, igmp
    Vlan14, uptime: 00:22:43, igmp
```

```
    Vlan16, uptime: 00:22:45, igmp
    Vlan12, uptime: 00:22:56, igmp
    Vlan18, uptime: 00:22:58, igmp
    Vlan15, uptime: 00:22:58, igmp
    Vlan17, uptime: 00:22:58, igmp
    Vlan11, uptime: 00:22:58, igmp
    Vlan13, uptime: 00:22:58, igmp
    Vlan19, uptime: 00:23:00, igmp
    Vlan2010, uptime: 00:23:54, igmp
    Vlan2006, uptime: 00:23:54, igmp
    Vlan2001, uptime: 00:23:54, igmp
    Vlan2008, uptime: 00:23:57, igmp
    Vlan2009, uptime: 00:23:57, igmp
    Vlan2003, uptime: 00:23:57, igmp
    Vlan2007, uptime: 00:23:57, igmp
    Vlan2002, uptime: 00:23:58, igmp
    Vlan2004, uptime: 00:24:00, igmp
    Vlan2005, uptime: 00:24:01, igmp

(102.11.17.1/32, 230.102.0.1/32), uptime: 00:24:01, ip pim
  Incoming interface: port-channel4, RPF nbr: 40.101.3.17
  Outgoing interface list: (count: 0)


DC5-DC101-5# sh ip pim intern vpc rp
PIM vPC RPF-Source Cache for Context "default" - Chassis Role Secondary


Source: 102.11.17.1
  Pref/Metric: 110/1100
  Source role: secondary
  Forwarding state: Tie (not forwarding)
```

3.1.4.1.17 Double-Sided vPC Topology

A double-sided vPC topology superposes two layers of vPC domain and the bundle between vPC domain 1 and vPC domain 2 is by itself a vPC. The vPC domain at the bottom is used for active/active connectivity from end-point devices to the network access layer. The vPC domain at the top is used for active/active FHRP in the L2/L3 boundary aggregation layer.

Figure 15 Double-Sided vPC Topology

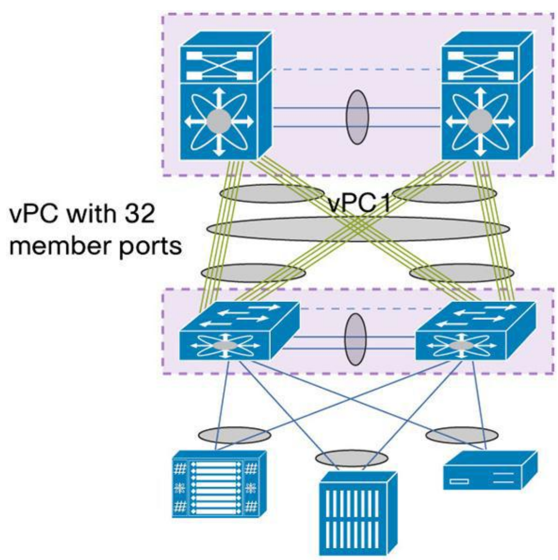

Benefits of double-sided vPC over single-sided vPC topology are listed below:

- Enables a larger Layer 2 domain.
- Provides a highly resilient architecture. In double-sided vPC, two access switches are connected to two aggregation switches whereas in single-sided vPC, one access switch is connected to two aggregation switches.
- Provides more bandwidth from the access to aggregation layer. Using a Cisco Nexus F1 Series modules line card for vPC and Cisco Nexus 5000 Series Switches with Release 4.1(3)N1(1a) or later, a vPC with 32 active member ports (that is, 320 Gbps) can be instantiated.

3.1.4.2 FabricPath

NVT FabricPath topology is designed to have four spines using Nexus 7000 at the aggregation layer. There are six Nexus 5000 leaf switches on access layer that are connected to all four spines. The FabricPath feature is only supported on the F-Series modules on the Nexus 7000. In DC1, spine switches consist of Nexus 7000 with Sup 1 and F1 linecards.

Because of the multiple forwarding engines (FEs) on the F-Series modules, the port pairs and port sets in the table below must be configured to be in the same VDC.

Nexus 7000 F Series Modules Port Pairs and Port Sets	
Port Pairs for F1 Modules	**Port Sets for F2 Modules**
Ports 1 and 2	Ports 1, 2, 3, 4
Ports 3 and 4	Ports 5, 6, 7, 8
Ports 5 and 6	Ports 9, 10, 11, 12
Ports 7 and 8	Ports 13, 14, 15, 16
Ports 9 and 10	Ports 17, 18, 19, 20
Ports 11 and 12	Ports 21, 22, 23, 24
Ports 13 and 14	Ports 25, 26, 27, 28
Ports 15 and 16	Ports 29, 30, 31, 32
Ports 17 and 18	Ports 33, 34, 35, 36
Ports 19 and 20	Ports 37, 38, 39, 40
Ports 21 and 22	Ports 41, 42, 43, 44
Ports 23 and 24	Ports 45, 46, 47, 48
Ports 25 and 26	
Ports 27 and 28	
Ports 29 and 30	
Ports 31 and 32	

NVT FabricPath Configuration:

```
dc102-703# show running-config FabricPath
```

```
!Command: show running-config FabricPath
!Time: Fri Mar  7 12:20:27 2014

version 5.2(1)N1(3)
feature-set FabricPath

logging level FabricPath isis 5

vlan 1,11-20,2001-2010,3001-3010
  mode FabricPath
FabricPath switch-id 251
logging level FabricPath switch-id 5
vpc domain 211
  FabricPath switch-id 1001
  FabricPath multicast load-balance


interface port-channel52
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel701
  switchport mode FabricPath

interface port-channel702
  switchport mode FabricPath

interface port-channel703
  switchport mode FabricPath

interface port-channel704
  switchport mode FabricPath

interface port-channel705
  switchport mode FabricPath

interface port-channel706
  switchport mode FabricPath

FabricPath domain default
  root-priority 109
FabricPath load-balance unicast include-vlan
FabricPath load-balance multicast rotate-amount 0x3 include-vlan
```

3.1.4.2.1 FabricPath Switch-IDs

Cisco FabricPath can assign switch IDs to all the devices in the network automatically; however, it is convenient to use a meaningful numbering scheme. During network troubleshooting, having a distinct numbering scheme allows for faster and easier switch role identification.

NVT has assigned switch IDs using the following scheme in the FabricPath domain network:

- The devices in the spine layer have been assigned an ID related to spine VDC naming: 251 to 254
- The devices in the leaf layer have been assigned an ID related to leaf device naming: 701 to 706
- The virtual switch for the domain has been assigned an ID: 1001-1002 and 2001-2003

Figure 16 NVT FabricPath POD Logical Topology

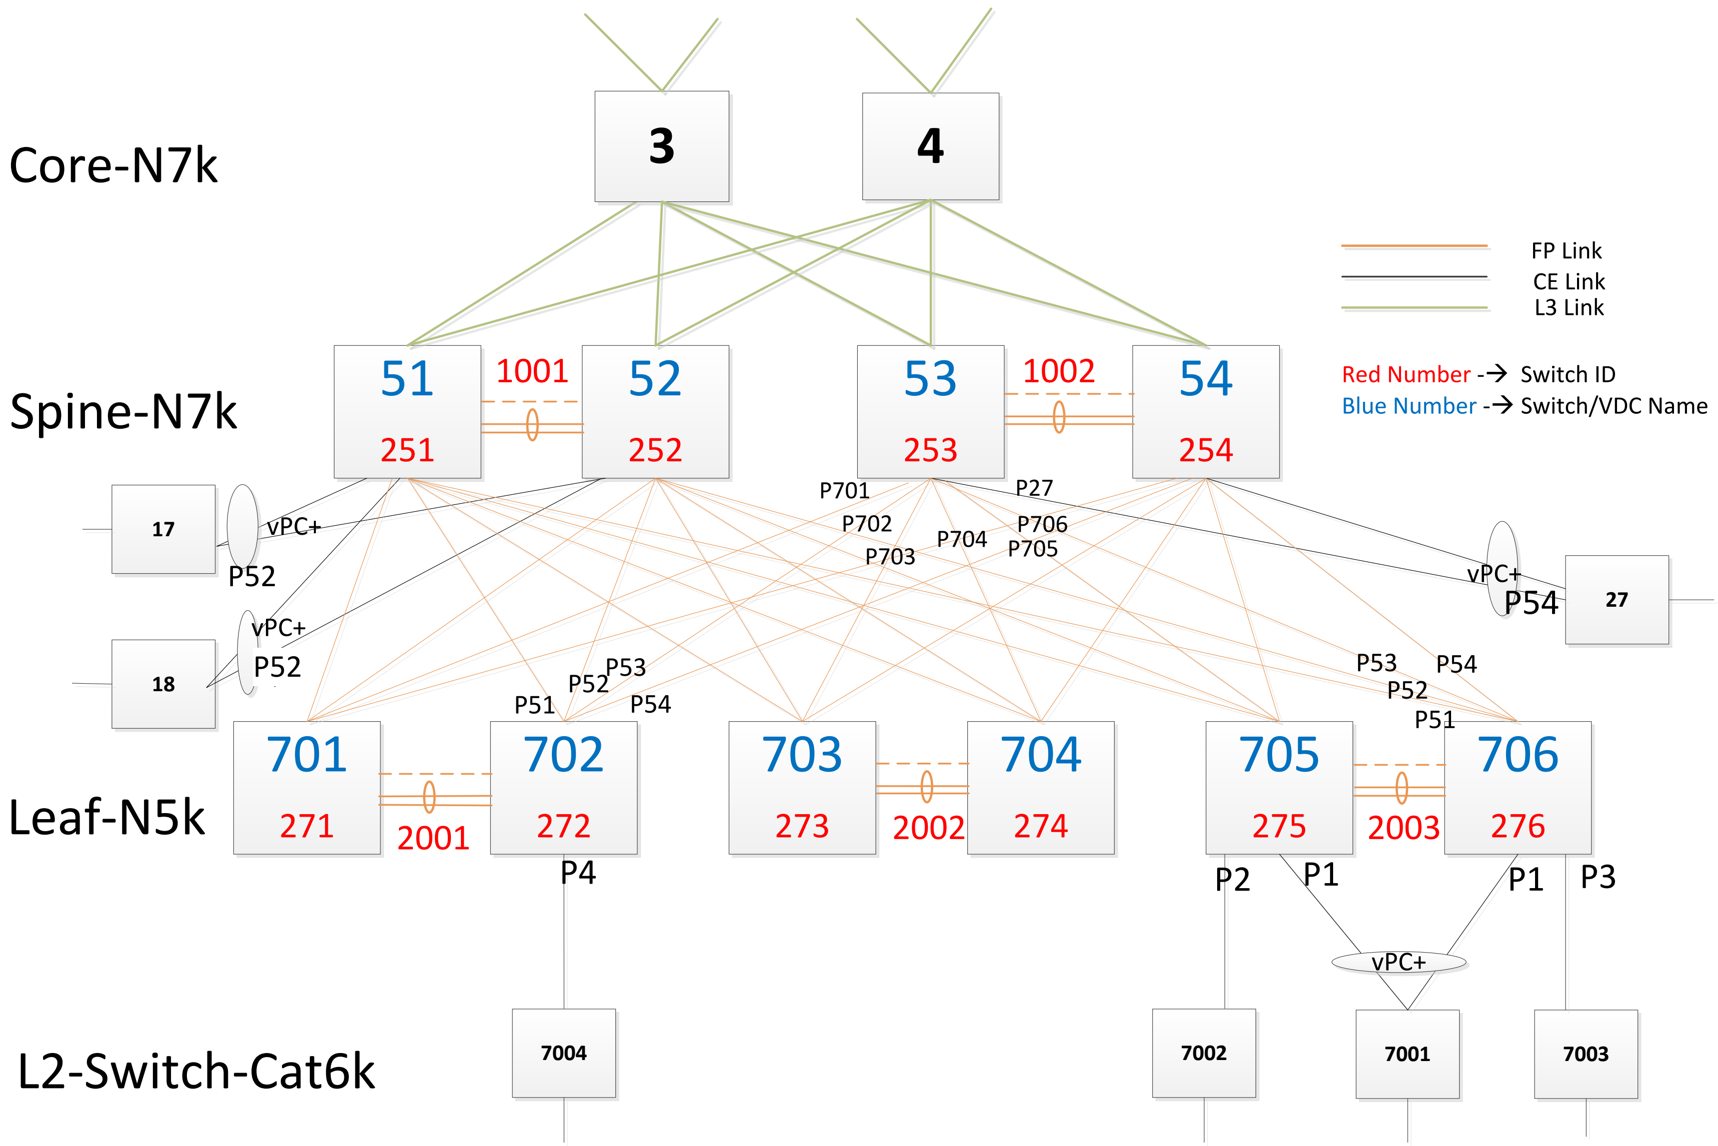

To Verify the FabricPath Switch ID:

```
DC5-DC102-51# show FabricPath switch-id local
Switch-Id: 151
System-Id: 0023.ac64.b2c3

DC5-DC102-51# show FabricPath switch-id
                        FABRICPATH SWITCH-ID TABLE
Legend: '*' - this system
=========================================================================
SWITCH-ID      SYSTEM-ID       FLAGS         STATE    STATIC  EMULATED
----------+----------------+------------+-----------+--------------------
 100         0023.ac64.b2c3    Primary     Confirmed    No      Yes
 100         0023.ac64.bbc3    Primary     Confirmed    No      Yes
*151         0023.ac64.b2c3    Primary     Confirmed    Yes     No
 152         0023.ac64.bbc3    Primary     Confirmed    Yes     No
 153         0023.ac64.b2c4    Primary     Confirmed    Yes     No
 154         0023.ac64.bbc4    Primary     Confirmed    Yes     No
 200         0023.ac64.b2c4    Primary     Confirmed    No      Yes
 200         0023.ac64.bbc4    Primary     Confirmed    No      Yes
 701         547f.eed1.d681    Primary     Confirmed    No      Yes
 701         547f.eed2.7741    Primary     Confirmed    No      Yes
 703         547f.eed1.d57c    Primary     Confirmed    No      Yes
 703         547f.eed2.7981    Primary     Confirmed    No      Yes
 705         547f.eed2.723c    Primary     Confirmed    No      Yes
 705         547f.eed2.757c    Primary     Confirmed    No      Yes
 1692        547f.eed2.723c    Primary     Confirmed    No      No
 1701        547f.eed1.d681    Primary     Confirmed    Yes     No
 1702        547f.eed2.7741    Primary     Confirmed    Yes     No
 1703        547f.eed2.7981    Primary     Confirmed    Yes     No
 1704        547f.eed1.d57c    Primary     Confirmed    Yes     No
 3402        547f.eed2.757c    Primary     Confirmed    No      No
Total Switch-ids: 20
```

3.1.4.2.2 FabricPath VLANs

Cisco FabricPath VLANs should be consistently defined on all the Cisco FabricPath switches in a particular FabricPath topology.

To Verify the FabricPath VLANs:

```
DC6-DC102-54# show FabricPath isis vlan-range
FabricPath IS-IS domain: default
MT-0
Vlans configured:
1, 10-20, 2001-2010, 3001-3010, 3951-3960
```

3.1.4.2.3 FabricPath Core Port

The configuration of a FabricPath core port is performed with the command *switchport mode FabricPath*. The FabricPath core port exchanges topology info through L2 ISIS adjacency and forwarding based on the Switch ID Table.

To Verify the FabricPath Interface:

```
DC6-DC102-54# show FabricPath isis interface port-channel 701
FabricPath IS-IS domain: default
Interface: port-channel701
```

```
  Status: protocol-up/link-up/admin-up
  Index: 0x0001, Local Circuit ID: 0x01, Circuit Type: L1
  No authentication type/keychain configured
  Authentication check specified
  Extended Local Circuit ID: 0x160002BC, P2P Circuit ID: 0000.0000.0000.00
  Retx interval: 5, Retx throttle interval: 66 ms
  LSP interval: 33 ms, MTU: 1500
  P2P Adjs: 1, AdjsUp: 1, Priority 64
  Hello Interval: 10, Multi: 3, Next IIH: 00:00:01
  Level   Adjs   AdjsUp  Metric   CSNP  Next CSNP  Last LSP ID
  1        1       1      20     60  00:00:09   ffff.ffff.ffff.ff-ff
  Topologies enabled:
    Topology Metric  MetricConfig Forwarding
    0        20      no           UP

DC6-DC102-54# show FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel53 P2P   7     Up/Ready    0x01/L1   1500 200     64         1/1
port-channel701 P2P   1     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel702 P2P   2     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel703 P2P   3     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel704 P2P   4     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel705 P2P   5     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel706 P2P   6     Up/Ready   0x01/L1   1500 20      64          1/1
```

3.1.4.2.4 FabricPath Metric

Cisco FabricPath ISIS calculates the preferred path to any switch-ID based on the metric to any given destination. The metric is as follows:

- 1-Gbps Ethernet links have a cost of 400
- 10-Gigabit Ethernet links have a cost of 40
- 20-Gbps have a cost of 20

NVT has set a higher ISIS metric on vPC peer links between the spine switches to prevent traffic from flowing through the vPC peer links.

To Verify the FabricPath ISIS Metric, use the Following Commands:

```
DC6-DC102-54# show FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel53 P2P   7     Up/Ready    0x01/L1   1500 200     64         1/1
port-channel701 P2P   1     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel702 P2P   2     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel703 P2P   3     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel704 P2P   4     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel705 P2P   5     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel706 P2P   6     Up/Ready   0x01/L1   1500 20      64          1/1
```

3.1.4.2.5 Root for FabricPath Multi-Destination Trees

In FabricPath, multicast, broadcast and flooded traffic are forwarded along a multi-destination tree. FabricPath allows for multiple multi-destination trees in order to achieve traffic load balancing for multi-destination frames.

Two multi-destination trees are defined in Cisco FabricPath network by default, and multi-destination traffic is mapped to either of those trees for load-balancing purposes. The root of those multi-destination trees in the network should be explicitly set so as to provide an optimal topology.

Cisco FabricPath Intermediate Switch-to-Intermediate Switch (IS-IS) Protocol elects the switch with the highest configured root priority as the root for multi-destination tree 1. The switch with the second-highest root priority becomes the root for multi-destination tree 2. If there is no root priority configured, the other two parameters will be compared, system ID and switch ID, with higher values being better in all cases.

NVT has set the roots of the two multi-destination trees on two spine switches, one from each pair of vPC+ switches. If either of those switches fails, a replacement root would be elected out of all the FabricPath domain switches. This backup root should be configured in advance so that the system falls back to a predetermined topology in a failure scenario.

The Figure 17 shows the NVT FabricPath Root design for the multi-destination trees. Spine 54 has the highest root priority and is selected as the root of FTag 1 and Spine 52 has the second highest root priority and is selected as root of FTag 2.

Figure 17 NVT FabricPath Root Design for the Multi-Destination Trees

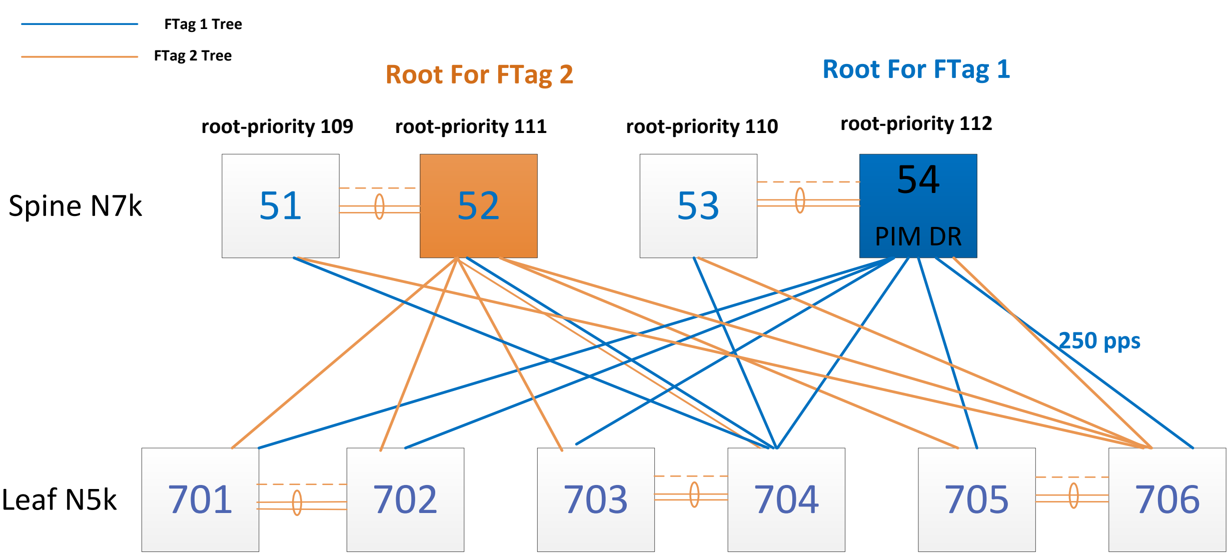

FTag trees are used as follows:

- FTag1 tree is used for unknown unicast, broadcast, and multicast.
- FTag2 tree is used only for multicast traffic.

To Verify FabricPath Multi-destination Tree Root:

```
DC6-DC102-54# show FabricPath isis topology summary
FabricPath IS-IS domain: default FabricPath IS-IS Topology Summary
MT-0
```

```
  Configured interfaces:  port-channel53  port-channel701  port-channel702  port-channel703  port-
channel704  port-channel705  port-channel706
  Number of trees: 2
    Tree id: 1, ftag: 1 [transit-traffic-only], root system: 0023.ac64.b2c3, 151
    Tree id: 2, ftag: 2, root system: 0023.ac64.b2c4, 153
```

To Verify which Multicast FTag Tree is Used in N7K:

```
DC6-DC102-54# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 102.11.27.1 dst-ip
130.102.0.1 vlan 12 module 7
128b Hash Key generated : 00 00 06 60 b1 b0 18 26 60 00 10 00 00 00 00 0c
0xc0
        FTAG SELECTED IS : 2

DC5-DC102-53# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 102.11.27.1 dst-ip
130.102.0.1 vlan 12 module 7
128b Hash Key generated : 18 26 60 00 10 00 00 00 00 0c 00 00 06 60 b1 b0
0xdc
        FTAG SELECTED IS : 1
```

To Verify which Multicast FTag Tree is Used in N5K:

```
 dc102-706# sh FabricPath load-balance multicast ftag-selected vlan 12 macg 0100.5e4d.0002

 Ftag selected : 1

=====================================

 Vlan : 12 (int_vlan : 39)
 Macg : 0100.5e4d.0002

 Hash-key : 0x00270000 00000000
 Hash-val : 210
 Num_trees : 2

=====================================
```

3.1.4.2.6 vPC+ for FabricPath

The NVT testbed is designed to have 2 pairs of vPC+ peers on the FabricPath spine and 3 pairs of vPC+ peers on the FabricPath leaf. The vPC+ peer-link must be configured as a FabricPath core link.

NVT FabricPath vPC+ Configuration:

```
N7K aggregation VDC 5:

!vPC+ configuration
feature vpc
vpc domain 111
  peer-switch
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  dual-active exclude interface-vlan 1,11-20,2001-
2010
  track 10
  auto-recovery
  FabricPath switch-id 100
  ip arp synchronize

!vPC+ member configuration
interface port-channel17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  vpc 17

!vPC+ peer link configuration
interface port-channel52
  switchport
  switchport mode FabricPath
  spanning-tree port type network
  vpc peer-link
  FabricPath isis metric 200

!vPC+ peer keepalive configuration
interface Ethernet1/19
  vrf member vpc-keepalive
  ip address 1.1.1.1/24
  no shutdown
```

```
N7K aggregation VDC 6:

!vPC+ configuration
feature vpc
vpc domain 111
  peer-switch
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  dual-active exclude interface-vlan 1,11-20,2001-
2010
  track 10
  auto-recovery
  FabricPath switch-id 100
  ip arp synchronize

!vPC+ member configuration
interface port-channel17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010,3951-3960
  vpc 17

!vPC+ peer link configuration
interface port-channel51
  switchport
  switchport mode FabricPath
  spanning-tree port type network
  vpc peer-link
  FabricPath isis metric 200

!vPC+ peer keepalive configuration
interface Ethernet4/16
  vrf member vpc-keepalive
  ip address 1.1.1.2/24
  no shutdown
```

To Verify the vPC+:

```
DC5-DC102-51# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 111
vPC+ switch id                    : 100
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
vPC FabricPath status             : peer is reachable through FabricPath

Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary, operational primary
Number of vPCs configured         : 2
Track object                      : 10
```

```
Peer Gateway                          : Disabled
Dual-active excluded VLANs            : 1,11-20,2001-2010
Graceful Consistency Check            : Enabled
Auto-recovery status                  : Enabled (timeout = 240 seconds)
FabricPath load balancing             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po52   up     1,5-7,10-20,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason           Active vlans  vPC+ Attribute
--   ----   ------ ----------- ------           ------------  --------------
17   Po17   up     success     success          1,11-20,2001- DF: Yes, FP
                                                2010,3001-301 MAC:
                                                0             100.11.4513
18   Po18   up     success     success          1,11-20,2001- DF: Yes, FP
                                                2010,3001-301 MAC:
                                                0             100.12.4513
                                                   0    DC5-DC102-51#
```

3.1.4.2.6.1 HSRP Active/Active with vPC+

Figure 18 HSRP Active/Active with vPC+

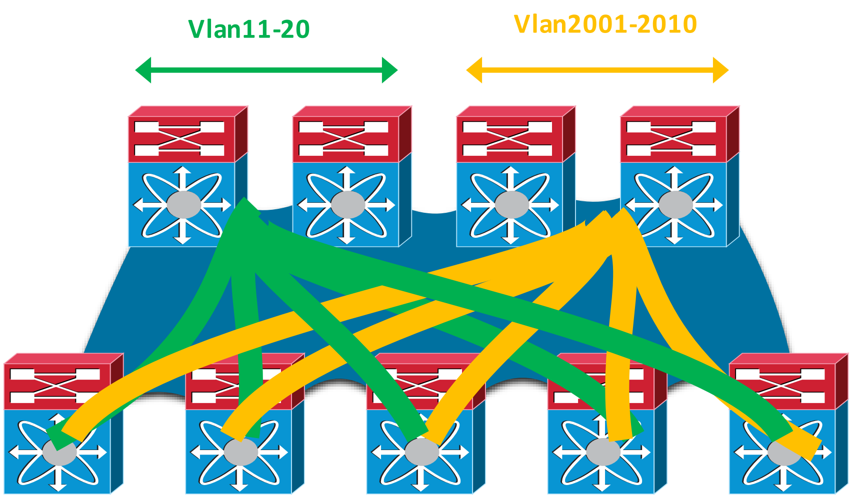

NVT has split HSRP VLANs among four spines with half the VLANs running HSRP between the first pair of spines and the other half running HSRP between the other pair of spines.

NVT spine HSRP configuration is as shown below; two HSRP groups with authentication and priority are configured for each VLAN:

```
DC5-DC102-51# show running-config interface vlan 11

!Command: show running-config interface Vlan11
!Time: Fri Mar  7 12:22:41 2014
```

```
version 6.2(6)

interface Vlan11
  no ip redirects
  ip address 102.11.0.51/16
 interface Vlan11
  no ip redirects
  ip address 102.11.0.51/16
  ip address 102.111.0.51/16 secondary
  ipv6 address 2001:1:102:11::51/64
  no ip ospf passive-interface
  ip router ospf 1 area 0.0.0.102
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 102.11.0.1
  hsrp 2
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 102.111.0.1
```

To Verify HSRP Peers and Virtual MAC Address on Nexus 7000 Spine:

```
DC5-DC102-51# show hsrp interface vlan 11
Vlan11 - Group 1 (HSRP-V2) (IPv4)
  Local state is Standby, priority 100 (Cfged 100), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 100
  Preemption Delay (Seconds) Minimum:120
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 1.434000 sec(s)
  Virtual IP address is 102.11.0.1 (Cfged)
  Active router is 102.11.0.52, priority 100 expires in 2.199000 sec(s)
  Standby router is local
  Authentication MD5, key-string "cisco"
  Virtual mac address is 0000.0c9f.f001 (Default MAC)
  4 state changes, last state change 01:47:34
  IP redundancy name is hsrp-Vlan11-1 (default)

Vlan11 - Group 2 (HSRP-V2) (IPv4)
  Local state is Standby, priority 100 (Cfged 100), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 100
  Preemption Delay (Seconds) Minimum:120
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 2.349000 sec(s)
  Virtual IP address is 102.111.0.1 (Cfged)
  Active router is 102.11.0.52, priority 100 expires in 1.616000 sec(s)
  Standby router is local
  Authentication MD5, key-string "cisco"
  Virtual mac address is 0000.0c9f.f002 (Default MAC)
  4 state changes, last state change 01:47:35
  IP redundancy name is hsrp-Vlan11-2 (default)


DC5-DC102-51# show hsrp brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr      Standby addr     Group addr
Vlan11      1   100  P Standby  102.11.0.52      local            102.11.0.1      (conf)
Vlan11      2   100  P Standby  102.11.0.52      local            102.111.0.1     (conf)

DC102-51# sh mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
```

```
       age - seconds since last seen,+ - primary entry using vPC Peer-Link,
       (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G 11       0000.0c9f.f001    static       -       F    F  sup-eth1(R)
G 11       0000.0c9f.f002    static       -       F    F  sup-eth1(R)
```

To Verify HSRP Virtual MAC on Nexus 5000 Edge Switches MAC Table:

```
dc102-705# sh mac address-table vlan 11
Legend:
       * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
       age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0000.0c9f.f001    dynamic   0          F    F  100.0.1054
* 11       0000.0c9f.f002    dynamic   0          F    F  100.0.1054
```

3.1.4.2.6.2 vPC+ Dual-Active Exclude

As a result of declaring the link that connects the spines as a vPC peer-link, the default behavior of vPC applies; if the peer-link goes down, the SVIs on the vPC secondary device are shut down.

In the context of FabricPath designs, this behavior is not beneficial, because the FabricPath links are still available, and there is no good reason to shut down the SVIs on the secondary. It is thus recommended to configure *dual-active exclude* for all the vPC+ vlans.

To Verify Dual-Active Exclude VLAN:

```
DC5-DC102-51# show vpc
Legend:
              (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 111
vPC+ switch id                    : 100
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
vPC FabricPath status             : peer is reachable through FabricPath
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary, operational primary
Number of vPCs configured         : 2
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : 1,11-20,2001-2010
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)
FabricPath load balancing         : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po52   up     1,5-7,10-20,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason           Active vlans  vPC+ Attribute
--   ----   ------ ----------- ------           ------------  --------------
17   Po17   up     success     success          1,11-20,2001- DF: Yes, FP
                                                2010,3001-301 MAC:
```

					0	100.11.4513
18	Po18	up	success	success	1,11-20,2001-2010,3001-3010	DF: Yes, FP MAC: 100.12.4513

3.1.4.2.7 FabricPath Region as Spanning Tree Root of the Network

On all the Cisco FabricPath switches that have Classic Ethernet ports, configure the same root priority using the *spanning-tree pseudo-information* command shown below. Verify that the root priority is the best (lowest) in the network so that the Cisco FabricPath region is the root of the spanning tree. If the Classic Ethernet edge ports receive a superior Bridge Protocol Data Unit (BPDU), those ports will be blocked from forwarding traffic. Also, those Classic Ethernet edge ports connected to the same Layer 2 non Cisco FabricPath domain, should be configured with the spanning-tree domain number. This approach will allow proper BPDU Propagation through the Cisco FabricPath network and help ensure a loop-free environment within that Layer 2 domain.

```
DC6-DC102-54(config)# spanning-tree pseudo-information
DC6-DC102-54(config-pseudo)# vlan 11-20 root priority 4096
```

3.1.4.2.8 Routed Multicast in FabricPath vPC+

PIM is enabled on the four Nexus 7000 spine VDCs with FabricPath VLANs configured under SVIs. It follows the same rules as all other non-FabricPath PODs. DC102 has defined all four spines as an auto-RP with Anycast RP/MSDP configured. From an operational perspective, it is advisable to align the PIM designated router (DR) priority with the HSRP primary.

3.1.4.3 FabricPath Load-Balancing and Verification

3.1.4.3.1 FabricPath Unicast Load-Balancing and Verification

Cisco NX-OS FabricPath unicast Layer 2 ISIS ECMP is on by default.

The default FabricPath unicast load balancing mechanism on the Nexus 7000 with F1/M1 line cards and the Nexus 5000 uses Layer 2/Layer 3/Layer 4 source and destination addresses and VLAN with symmetric hashing. To avoid hash polarization, each Cisco FabricPath switch automatically rotates the hash string by a number of bytes based on the system MAC address.

NVT has changed Nexus 7000 spine FabricPath unicast load-balancing mechanism using the following command and kept the Nexus 5000 FabricPath unicast load-balance as default.

F1/M1 VDC

```
! Change FabricPath load-balance on F1/M1VDC

DC102-52(config)# FabricPath load-balance source-destination
```

```
! Change FabricPath load-balance unicast on spine F1/M1 VDC

DC102-52(config)# FabricPath load-balance unicast layer3 include-vlan
```

```
! Verify Nexus 7000 spine FabricPath load-balance after modify

DC102-52(config)# sh FabricPath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: L3
```

```
Hash Control: Source-Destination
Rotate amount: 14 bytes
Use VLAN: TRUE


Ftag load-balancing configuration:
Hash Control: Source-Destination
Rotate amount: 14 bytes
Use VLAN: TRUE
```

In the NVT FabricPath network topology there are four equal cost paths from one leaf switch to any other leaf switch, except its vPC+ peer.

To Verify the FabricPath unicast ECMP path and load-balancing in leaf switch Nexus 5000, use the following commands.

Display Information About All FabricPath Topology Interfaces:

```
DC102-705# sh FabricPath topology interface
Interface          Topo-Description               Topo-ID    Topo-IF-State
------------------ ------------------------------ ---------- -------------
port-channel51     0                              0               Up
port-channel52     0                              0               Up
port-channel53     0                              0               Up
port-channel54     0                              0               Up
port-channel706    0                              0               Up
```

Display All FabricPath IS-IS Adjacency Information:

```
dc102-705# show FabricPath isis adjacency
FabricPath IS-IS domain: default FabricPath IS-IS adjacency database:
System ID       SNPA            Level  State  Hold Time  Interface
DC5-DC102-51    N/A             1      UP     00:00:28   port-channel51
DC102-52        N/A             1      UP     00:00:33   port-channel52
DC5-DC102-53    N/A             1      UP     00:00:22   port-channel53
DC102-54        N/A             1      UP     00:00:26   port-channel54
dc102-706       N/A             1      UP     00:00:26   port-channel706
```

Display the FabricPath Layer 2 IS-IS Routing Table for Unicast Routes:

```
dc102-705# sh FabricPath isis route
FabricPath IS-IS domain: default MT-0
Topology 0, Tree 0, Swid routing table
100, L1
 via port-channel51, metric 20
 via port-channel52, metric 20
151, L1
 via port-channel51, metric 20
152, L1
 via port-channel52, metric 20
153, L1
 via port-channel53, metric 20
154, L1
 via port-channel54, metric 20
200, L1
 via port-channel53, metric 20
 via port-channel54, metric 20
701, L1
 via port-channel53, metric 40
 via port-channel54, metric 40
703, L1
 via port-channel53, metric 40
```

```
 via port-channel54, metric 40
705, L1
 via port-channel706, metric 20
1701, L1
 via port-channel53, metric 40
 via port-channel54, metric 40
1702, L1
 via port-channel53, metric 40
 via port-channel54, metric 40
1703, L1
 via port-channel53, metric 40
 via port-channel54, metric 40
1704, L1
 via port-channel53, metric 40
 via port-channel54, metric 40
3402, L1
 via port-channel706, metric 20
```

Display Unicast Routes to Switch-ID 271:

```
dc102-706# sh l2 route switchid 701
FabricPath Unicast Route Table
'a/b/c' denotes ftag/switch-id/subswitch-id
'[x/y]' denotes [admin distance/metric]
ftag 0 is local ftag
subswitch-id 0 is default subswitch-id


FabricPath Unicast Route Table for Topology-Default

1/701/0, number of next-hops: 2
        via Po53, [115/40], 0 day/s 00:25:02, isis_FabricPath-default
        via Po54, [115/40], 0 day/s 01:11:13, isis_FabricPath-default
```

Display FabricPath Unicast Ftag Information:

```
DC102-705# sh FabricPath topology ftag unicast
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1
```

Display which Path the FabricPath Unicast Load-balancing Utilizes for a Given Flow:

```
DC102-705# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 151 dst-ip 101.11.7.1
Missing params will be substituted by 0's.

crc8_hash: 213
This flow selects interface Po51

DC102-705# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 151 dst-ip 101.11.7.2

Missing params will be substituted by 0's.

crc8_hash: 227
This flow selects interface Po53
```

3.1.4.3.2 FabricPath Multicast Load-Balancing and Verification

In the NVT FabricPath topology excerpt shown in Figure 19, the multicast traffic source is located on the L2 switch, 27, and the receiver is located on the L2 switch, 7003. Multicast traffic that reaches the spine 53, selects FTag 1 and uses tree 1 to forward the multicast data to the receiver which is attached to the leaf switch, 706. Note that the multicast traffic is also forwarded to all other spines because of PIM neighborship.

Figure 19 FabricPath Ftag 1 Multi-Destination Tree

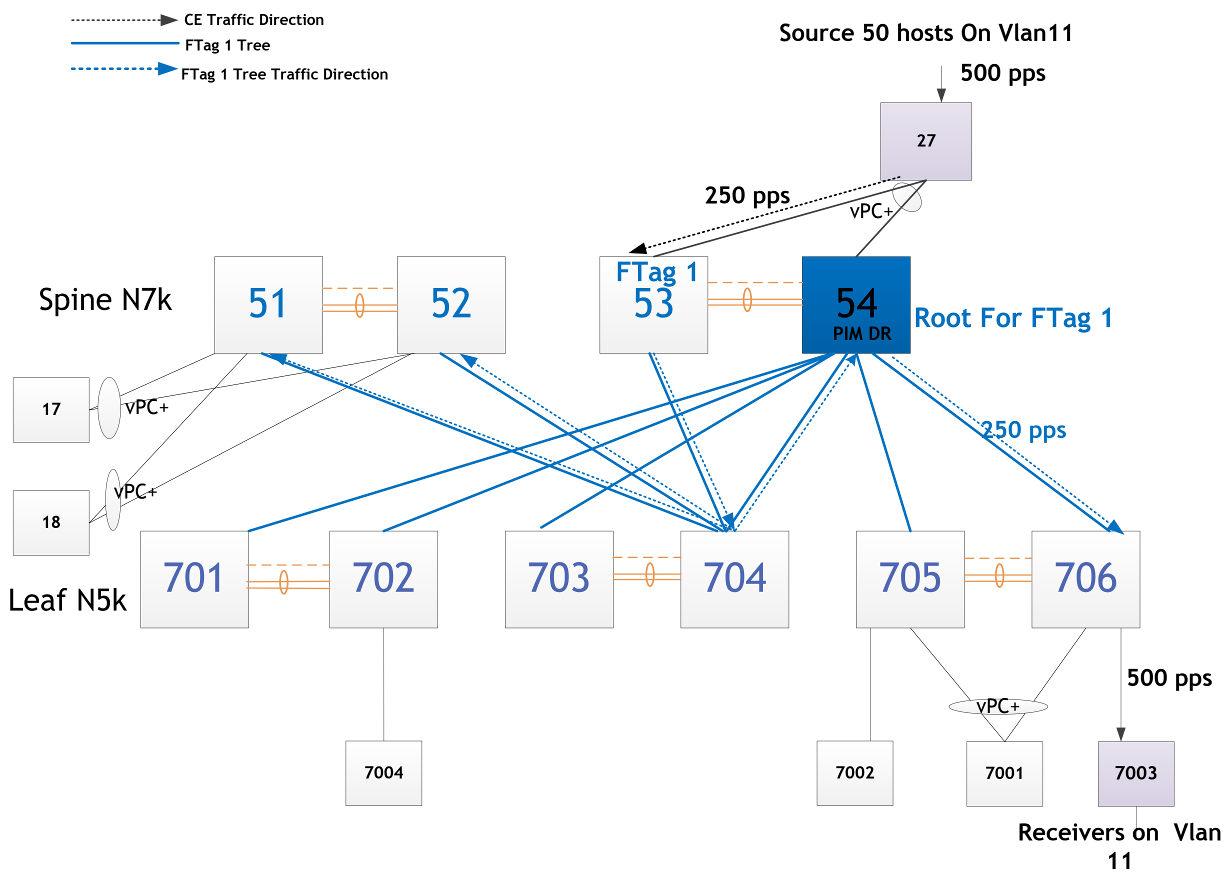

The Multicast traffic that reaches the spine 54, selects FTag 2 and uses tree 2 to forward the multicast data to the receiver which is attached to the leaf switch, 706. Note that this multicast traffic is also forwarded to all other spines because of PIM neighborship.

Figure 20 FabricPath Ftag 2 Multi-Destination Tree

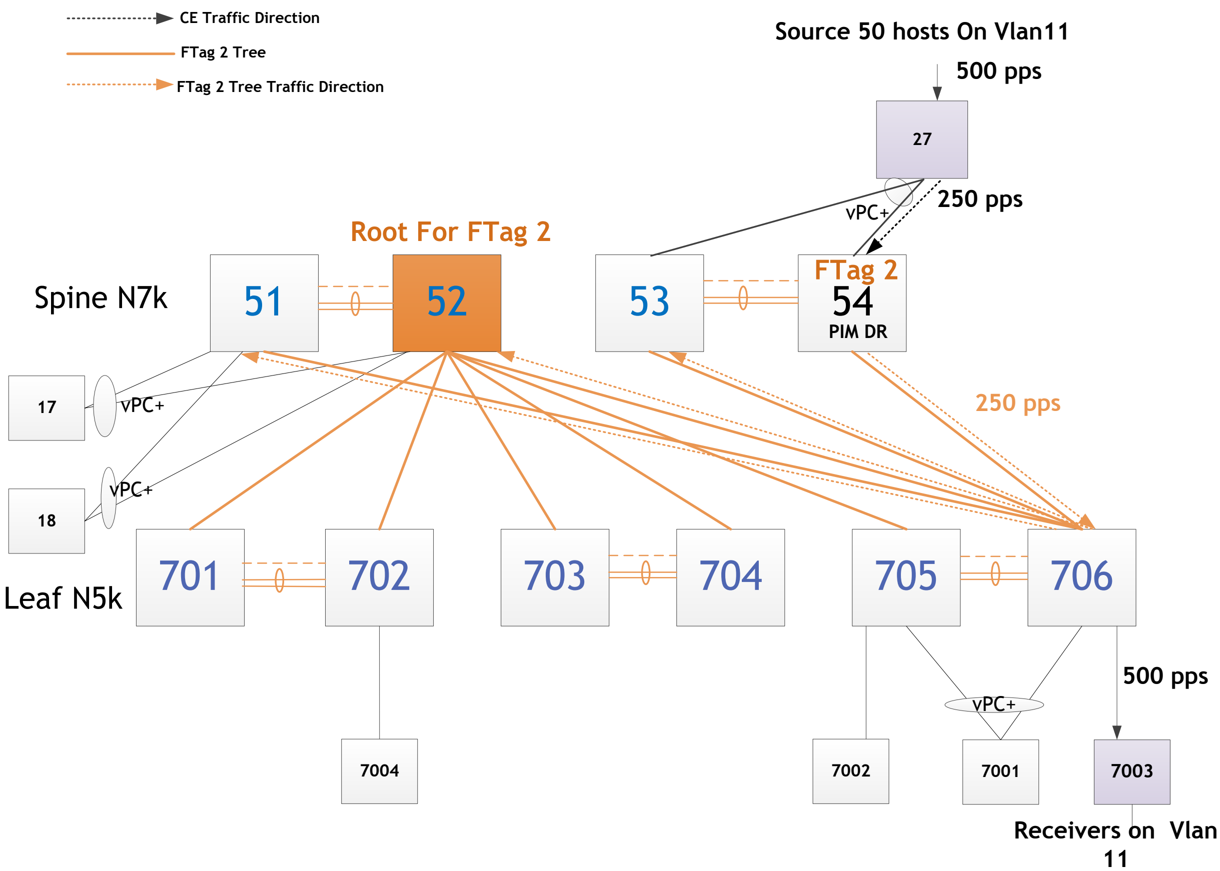

The hashing to either multi-destination tree is platform-dependent and the hash function is per flow. The default multicast load balancing mechanism for Nexus 7000 F1 VDC uses a symmetric hash input combining both Layer 3 (source and destination IP addresses) and Layer 4 (source and destination TCP and UDP port numbers, if present) information, as well as the VLAN ID. The default multicast load balancing mechanism for the Nexus 5000 uses symmetric hash with Layer 2/Layer 3/Layer 4 source and destination addresses, as well as VLAN ID.

NVT has kept the multicast load balancing mechanism on Nexus 7000 F1/M1 VDC and Nexus 5000 as default.

```
DC102-54# FabricPath load-balance multicast rotate-amount 0x5 include-vlan

DC102-54# sh run FabricPath all | in "multicast rotate"
FabricPath load-balance multicast rotate-amount 0x5 include-vlan

DC102-54# sh FabricPath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Hash Control: Symmetric
Rotate amount: 0 bytes
Use VLAN: TRUE
```

```
Ftag load-balancing configuration:
Hash Control: Symmetric
Rotate amount: 0 bytes
Use VLAN: TRUE

DC102-54#
```

To Verify the FabricPath multicast load-balancing path for a given multicast group in Nexus 7000, use the following commands.

Display the IP Multicast Routes for VLAN 11, Group 230.102.0.1:

```
DC102-54# sh FabricPath isis ip mroute vlan 11 group 230.102.0.1
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.102.0.1)
  Outgoing interface list: (count: 6)
    SWID: 0x97 (151)
    SWID: 0x98 (152)
    SWID: 0x99 (153)
    SWID: 0x69c (1692)
    SWID: 0x6a6 (1702)
    SWID: 0xd4a (3402)
```

Display FabricPath Multicast Routes for VLAN 11:

```
DC102-54# show FabricPath mroute vlan 11

(vlan/11, 0.0.0.0, 230.1.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:16:42, igmp
  Switch-id 151, uptime: 00:07:49, isis
  Switch-id 152, uptime: 00:16:47, isis
  Switch-id 153, uptime: 00:05:56, isis
  Switch-id 1692, uptime: 00:16:44, isis
  Switch-id 1702, uptime: 00:16:44, isis
  Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.2.0.1), uptime: 00:16:44, isis
 Outgoing interface list: (count: 5)
  Switch-id 151, uptime: 00:07:49, isis
  Switch-id 152, uptime: 00:16:41, isis
  Switch-id 1692, uptime: 00:16:44, isis
  Switch-id 1702, uptime: 00:16:44, isis
  Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.101.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:16:42, igmp
  Switch-id 151, uptime: 00:07:52, isis
  Switch-id 152, uptime: 00:16:47, isis
  Switch-id 153, uptime: 00:05:49, isis
  Switch-id 1692, uptime: 00:16:44, isis
  Switch-id 1702, uptime: 00:16:44, isis
  Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.102.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:16:42, igmp
  Switch-id 151, uptime: 00:07:54, isis
  Switch-id 152, uptime: 00:16:47, isis
  Switch-id 153, uptime: 00:05:49, isis
  Switch-id 1692, uptime: 00:16:44, isis
```

```
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.103.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
   Interface port-channel27, uptime: 00:16:42, igmp
   Switch-id 151, uptime: 00:07:57, isis
   Switch-id 152, uptime: 00:16:47, isis
   Switch-id 153, uptime: 00:05:49, isis
   Switch-id 1692, uptime: 00:16:44, isis
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.104.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
   Interface port-channel27, uptime: 00:16:42, igmp
   Switch-id 151, uptime: 00:07:49, isis
   Switch-id 152, uptime: 00:16:47, isis
   Switch-id 153, uptime: 00:05:56, isis
   Switch-id 1692, uptime: 00:16:44, isis
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.105.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
   Interface port-channel27, uptime: 00:16:42, igmp
   Switch-id 151, uptime: 00:07:49, isis
   Switch-id 152, uptime: 00:16:47, isis
   Switch-id 153, uptime: 00:05:49, isis
   Switch-id 1692, uptime: 00:16:44, isis
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.106.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
   Interface port-channel27, uptime: 00:16:42, igmp
   Switch-id 151, uptime: 00:07:52, isis
   Switch-id 152, uptime: 00:16:47, isis
   Switch-id 153, uptime: 00:05:56, isis
   Switch-id 1692, uptime: 00:16:44, isis
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, 0.0.0.0, 230.107.0.1), uptime: 00:16:47, isis igmp
 Outgoing interface list: (count: 7)
   Interface port-channel27, uptime: 00:16:42, igmp
   Switch-id 151, uptime: 00:07:41, isis
   Switch-id 152, uptime: 00:16:47, isis
   Switch-id 153, uptime: 00:05:49, isis
   Switch-id 1692, uptime: 00:16:44, isis
   Switch-id 1702, uptime: 00:16:44, isis
   Switch-id 3402, uptime: 00:16:44, isis

(vlan/11, *, *), Flood, uptime: 00:28:57, isis
 Outgoing interface list: (count: 9)
   Switch-id 151, uptime: 00:08:30, isis
   Switch-id 152, uptime: 00:28:57, isis
   Switch-id 153, uptime: 00:06:24, isis
   Switch-id 1692, uptime: 00:28:57, isis
   Switch-id 1701, uptime: 00:28:57, isis
   Switch-id 1702, uptime: 00:28:57, isis
   Switch-id 1703, uptime: 00:28:57, isis
   Switch-id 1704, uptime: 00:28:57, isis
   Switch-id 3402, uptime: 00:28:57, isis

(vlan/11, *, *), Router ports (OMF), uptime: 00:29:25, isis igmp
 Outgoing interface list: (count: 4)
   Switch-id 151, uptime: 00:08:21, isis
   Switch-id 152, uptime: 00:28:57, isis
```

```
  Switch-id 153, uptime: 00:06:24, isis
  Interface Vlan11, [SVI] uptime: 00:28:57, igmp
```

Display FabricPath Topology FTag Information:

```
DC102-54# sh FabricPath topology ftag multicast
Topo-Description         Topo-ID    Graph-ID  Ftag
------------------------ ---------- --------- ---------
0                        0          1         1
0                        0          2         2

DC102-54# sh FabricPath topology ftag active
Topo-Description         Topo-ID    Graph-ID  Ftag
------------------------ ---------- --------- ---------
0                        0          2         2
```

Display FabricPath Multicast Load-balancing Information:

```
DC6-DC102-54# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 102.11.27.1 dst-ip
130.102.0.1 vlan 12 module 7
128b Hash Key generated : 00 00 06 60 b1 b0 18 26 60 00 10 00 00 00 00 0c
0xc0
        FTAG SELECTED IS : 2
```

Display FabricPath Multicast Route for VLAN 11, Ftag 2:

```
DC102-54# sh FabricPath mroute vlan 11 ftag 2

(ftag/2, vlan/11, 0.0.0.0, 230.1.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.2.0.1), uptime: 00:26:39, isis
 Outgoing interface list: (count: 5)
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.101.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.102.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
```

```
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.103.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.104.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.105.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.106.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, 0.0.0.0, 230.107.0.1), uptime: 00:26:41, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 00:26:36, igmp
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, *, *), Flood, uptime: 00:38:51, isis
 Outgoing interface list: (count: 9)
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:33:54, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis

(ftag/2, vlan/11, *, *), Router ports (OMF), uptime: 00:39:19, isis igmp
 Outgoing interface list: (count: 4)
```

```
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:18:23, isis
  Interface port-channel701, uptime: 00:16:19, isis
  Interface Vlan11, [SVI] uptime: 00:38:51, igmp
```

To Verify the traffic path for a given multicast group in Nexus 5000 leaf switch, use the following commands.

Display the IP Multicast Routes for VLAN 11, Group 230.102.0.1:

```
dc102-706# sh FabricPath isis ip mroute vlan 11 group 230.102.0.1
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.102.0.1)
  Outgoing interface list: (count: 6)
    SWID: 0x97 (151)
    SWID: 0x98 (152)
    SWID: 0x99 (153)
    SWID: 0x9a (154)
    SWID: 0x69c (1692)
    SWID: 0x6a6 (1702)
```

Display FabricPath Multicast Routes for VLAN 11:

```
dc102-706# sh FabricPath mroute vlan 11

(vlan/11, 0.0.0.0, 230.1.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:31, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:38, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.2.0.1), uptime: 00:29:28, isis igmp
 Outgoing interface list: (count: 5)
  Switch-id 151, uptime: 00:20:31, isis
  Switch-id 152, uptime: 00:29:25, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis
  Interface port-channel3, uptime: 00:29:28, igmp

(vlan/11, 0.0.0.0, 230.101.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:35, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:31, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.102.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:36, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:31, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
```

```
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.103.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:39, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:31, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.104.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:31, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:38, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.105.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:31, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:31, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.106.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:35, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:38, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, 0.0.0.0, 230.107.0.1), uptime: 00:29:29, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:29:28, igmp
  Interface port-channel3, uptime: 00:29:28, igmp
  Switch-id 151, uptime: 00:20:24, isis
  Switch-id 152, uptime: 00:29:29, isis
  Switch-id 153, uptime: 00:18:31, isis
  Switch-id 154, uptime: 00:29:23, isis
  Switch-id 1692, uptime: 00:29:27, isis
  Switch-id 1702, uptime: 00:29:27, isis

(vlan/11, *, *), Flood, uptime: 2d21h, isis
 Outgoing interface list: (count: 9)
  Switch-id 151, uptime: 00:21:12, isis
  Switch-id 152, uptime: 00:43:45, isis
  Switch-id 153, uptime: 00:19:06, isis
  Switch-id 154, uptime: 00:41:39, isis
  Switch-id 1692, uptime: 2d21h, isis
  Switch-id 1701, uptime: 2d21h, isis
  Switch-id 1702, uptime: 2d21h, isis
  Switch-id 1703, uptime: 2d21h, isis
  Switch-id 1704, uptime: 2d21h, isis
```

```
(vlan/11, *, *), Router ports (OMF), uptime: 2d21h, isis
 Outgoing interface list: (count: 4)
  Switch-id 151, uptime: 00:21:04, isis
  Switch-id 152, uptime: 00:43:38, isis
  Switch-id 153, uptime: 00:19:06, isis
  Switch-id 154, uptime: 00:41:39, isis
```

Display FabricPath Topology FTag Information:

```
dc102-706# sh FabricPath topology ftag multicast
Topo-Description         Topo-ID    Graph-ID  Ftag
------------------------ ---------- --------- ---------
0                        0          1         1
0                        0          2         2

dc102-706# sh FabricPath topology ftag active
Topo-Description         Topo-ID    Graph-ID  Ftag
------------------------ ---------- --------- ---------
0                        0          2         2
```

Display FabricPath Multicast Load-balancing Information:

```
dc102-706# sh FabricPath load-balance multicast ftag-selected vlan 11 macg 0100.5e4d.0001

 Ftag selected : 1

====================================

 Vlan : 11 (int_vlan : 40)
 Macg : 0100.5e4d.0001

 Hash-key : 0x00280000 00000000
 Hash-val : 240
 Num_trees : 2

====================================
```

Display FabricPath Multicast Route for VLAN 11, ftag 1:

```
dc102-706# sh FabricPath mroute vlan 11 ftag 1

(ftag/1, vlan/11, 0.0.0.0, 230.1.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.2.0.1), uptime: 00:46:31, isis igmp
 Outgoing interface list: (count: 5)
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel3, uptime: 00:46:31, igmp

(ftag/1, vlan/11, 0.0.0.0, 230.101.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
```

```
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.102.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.103.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.104.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.105.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.106.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:32, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, 0.0.0.0, 230.107.0.1), uptime: 00:46:32, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel1, uptime: 00:46:31, igmp
  Interface port-channel3, uptime: 00:46:31, igmp
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
```

```
(ftag/1, vlan/11, *, *), Flood, uptime: 2d21h, isis
 Outgoing interface list: (count: 9)
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis

(ftag/1, vlan/11, *, *), Router ports (OMF), uptime: 2d21h, isis
 Outgoing interface list: (count: 4)
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:38:16, isis
  Interface port-channel51, uptime: 00:36:10, isis
  Interface port-channel51, uptime: 00:38:16, isis
```

3.1.5 Fabric Extenders (FEX)

The Fabric Extender integrates with its parent switch, which is a Cisco Nexus Series device, to allow automatic provisioning and configuration taken from the settings on the parent device.

The Fabric Interface is an uplink port that is designated for connection from the Fabric Extender to its parent switch. A fabric interface cannot be used for any other purpose. It must be directly connected to the parent switch. Multiple fabric interfaces can be combined together to form a port-channel fabric interface. Beginning with Cisco NX-OS Release 6.1(3), a minimum number of links for the FEX fabric port channel can be configured so that when a certain number of FEX fabric port-channel member ports go down, the host-facing interfaces of the FEX are suspended.

The host interfaces are Ethernet host interfaces for connection to a server or host system.

```
DC6-DC101-6# show running-config fex

!Command: show running-config fex
!Time: Fri Mar  7 12:25:04 2014

version 6.2(6)
feature-set fex

fex 101
  pinning max-links 1
  description FEX0101

! Port-channel fabric interface
interface port-channel101
  switchport
  switchport mode fex-fabric
  fex associate 101
  port-channel min-links 2

interface Ethernet10/1
  switchport
  switchport mode fex-fabric
  fex associate 101
  channel-group 101
  no shutdown

! Port-channel host interface
interface port-channel201
```

```
  switchport
  switchport access vlan 11
  spanning-tree port type edge
  spanning-tree bpdufilter enable
  flowcontrol send on
  vpc 201

interface Ethernet101/1/1
  switchport
  switchport access vlan 11
  logging event port link-status
  channel-group 201 mode active
  no shutdown
```

Display the Fabric Extenders Attached to the System:

```
DC101-6# sh fex
  FEX         FEX           FEX                       FEX
Number    Description      State            Model            Serial
------------------------------------------------------------------------
101        FEX0101               Online     N2K-C2224TP-1GE   SSI15480E4B
102        FEX0102               Online   N2K-C2248TP-E-1GE   SSI154005BN
103        FEX0103               Online     N2K-C2248TP-1GE   SSI161509VH
104        FEX0104               Online    N2K-C2232PP-10GE   SSI160700MV
105        FEX0105               Online    N2K-C2232PP-10GE   SSI16070CM8
```

Since the FEX host interfaces are supposed to be connected directly to hosts, certain defaults should be noted as shown below. Also, CDP is not supported on the Fabric Extenders connected to a Nexus 7000 parent switch.

```
DC6-DC101-6# show run int e101/1/1 all
interface Ethernet101/1/1
  no description
  lacp port-priority 32768
  lacp rate normal
  lldp transmit
  lldp receive
  switchport
  switchport mode access
  no switchport dot1q ethertype
  switchport access vlan 11
  switchport trunk native vlan 1
  switchport trunk allowed vlan 1-4094
  spanning-tree port-priority 128
  spanning-tree cost auto
  spanning-tree link-type auto
  spanning-tree port type edge
  spanning-tree bpduguard enable
  no spanning-tree bpdufilter
  speed auto
  duplex auto
  flowcontrol receive off
  flowcontrol send on
  link debounce time 100
  no beacon
  delay 1
  snmp trap link-status
  logging event port link-status
  logging event port trunk-status default
  medium broadcast
  channel-group 201 mode active
    lacp suspend-individual
  no ip dhcp snooping trust
  no ip dhcp snooping limit rate
  no ip arp inspection trust
```

```
  ip arp inspection limit rate 15 burst interval 5
  no ip verify source dhcp-snooping-vlan
  no shutdown
```

3.1.6 Unified Computing System (UCS) Overview

Cisco Unified Computing System (UCS) combines computing, networking, management, virtualization and storage access into a single integrated architecture.

3.1.6.1 UCS Management and Monitoring

Cisco Unified Computing System Manager (UCSM) is the management system for all components in a Cisco UCS domain and runs on the fabric interconnect (FI). Any of the interfaces available with this management service can be used to access, configure, administer and monitor the network and blade resources for all chassis connected to the Fabric Interconnect (FI).

Cisco UCS Manager includes the following user interfaces that can be used to manage a Cisco UCS domain:

- Cisco UCS Manager GUI
- Cisco UCS Manager CLI

NVT has provisioned out-of-band (OOB) networks for network infrastructure and virtual machine management.

3.1.6.1.1 Image Upgrade

NVT has configured NTP and time zones to ensure that the clocks on all UCS infrastructure and chassis components are synchronized.

The following issues may be encountered if the clocks are not synchronized:

- IOM may freeze during image upgrade (CSCuh25709/CSCuh25841/CSCuh87431).
    - To be addressed as part of feature enhancement: CSCtg28246 - *System / Fabric A&B Clock Set and Synchronization*.
    - Workaround: Manually OIR the failed IOM
- The upgrade procedure continuously retries step “Deploy Poll Activate Of Local FI” during the firmware install process (CSCui13535).
    - Workaround: Manually issue the install firmware command again with the ‘Force’ option enabled.

3.1.6.1.2 Syslogs

NVT has configured syslog to report to a centralized server.

Verification of Syslogs through the UCSM CLI:

```
UCS-FI-106-01-A# scope monitoring
UCS-FI-106-01-A /monitoring # show syslog
```

```
console
    state: Enabled
    level: Alerts

monitor
    state: Enabled
    level: Information

file
    state: Enabled
    level: Information
    name:  UCS-FI-106-01
    size:  4194304

remote destinations
    Name     Hostname             State    Level         Facility
    -------- -------------------- -------- ------------- --------
    Server 1 172.28.92.10         Enabled  Information   Local6
    Server 2 none                 Disabled Critical      Local7
    Server 3 none                 Disabled Critical      Local7

sources
    faults: Enabled
    audits: Enabled
    events: Enabled


UCS-FI-106-01-A /monitoring #
```

3.1.6.2 UCS Blade Management

In order to provision blade servers, service profiles need to be defined using policies and resource pools.

3.1.6.2.1 Service Profiles and Blade Policies

Service profiles must be created in order to provision compute services on the blade servers. All service profiles are configured with two static vNICs: vNIC0 for management and vNIC1 for data-plane traffic. Service profiles are made up of a set of policies and address pools, including the following:

- **Local Disk Policy** – NVT has configured RAID 1 when local disks are present. Any modifications to the disk policy may result in data loss.
- **BIOS Policy** – NVT has configured the BIOS policy enabling Virtualization Technology (VT) and Intel Directed IO for higher performance on the virtual machines deployed on the blade servers. These are required in order to leverage the performance advantage of VM-FEX.
- **Maintenance Policy** – NVT has configured the "UserACK" policy option instead of "Immediate". When this option is selected, the blade server is not immediately rebooted when changes to the service profile are made. Instead, the service profile will show the pending changes in its status field, and will wait for the administrator to manually acknowledge the changes to reboot the blade server.
- **Dynamic vNIC Connection Policy** – NVT has allocated 50 dynamic vNICs per blade server. Each of these dynamic vNIC connection policies has been configured with a "VMWarePassThru" adapter policy for performance and the "Protected" option to enable failover.
- **MAC Address Pools** – NVT has configured MAC Address Pools as part of the Service Profiles in order to provide addresses to the hypervisor or bare metal OS on the blade server.

Service profile templates facilitate the reuse and rapid-deployment of service profiles. There are two types of templates supported:

- **Initial template** – Service profiles created from an initial template inherit all the properties of the template. After the creation of a service profile from the template, any changes to the template no longer affect the replicated service profiles.
- **Updating templates** – Service profiles created from an updating template inherit all the properties of the template and remain connected to the template. Any changes to the template automatically update the service profiles created from the template.

NVT makes use of Updating templates in order to quickly propagate service profile changes to facilitate test configurations.

3.1.6.3 UCS Uplink Port Infrastructure

Any port on the Fabric Interconnect can be configured as either an uplink port or a server port. NVT makes use of End-Host Mode on the uplink ports.

Verification of End-Host Mode through the UCSM CLI:

```
UCS-FI-106-01-A# scope eth-uplink
UCS-FI-106-01-A /eth-uplink # show detail

Ethernet Uplink:
    Mode: End Host
    MAC Table Aging Time (dd:hh:mm:ss): Mode Default
    VLAN Port Count Optimization: Disabled
    Current Task:
UCS-FI-106-01-A /eth-uplink #
```

In End-Host mode, a single uplink port/port channel on each FI is chosen to be the receiver for broadcast, multicast and unknown-unicast traffic on all VLANs. This port is called the G-pinned port and is selected by the system.

Verification of the G-Pinned Port in UCSM CLI:

```
UCS-FI-106-01-A(nxos)# show platform software enm internal info vlandb id 11

vlan_id 11
-------------
Designated receiver: Po71
Membership:
Po71
UCS-FI-106-01-A(nxos)#
```

3.1.6.3.1 Uplink Port-Channels

Cisco UCS uses Link Aggregation Control Protocol (LACP) to bundle the uplink ports into a port channel. In order to maximize throughput from the FIs while also guaranteeing both high-availability and load-sharing to the upstream switches, NVT has configured up to eight ports per uplink port-channel.

NVT uses static pinning to assign VM data traffic to specific uplink port-channels. This configuration is done using LAN Pin Groups.

Verification of LAN Pin Groups through the UCSM CLI:

```
UCS-FI-106-01-A# scope eth-uplink
UCS-FI-106-01-A /eth-uplink # show pin-group expand
```

```
Ethernet Pin Group:
    Name: DC106-5-6

    Ethernet Pin Target:
        Fabric Endpoint
        ------ --------
        A      fabric/lan/A/pc-71
        B      fabric/lan/B/pc-72

    Name: Management

    Ethernet Pin Target:
        Fabric Endpoint
        ------ --------
        A      fabric/lan/A/phys-slot-1-port-32
        B      fabric/lan/B/phys-slot-1-port-32
UCS-FI-106-01-A /eth-uplink #
```

3.1.6.3.2 VLAN Configuration

NVT has configured *common/global* VLANs spanning across both Fabric Interconnects in a cluster. Note that VLANs with IDs from 3968 to 4043 and 4094 are reserved and cannot be created for data traffic.

Display Reserved VLANs on the FIs:

```
UCS-FI-106-01-A(nxos)# show vlan internal usage

VLAN         DESCRIPTION
---------    -------------------------------------------------------
3968-4031    Multicast
4032         Online diagnostics vlan1
4033         Online diagnostics vlan2
4034         Online diagnostics vlan3
4035         Online diagnostics vlan4
4036-4043    Reserved
4094         Reserved

UCS-FI-106-01-A(nxos)#
```

Verification of VLANs through the UCSM CLI:

```
UCS-FI-106-01-A(nxos)# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po71

Remote SPAN VLAN
----------------
Disabled

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------



UCS-FI-106-01-A(nxos)#
```

3.1.6.3.2.1 VLAN Groups

NVT configured an out-of-band management domain on separate VLAN groups for all deployed FIs and VMs. VLANs 2 and 3 are associated for management network interfaces. VLANs 11-20 and 2001-2010 are associated with data plane network interfaces.

Verification of VLAN Groups through the UCSM CLI:

```
UCS-FI-106-01-A# scope eth-uplink
UCS-FI-106-01-A /eth-uplink # show vlan-group

    Network Group:

    Name                 Size   Native VLAN Name     Native VLAN
    -------------------- ------ -------------------- -----------
    Data_Uplink          20

    Management_Uplink    2      vlan2                fabric/lan/net-vlan2
UCS-FI-106-01-A /eth-uplink #
```

3.1.6.4 UCS Server Port Infrastructure

In order to obtain maximum throughput from the IOMs, NVT has utilized eight connections from the FI to the IOM. All links from an individual IOM must connect to the same FI because intercrossed connections are not supported.

3.1.6.4.1 Chassis Discovery Policy with Port Channels

NVT has configured the minimum number of links needed to discover the chassis and set the Link Grouping Preference to Port Channel.

3.1.6.5 UCS Distributed Virtual Switches (DVS)

NVT has enabled VM-FEX and all inter-VLAN traffic is forwarded by the FI to the upstream gateway switch for routing.

NVT has configured DirectPath I/O to increase performance from the VMs through the hypervisor.

Distributed virtual switches created by UCSM cannot span across multiple FI clusters. The UCSM running on a FI cluster can only create and manage distributed virtual switches within that cluster (CSCuh38886).

3.1.6.5.1 Port Profiles

NVT has configured port profiles for each *common/global* VLAN so that all VM interfaces can be logically separated by VLAN ID.

Verification of Port Profiles through the UCSM CLI:

```
UCS-FI-106-01-A(nxos)# show port-profile brief
----------------------------------------------------------
Port                Profile  Conf   Eval   Assigned  Child
Profile              State   Items  Items  Intfs     Profs
----------------------------------------------------------
UCS_Vlan11          1        8      8      0         0
UCS_Vlan12          1        8      8      0         0
UCS_Vlan13          1        8      8      0         0
UCS_Vlan14          1        8      8      0         0
UCS_Vlan15          1        8      8      0         0
UCS_Vlan16          1        8      8      0         0
UCS_Vlan17          1        8      8      0         0
UCS_Vlan18          1        8      8      0         0
UCS_Vlan19          1        8      8      0         0
UCS_Vlan3           1        7      7      0         0
ucsm_internal_rackserver_portprofile 1        3      3      0         0
UCS-FI-106-01-A(nxos)#
```

3.2 DC2 NVT Network Implementation and Configuration

3.2.1 Configuration of Platform Specific Features

3.2.1.1 Licensing

Feature-based licenses enable specific feature sets for the physical device. Any feature not included in a license package is bundled with the Cisco NX-OS software.

License Usage on Nexus 7000 in DC2:

```
N7K-aggregation# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
MPLS_PKG                      Yes   -   In use Never        -
STORAGE-ENT                   Yes   -   Unused Never        -
VDC_LICENSES                  Yes   4   In use Never        -
ENTERPRISE_PKG                No    -   Unused              -
FCOE-N7K-F132XP               No    0   Unused              -
FCOE-N7K-F248XP               Yes   1   Rsrved Never        -
ENHANCED_LAYER2_PKG           Yes   -   Unused Never        -
SCALABLE_SERVICES_PKG         Yes   -   In use Never        -
TRANSPORT_SERVICES_PKG        Yes   -   In use Never        -
LAN_ADVANCED_SERVICES_PKG     Yes   -   In use Never        -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   In use Never        -
--------------------------------------------------------------------------------

N7k-core# sh license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
MPLS_PKG                      Yes   -   In use Never        -
STORAGE-ENT                   No    -   Unused              -
VDC_LICENSES                  Yes   4   Unused Never        -
ENTERPRISE_PKG                No    -   Unused              -
FCOE-N7K-F132XP               No    0   Unused              -
FCOE-N7K-F248XP               No    0   Unused              -
ENHANCED_LAYER2_PKG           No    -   Unused              -
SCALABLE_SERVICES_PKG         No    -   Unused              -
TRANSPORT_SERVICES_PKG        No    -   Unused              -
LAN_ADVANCED_SERVICES_PKG     No    -   Unused              -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   In use Never        -
--------------------------------------------------------------------------------
```

License Usage on Nexus 7700 in DC2:

```
N7700# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
STORAGE-ENT                   Yes   -   Unused Never        -
VDC_LICENSES                  Yes  12   In use Never        -
FCOE-N7K-F248XP               No    0   Unused             Grace 119D 0H
ENHANCED_LAYER2_PKG           Yes   -   In use Never        -
TRANSPORT_SERVICES_PKG        No    -   Unused              -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   Unused Never        -
--------------------------------------------------------------------------------
```

License Usage on Nexus 5000 in DC2:

```
N5K# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
```

```
FCOE_NPV_PKG                  No    -   Unused              -
FM_SERVER_PKG                 No    -   Unused              -
ENTERPRISE_PKG                No    -   Unused              -
FC_FEATURES_PKG               No    -   Unused              -
VMFEX_FEATURE_PKG             No    -   Unused Never        -
ENHANCED_LAYER2_PKG           Yes   -   In use Never        -
LAN_BASE_SERVICES_PKG         Yes   -   In use Never        -
LAN_ENTERPRISE_SERVICES_PKG   No    -   Unused              -
--------------------------------------------------------------------------------
```

License Usage on Nexus 6000 in DC2:

```
N6K# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
FCOE_NPV_PKG                  No    -   Unused              -
FM_SERVER_PKG                 No    -   Unused              -
ENTERPRISE_PKG                No    -   Unused              -
FC_FEATURES_PKG               No    -   Unused              -
VMFEX_FEATURE_PKG             Yes   -   Unused Never        -
ENHANCED_LAYER2_PKG           Yes   -   In use Never        -
LAN_BASE_SERVICES_PKG         Yes   -   In use Never        -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   Unused Never        -
--------------------------------------------------------------------------------
```

License Usage on Nexus 3548 in DC2:

```
N3548# sh license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
LAN_BASE_SERVICES_PKG         Yes   -   In use Never        -
ALGO_BOOST_SERVICES_PKG       No    -   Unused              -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   Unused Never        -
--------------------------------------------------------------------------------
```

3.2.1.2 Out-of-Band Management Network

DC2 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic. Specifically, DC2 testbed makes use of the mgmt0 ports on the Nexus devices on a separate management VRF.

Configuration:

```
interface mgmt0
  vrf member management
  ip address 10.2.2.15/16
```

3.2.1.3 Common Configurations

3.2.1.3.1 SSH and TACACS+

SSH is enabled in DC2 testbed to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration:

```
feature tacacs+


ip tacacs source-interface mgmt 0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
```

```
    server 172.28.92.17
    use-vrf management

dc2-4# sh ssh server
ssh version 2 is enabled

dc2-4# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
admin    ttyS0        Feb 17 09:56 17:16        7323
interop  pts/0        Feb 24 11:02   .          8402 (172.28.92.47) session=ssh *
```

3.2.1.3.2 CDP and LLDP

CDP and LLDP are pervasively used on the DC2 testbed for inter-device discovery. LLDP is used where CDP is not supported on links to UCS.

```
dc2-4# sh run cdp all

!Command: show running-config cdp all
!Time: Fri Feb 21 16:33:26 2014

version 6.2(8)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
no cdp format device-id system-name


interface Ethernet1/1
  cdp enable


dc2-4# sh cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform      Port ID
mgmt-sw2.interop.cisco.com
                    mgmt0          148    R S I     WS-C6509-E    Gig1/6
DC201-5.interop.cisco.com(TBM14343038)
                    Eth1/3         178    R S s     N7K-C7010     Eth1/2
DC201-6.interop.cisco.com(JAF1431DMTE)
                    Eth1/7         178    R S s     N7K-C7010     Eth2/1
DC202-51.interop.cisco.com(TBM14343038)
                    Eth1/11        178    R S s     N7K-C7010     Eth1/18
DC202-52.interop.cisco.com(JAF1431DMTE)
                    Eth1/15        176    R S s     N7K-C7010     Eth1/18
```

```
dc2-4# sh run lldp all

feature lldp

lldp holdtime 120
lldp reinit 2
lldp timer 30
lldp tlv-select port-description
lldp tlv-select system-name
lldp tlv-select system-description
lldp tlv-select system-capabilities
lldp tlv-select management-address
lldp tlv-select dcbxp
lldp tlv-select port-vlan
```

```
interface Ethernet1/1
  lldp transmit
  lldp receive

interface Ethernet1/2
  lldp transmit
  lldp receive
-

dc2-4# sh lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID           Local Intf      Hold-time  Capability  Port ID
DC201-6             Eth1/7          120        BR          Eth2/1
DC202-51            Eth1/11         120        BR          Eth1/18
DC202-52            Eth1/15         120        BR          Eth1/18
-
```

3.2.1.3.3 Syslog

Syslog is used to record all network events on the DC2 test bed. Whenever possible, NVT uses a separate management VRF for syslog.

Configuration:

```
logging server syslog.interop.cisco.com 5 use-vrf management facility local6

N7K# sh logging server
Logging server:                 enabled
{syslog.interop.cisco.com}
        server severity:        notifications
        server facility:        local6
        server VRF:             management
```

3.2.1.3.4 SNMP

SNMP is used for system monitoring in DC2 on Nexus 7000 switches. Scripts are used to poll the systems asynchronously during the course of all DC2 test execution.

Configuration:

```
snmp-server user test network-operator
snmp-server user admin network-admin auth md5 0xadaa2472f13e36349c13755305a1865b priv
0xadaa2472f13e36349c13755305a1865b localizedkey
snmp-server user snmpv3 network-operator auth md5 0x98174bd80aa4cffb9d4d3ddda1e83511 localizedkey
snmp-server user snmpv3 vdc-admin
snmp-server host 172.28.92.51 traps version 2c public
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-operator
snmp-server community private group network-admin
snmp-server community cisco group network-admin
snmp-server community interop group network-operator

dc2-4# sh snmp trap
--------------------------------------------------------------------------------
Trap type                       Description                  Enabled
--------------------------------------------------------------------------------
entity              : entity_mib_change                             Yes
```

```
entity                  : entity_module_status_change                  Yes
entity                  : entity_power_status_change                   Yes
entity                  : entity_module_inserted                       Yes
entity                  : entity_module_removed                        Yes
entity                  : entity_unrecognised_module                   Yes
entity                  : entity_fan_status_change                     Yes
entity                  : entity_power_out_change                      Yes
link                    : linkDown                                     Yes
link                    : linkUp                                       Yes
link                    : extended-linkDown                            Yes
link                    : extended-linkUp                              Yes
link                    : cieLinkDown                                  Yes
link                    : cieLinkUp                                    Yes
link                    : connUnitPortStatusChange                     Yes
link                    : delayed-link-state-change                    Yes
callhome                : event-notify                                 No
callhome                : smtp-send-fail                               No
cfs                     : state-change-notif                           No
cfs                     : merge-failure                                No
rf                      : redundancy_framework                         Yes
aaa                     : server-state-change                          No
license                 : notify-license-expiry                        Yes
license                 : notify-no-license-for-feature                Yes
license                 : notify-licensefile-missing                   Yes
license                 : notify-license-expiry-warning                Yes
hsrp                    : state-change                                 No
upgrade                 : UpgradeOpNotifyOnCompletion                  Yes
upgrade                 : UpgradeJobStatusNotify                       Yes
feature-control         : FeatureOpStatusChange                        No
sysmgr                  : cseFailSwCoreNotifyExtended                  No
rmon                    : risingAlarm                                  Yes
rmon                    : fallingAlarm                                 Yes
rmon                    : hcRisingAlarm                                Yes
rmon                    : hcFallingAlarm                               Yes
```

3.2.1.3.5 NTP

NTP is used to synchronize the clocks on all DC2 devices to provide consistent timestamps on all network logs and events.

Configuration:

```
ntp distribute
ntp server 172.28.92.1 use-vrf management
ntp commit


dc2-3# sh ntp status
Distribution : Enabled
Last operational state: No session

dc2-3# sh ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                 st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                 8   16      17   0.00121 management
```

3.2.1.3.6 SPAN

SPAN has been enabled on Nexus 7000 DC2 switches to provide packet captures to assist in network debugging.

Configuration:

```
monitor session 1
  source interface port-channel36 both
  destination interface Ethernet1/15
  destination interface Ethernet1/32
  no shut

DC2-4# sh monitor session 1
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Po36
    tx            : Po36
    both          : Po36
source VLANs      :
    rx            :
    tx            :
    both          :
source exception  :
filter VLANs      : filter not specified
destination ports : Eth1/15           Eth1/32

Feature       Enabled   Value   Modules Supported       Modules Not-Supported
-----------------------------------------------------------------------------
MTU-Trunc     No
rate-limit-rx No
rate-limit-tx No
Sampling      No
MCBE          No
L3-TX         -         -       1  2  5  7              -
RB span       No


Legend:
  MCBE  = Multicast Best Effort
  L3-TX = L3 Multicast Egress SPAN
  ExSP-X = Exception Span for type X (L3, FP, or misc)
```

3.2.1.3.7 DNS

DNS has been enabled to provide name lookup in this network.

Configuration:

```
ip domain-lookup
ip domain-name interop.cisco.com
ip domain-list cisco.com
ip domain-list interop.cisco.com
ip name-server 172.28.92.9 172.28.92.10

dc2-3# ping karo
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=61 time=0.789 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=61 time=0.903 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=61 time=0.743 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=61 time=0.854 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=61 time=0.721 ms

--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
```

3.2.1.3.8 UDLD

UDLD is used to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a device detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

Configuration:

```
feature udld

udld aggressive

dc2-4# sh udld neighbors
Port            Device Name    Device ID   Port ID         Neighbor State
-----------------------------------------------------------------------
Ethernet1/3     TBM14343038    1           Ethernet1/2     bidirectional
Ethernet1/7     JAF1431DMTE    1           Ethernet2/1     bidirectional
Ethernet1/11    TBM14343038    1           Ethernet1/18    bidirectional
Ethernet1/15    JAF1431DMTE    1           Ethernet1/18    bidirectional
Ethernet1/19    TBM14343038    1           Ethernet1/34    bidirectional
Ethernet1/23    JAF1431DMTE    1           Ethernet1/34    bidirectional
Ethernet1/25    01EF6E2680     1           Te2/1/2         bidirectional
Ethernet1/26    01EF6E2680     1           Te1/1/2         bidirectional
Ethernet1/29    01EF6E26C0     1           Te1/2           bidirectional
Ethernet1/30    01EF645D40     1           Te1/2           bidirectional
Ethernet1/33    FOX10491DY3    1           Te1/1           bidirectional
```

3.2.1.3.9 DHCP Relay

DHCP relay is enabled on the aggregation layer to provide IP address services to hypervisors and VMs running on UCS systems.

Configuration:

```
feature dhcp

service dhcp
ip dhcp relay
ipv6 dhcp relay

interface Vlan10
  ip dhcp relay address 94.253.253.2
  ip dhcp relay address 94.1.1.2

interface Vlan11
  ip dhcp relay address 94.253.253.2
  ip dhcp relay address 94.1.1.2

DC201-5# sh ip dhcp relay
DHCP relay service is enabled
Insertion of option 82 is disabled
Insertion of VPN suboptions is disabled
Insertion of cisco suboptions is disabled
Global smart-relay is disabled

Smart-relay is enabled on the following interfaces:
------------------------------------------------------
Subnet-broadcast is enabled on the following interfaces:
------------------------------------------------------
Helper addresses are configured on the following interfaces:
 Interface        Relay Address     VRF Name
 -------------    -------------     --------
 Vlan10            94.253.253.2
 Vlan10            94.1.1.2
 Vlan11            94.253.253.2
 Vlan11            94.1.1.2
 Vlan12            94.253.253.2
```

```
 Vlan12             94.1.1.2
 Vlan13             94.253.253.2
 Vlan13             94.1.1.2
```

3.2.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

When the switch comes up for the first time, there are multiple CoPP configuration templates that are presented: *strict, moderate, lenient* and *dense*. For DC2, the *lenient* template is configured.

Default Lenient CoPP on Nexus 7000 for Software Release 6.2.x as Used in DC2:

```
copp profile lenient

N7K# sh policy-map type control-plane name copp-system-p-policy-lenient

policy-map type control-plane copp-system-p-policy-lenient
    class copp-system-p-class-critical
      set cos 7
      police cir 36000 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-important
      set cos 6
      police cir 1400 kbps bc 1500 ms
        conform transmit violate drop
    class copp-system-p-class-multicast-router
      set cos 6
      police cir 2600 kbps bc 1000 ms
        conform transmit violate drop
    class copp-system-p-class-management
      set cos 2
      police cir 10000 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-multicast-host
      set cos 1
      police cir 1000 kbps bc 1000 ms
        conform transmit violate drop
    class copp-system-p-class-redirect
      set cos 1
      police cir 280 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-normal
      set cos 1
      police cir 680 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-ndp
      set cos 6
      police cir 680 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-normal-dhcp
      set cos 1
      police cir 1500 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-normal-dhcp-relay-response
      set cos 1
      police cir 1800 kbps bc 750 ms
        conform transmit violate drop
    class copp-system-p-class-exception
      set cos 1
      police cir 360 kbps bc 375 ms
        conform transmit violate drop
    class copp-system-p-class-monitoring
      set cos 1
```

```
      police cir 130 kbps bc 1500 ms
        conform transmit violate drop
    class copp-system-p-class-l2-unpoliced
      police cir 8 gbps bc 5 mbytes
        conform transmit violate transmit
    class copp-system-p-class-undesirable
      set cos 0
      police cir 32 kbps bc 375 ms
        conform drop violate drop
    class copp-system-p-class-fcoe
      set cos 6
      police cir 1060 kbps bc 1500 ms
        conform transmit violate drop
    class copp-system-p-class-l2-default
      police cir 1 kbps bc 375 ms
        conform transmit violate drop
    class class-default
      set cos 0
      police cir 1 kbps bc 250 ms
        conform transmit violate drop
```

Default CoPP on Nexus 5000 as Used in DC2:

```
N5K# show policy-map type control-plane name copp-system-policy-default

  policy-map type control-plane copp-system-policy-default
    class copp-system-class-igmp
      police cir 1024 kbps bc 65535 bytes
    class copp-system-class-pim-hello
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bridging
      police cir 20000 kbps bc 4800000 bytes
    class copp-system-class-arp
      police cir 1024 kbps bc 3600000 bytes
    class copp-system-class-dhcp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-mgmt
      police cir 12000 kbps bc 4800000 bytes
    class copp-system-class-lacp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-lldp
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-udld
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-isis
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-msdp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-cdp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-fip
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bgp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-eigrp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-exception
      police cir 64 kbps bc 4800000 bytes
    class copp-system-class-glean
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-hsrp-vrrp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-icmp-echo
      police cir 64 kbps bc 3600000 bytes
    class copp-system-class-ospf
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-pim-register
```

```
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-rip
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-l3dest-miss
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-mcast-miss
      police cir 256 kbps bc 3200000 bytes
    class copp-system-class-excp-ip-frag
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-same-if
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-ttl
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-default
      police cir 512 kbps bc 6400000 bytes
```

Default CoPP on Nexus 6000 as Used in DC2:

```
N6K# show policy-map type control-plane name copp-system-policy-default

  policy-map type control-plane copp-system-policy-default
    class copp-system-class-igmp
      police cir 1024 kbps bc 65535 bytes
    class copp-system-class-pim-hello
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bridging
      police cir 20000 kbps bc 4800000 bytes
    class copp-system-class-arp
      police cir 1024 kbps bc 3600000 bytes
    class copp-system-class-dhcp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-mgmt
      police cir 12000 kbps bc 4800000 bytes
    class copp-system-class-lacp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-lldp
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-udld
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-isis
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-msdp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-cdp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-fip
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bgp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-eigrp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-exception
      police cir 64 kbps bc 4800000 bytes
    class copp-system-class-glean
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-hsrp-vrrp
      police cir 1024 kbps bc 256000 bytes
    class copp-system-class-icmp-echo
      police cir 64 kbps bc 3600000 bytes
    class copp-system-class-ospf
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-pim-register
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-rip
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-l3dest-miss
      police cir 64 kbps bc 16000 bytes
```

```
    class copp-system-class-mcast-miss
      police cir 256 kbps bc 3200000 bytes
    class copp-system-class-excp-ip-frag
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-same-if
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-ttl
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-default
      police cir 512 kbps bc 6400000 bytes
    class copp-system-class-rpf-fail
      police cir 512 kbps bc 3200000 bytes
    class copp-system-class-mcast-last-hop
      police cir 512 kbps bc 3200000 bytes
    class copp-system-class-bfd
      police cir 9600 kbps bc 4800000 bytes
```

3.2.1.5 Rate Limiters

Rate limiters are an additional set of features on Nexus 7000 to prevent undesirable packets from overwhelming the CPU on the supervisor module.

Default Values:

```
dc2-4# sh hardware rate-limiter

Units for Config: packets per second
Allowed, Dropped & Total: aggregated since last clear counters


Module: 1
  R-L Class           Config          Allowed         Dropped            Total
 +------------------+--------+---------------+---------------+-----------------+
  L3 mtu                 500               0               0                0
  L3 ttl                 500             106               0              106
  L3 control           10000               0               0                0
  L3 glean               100               8               0                8
  L3 mcast dirconn    Disable
  L3 mcast loc-grp      3000               0               0                0
  L3 mcast rpf-leak      500               0               0                0
  L2 storm-ctrl       Disable
  access-list-log        100               0               0                0
  copy                 30000         3599100               0          3599100
  receive              30000         1961830               0          1961830
  L2 port-sec            500               0               0                0
  L2 mcast-snoop       10000               0               0                0
  L2 vpc-low            4000               0               0                0
  L2 l2pt                500               0               0                0
  f1 rl-1               4500                               0
  f1 rl-2               1000                               0
  f1 rl-3               1000                               0
  f1 rl-4                100                               0
  f1 rl-5               1500                               0
  L2 vpc-peer-gw        5000               0               0                0
  L2 lisp-map-cache     5000               0               0                0
  L2 dpss                100               0               0                0
  L3 glean-fast          100               0               0                0
  L2 otv                 100               0               0                0
  L2 netflow             500               0               0                0

  Port group with configuration same as default configuration
       Eth1/1-4       Eth1/5-8      Eth1/9-12     Eth1/13-16
     Eth1/17-20     Eth1/21-24     Eth1/25-28     Eth1/29-32
     Eth1/33-36     Eth1/37-40     Eth1/41-44     Eth1/45-48
```

3.2.1.1 VDCs and Resource Allocation

VDCs on the Nexus 7000 are used in the DC2 testbed to partition a single physical device into multiple logical devices that provide fault isolation, management isolation, address allocation isolation, service differentiation domains, and adaptive resource management.

```
DC6-sup2# show vdc

Switchwide mode is m1 f1 m1xl f2 m2xl f2e f3
vdc_id  vdc_name                    state        mac                  type       lc
------  --------                    -----        ----------           ---------  ------
1       DC6-sup2                    active       f8:66:f2:07:25:41    Ethernet   m1 f1 m1xl m2xl
2       DC201-6                     active       f8:66:f2:07:25:42    Ethernet   f2 f2e
3       DC202-52                    active       f8:66:f2:07:25:43    Ethernet   f2 f2e
4       DC202-54                    active       f8:66:f2:07:25:44    Ethernet   f2 f2e
```

Resource allocation for VDC's is done from the main VDC based on the requirements. The configuration used in the DC2 testbed is as shown below.

The Following Command Can Be Used to Help Estimate the VDC Resource Allocation:

```
DC5-sup2# show routing memory estimate routes 68000 next-hops 2
Shared memory estimates:
  Current max     96 MB;  33388 routes with 32 nhs
  Current max     96 MB;  29577 routes with 32 IPv6 nhs
          in-use   1 MB;     38 routes with  1 nhs (average)
          in-use   1 MB;     38 routes with  0 IPv6 nhs (average)
  Configured max  96 MB;  33388 routes with 32 nhs
  Configured max  96 MB;  29577 routes with 32 IPv6 nhs
  Estimate memory with fixed overhead:  26 MB;  68000 routes with  2 nhs and  0
```

Configuration:

```
vdc DC6-sup2 id 1
  limit-resource module-type m1 f1 m1xl m2xl
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  cpu-share 5
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
vdc DC201-6 id 2
  limit-resource module-type f2 f2e
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  cpu-share 5
  allocate interface Ethernet1/1-16
  allocate interface Ethernet2/1-16
  allocate interface Ethernet3/1-16
  boot-order 1
```

```
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
vdc DC202-52 id 3
  limit-resource module-type f2 f2e
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  cpu-share 5
  allocate interface Ethernet1/17-32
  allocate interface Ethernet2/17-32
  allocate interface Ethernet3/17-32
  allocate interface Ethernet7/1-48
  allocate interface Ethernet8/1-48
  allocate interface Ethernet9/1-48
  allocate interface Ethernet10/17-32
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
vdc DC202-54 id 4
  limit-resource module-type f2 f2e
  allow feature-set FabricPath
  allow feature-set fex
  allow feature-set mpls
  cpu-share 5
  allocate interface Ethernet1/33-48
  allocate interface Ethernet2/33-48
  allocate interface Ethernet3/33-48
  allocate interface Ethernet10/9-16,Ethernet10/33-48
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
```

3.2.2 Image Upgrade and Downgrade

DC2 makes use of ISSU/D to upgrade/downgrade software images whenever possible.

On the Nexus 7000, to check if the process will be disruptive or not, perform: *show install all impact system <system_image_name> kickstart <kickstart_image_name>:*

```
N7K-sup2# sh install all impact kickstart n7000-s2-kickstart.6.2.6.bin system n7000-s2-dk9.6.2.6.bin
Installer will perform impact only check. Please wait.

Verifying image bootflash:/n7000-s2-kickstart.6.2.6.bin for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n7000-s2-dk9.6.2.6.bin for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[####################] 100% -- SUCCESS

Extracting "lc1n7k" version from image bootflash:/n7000-s2-dk9.6.2.6.bin.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n7000-s2-dk9.6.2.6.bin.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n7000-s2-dk9.6.2.6.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/n7000-s2-kickstart.6.2.6.bin.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
2014 Feb 25 23:07:36.019 DC6-sup2 %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by admin on vsh.5902
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes  non-disruptive       rolling
     3       yes  non-disruptive       rolling
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset
     7       yes  non-disruptive       rolling
     8       yes  non-disruptive       rolling
     9       yes  non-disruptive       rolling
    10       yes  non-disruptive       rolling



Images will be upgraded according to following table:
Module       Image                  Running-Version(pri:alt)           New-Version  Upg-Required
------  ----------  ----------------------------------------  --------------------  ------------
     1      lc1n7k                                    6.2(2a)                6.2(6)           yes
     1        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)            no
     2      lc1n7k                                    6.2(2a)                6.2(6)           yes
     2        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)            no
     3      lc1n7k                                    6.2(2a)                6.2(6)           yes
     3        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)            no
     4      lc1n7k                                    6.2(2a)                6.2(6)           yes
     4        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)            no
     5      system                                    6.2(2a)                6.2(6)           yes
     5   kickstart                                    6.2(2a)                6.2(6)           yes
     5        bios   v2.12.0(05/29/2013):v2.12.0(05/29/2013)   v2.12.0(05/29/2013)            no
```

```
     6      system                                  6.2(2a)           6.2(6)        yes
     6   kickstart                                  6.2(2a)           6.2(6)        yes
     6        bios  v2.12.0(05/29/2013):v2.12.0(05/29/2013)  v2.12.0(05/29/2013)     no
     7      lc1n7k                                  6.2(2a)           6.2(6)        yes
     7        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)     no
     8      lc1n7k                                  6.2(2a)           6.2(6)        yes
     8        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)     no
     9      lc1n7k                                  6.2(2a)           6.2(6)        yes
     9        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)     no
    10      lc1n7k                                  6.2(2a)           6.2(6)        yes
    10        bios      v2.0.32(12/16/13):v2.0.32(12/16/13)    v2.0.22(06/03/13)     no
```

Running the command *show incompatibility-all system <image-name***>** will show the incompatible configuration and the necessary steps needed achieve non-disruptive upgrade/downgrade:

```
N7K-sup2# sh incompatibility-all system bootflash:n7000-s2-dk9.6.2.6.bin

Checking incompatible configuration(s) for vdc 'DC6-sup2':
-----------------------------------------------------------
No incompatible configurations

Checking dynamic incompatibilities for vdc 'DC6-sup2':
-------------------------------------------------------
No incompatible configurations

Checking incompatible configuration(s) for vdc 'DC201-6':
----------------------------------------------------------
No incompatible configurations

Checking dynamic incompatibilities for vdc 'DC201-6':
------------------------------------------------------
No incompatible configurations

Checking incompatible configuration(s) for vdc 'DC202-52':
-----------------------------------------------------------
No incompatible configurations

Checking dynamic incompatibilities for vdc 'DC202-52':
-------------------------------------------------------
No incompatible configurations

Checking incompatible configuration(s) for vdc 'DC202-54':
-----------------------------------------------------------
No incompatible configurations

Checking dynamic incompatibilities for vdc 'DC202-54':
-------------------------------------------------------
No incompatible configurations
```

The following caveats apply to ISSU/D:

1. When performing a software release upgrade or downgrade without ISSU in a system with FEX, the host interface configurations on the FEX will be lost after the reload to activate the new image. An extra step is required to reapply the configuration after the FEX module is fully online (*CSCuh58086*). A future FEX pre-provisioning feature will take care of this issue (*CSCuh57942*).
2. When performing ISSU process with OTV configuration, the following error was encountered: Conversion function failed for service "otv" (error-id 0xFFFFFFFF)
   With OTV configured, ISSU will be disruptive and requires shutting down the overlay interface. An enhancement request has been filed to place a configuration compatibility check and throw a message to disallow the procedure until the overlay interface is shutdown (*CSCug73006*).

3.2.3 Routing Design Overview

3.2.3.1 Unicast

3.2.3.1.1 BGP Routing Design

From edge/core switches to public cloud, DC2 has enabled eBGP configuration to establish peering between data center autonomous systems and public cloud autonomous systems to exchange routing updates. BGP policy has been applied to the eBGP peering configuration to control route updates between peers.

DC2 has been configured with route maps to filter the redistribution of OSPF routes from the testbed into BGP. The filters are configured based on IP prefix matching.

NSF is a high availability feature on modular switches running NX-OS or IOS with a redundant supervisor. On the Nexus 7000, data packets are forwarded by the hardware forwarding engines on the linecards. These engines are programmed with information learned from the routing control plane running on the supervisors. If the active supervisor were to fail, the forwarding tables on the linecards are preserved. All interface states are also preserved while the standby supervisor takes over active control of the system. This high availability system prevents any drop in traffic during the failure of the active control plane.

BGP graceful restart is a BGP feature that prevents disruption to the control and data plane. It allows for the graceful recovery of BGP sessions after a peer has failed. When combined with the NSF feature, any GR capable peers connected to a switch going through supervisor switchover will continue to forward traffic seamlessly.

Nonstop Forwarding (NSF) and graceful restart (GR) for BGP are enabled by default on NX-OS. SSO/NSF and graceful restart must be explicitly enabled for the system and for BGP, respectively, for Catalyst 6500 and 4500 running IOS.

DC2 BGP Configuration:

```
feature bgp

router bgp 200
  router-id 40.2.0.15
  graceful-restart stalepath-time 360
  log-neighbor-changes
  address-family ipv4 unicast
    redistribute direct route-map CONN
    redistribute ospf 2 route-map CONN
    maximum-paths 8
    maximum-paths ibgp 8
  neighbor 40.90.201.11 remote-as 100090
    address-family ipv4 unicast
      prefix-list NO_SELF in
  neighbor 40.90.203.13 remote-as 100090
    address-family ipv4 unicast
      prefix-list NO_SELF in
```

3.2.3.1.2 OSPF Routing Design

OSPF has been chosen as the IGP routing protocol for DC2 testbed. OSPF has been deployed from Core to Aggregation to L3 Access in DC2 data center.

DC2 core switches are configured as backbone Area 0. Each aggregation-access block is configured as a different non-backbone area. The multi-area design reduces computational work for OSPF routers during a topology change.

DC2 OSPF configuration:

```
feature ospf
router ospf 2
  router-id 40.2.0.15
  redistribute bgp 200 route-map BGPCORE-TO-DC2
  log-adjacency-changes
  timers throttle spf 100 200 5000
  timers throttle lsa 50 100 300
  auto-cost reference-bandwidth 1000000

interface loopback0
  ip router ospf 2 area 0.0.0.0

interface loopback1
  ip router ospf 2 area 0.0.0.0

interface port-channel15
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 3 a667d47acc18ea6b
  ip ospf network point-to-point
  ip router ospf 2 area 0.0.0.201
```

3.2.3.1.2.1 OSPF Router-ID

Each switch in the OSPF routing domain is identified by a Router ID. DC2 testbed has been configured with a loopback interface IP address as OSPF Router-ID for each switch in the testbed to identify each OSPF instance. If there is no OSPF Router-ID, NX-OS will choose the available loopback IP address as OSPF Router-ID and if there is no loopback address available, NX-OS will choose the highest interface IP address as OSPF Router-ID. If the interface IP address is used as the OSPF Router-ID, it will cause routing re-convergence when that interface goes down.
Router-ID is configured per OSPF process instance. NVT DC2 testing only creates one instance per VDC.

To Verify the OSPF Router-ID:

```
dc2-3# show ip ospf

 Routing Process 2 with ID 40.2.0.15 VRF default
 Routing Process Instance Number 1

DC201-5# show ip ospf neighbors
 OSPF Process ID 2 VRF default
 Total number of neighbors: 7
 Neighbor ID     Pri State            Up Time  Address         Interface
 40.2.0.15         1 FULL/ -          1d19h    40.201.1.15     Po3
```

3.2.3.1.2.2 OSPF Reference Bandwidth

The default OSPF Auto-Cost reference bandwidth for calculating OSPF metric is 40Gbps for NX-OS and 100Mbps for IOS. The reference bandwidth should be configured to be the same across the entire network; DC2 has been configured with 100Gbps as the reference bandwidth.

To Verify OSPF Reference Bandwidth:

```
dc2-3# show ip ospf

 Routing Process 2 with ID 40.2.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 This router is an area border and autonomous system boundary.
 Redistributing External Routes from bgp-200
 Administrative distance 110
 Reference Bandwidth is 1000000 Mbps
```

3.2.3.1.2.3 OSPF Network Type

NVT has configured DC2 testbed with point-to-point OSPF Network Type on all interfaces between the core and aggregation switches. It removes the OSPF designated router and backup designated router (DR/BDR) election and reduces the OSPF neighbor adjacency negotiation process.

To Verify OSPF Point-to-Point OSPF Network:

```
dc2-3# show ip ospf interface po15
 port-channel15 is up, line protocol is up
    IP address 40.201.1.15/24, Process ID 2 VRF default, area 0.0.0.201
    Enabled by interface configuration
    State P2P, Network type P2P, cost 50
    BFD is enabled
    Index 1, Transmit delay 1 sec
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:08
      LSU timer due in 00:00:00
    Message-digest authentication, using key id 1
    Number of opaque link LSAs: 0, checksum sum 0

dc2-3# sh ip ospf neighbors
 OSPF Process ID 2 VRF default
 Total number of neighbors: 13
 Neighbor ID     Pri State            Up Time  Address         Interface
 40.201.0.19       1 FULL/ -          03:28:34 40.201.1.19     Po15
 40.201.0.21       1 FULL/ -          1d20h    40.201.2.21     Po16
```

3.2.3.1.2.4 OSPF Authentication

Cisco NX-OS supports two authentication methods, simple password authentication and MD5 authentication digest. Authentication can be configured for an OSPFv2 area or per interface.

DC2 has been configured with MD5 authentication for each interface.

To Verify OSPF Authentication:

```
dc2-3# show ip ospf interface p15
 port-channel15 is up, line protocol is up
    IP address 40.201.1.15/24, Process ID 2 VRF default, area 0.0.0.201
    Enabled by interface configuration
```

```
    State P2P, Network type P2P, cost 50
    BFD is enabled
    Index 1, Transmit delay 1 sec
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:01
    Message-digest authentication, using key id 1
    Number of opaque link LSAs: 0, checksum sum 0
```

3.2.3.1.2.5 Route Redistribution

Route redistribution is configured on the Core/Edge switches for DC2 to learn routes from BGP. Route maps are used to control which external routes are redistributed. DC2 has been configured IP prefix-list to filter IP addresses.

3.2.3.1.2.6 OSPF High Availability and Graceful Restart

Cisco provides multilevel high-availability architecture for OSPF: Non Stop Routing (NSR) and Graceful Restart (GR) with NSF.

With NSR, OSPF preserves the running state of the protocol data and sessions in persistent memory. If the OSPF application fails or needs to be restarted for any reason, it will restart from the preserved state to ensure that there is no disruption seen by any of its OSPF peers. The internal applications that manage the routing table and hardware forwarding tables will also not experience any failure, allowing for non-disruptive OSPF process restarts.

OSPF GR and NSF allow for non-disruptive failure of the supervisor on Cisco modular switches. On the Nexus 7000, the hardware routing engines are programed per linecard. On active supervisor failure, the forwarding tables on the linecards are preserved while the standby supervisor takes over active control of the system. There is no disruption to packet forwarding during this process. GR prevents OSPF peers from restarting during a supervisor failure; thus, preserving their packet forwarding states. The combination of OSPF GR and SSO/NSF allows the entire network to continue operating seamlessly during a supervisor failure.

OSPF NSR and graceful restart are enabled by default on NX-OS. SSO/NSF and graceful restart must be explicitly enabled for the system and for OSPF, respectively, for Catalyst 6500 and 4500 running IOS.

To Verify OSPF Graceful Restart:

```
dc2-3# sh ip ospf
 Routing Process 2 with ID 40.2.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
```

3.2.3.1.2.7 Passive Interfaces

All servers/hosts facing SVIs (Switched Virtual Interfaces) are configured as OSPF passive interfaces. This is to ensure that server farm subnets are advertised into OSPF, while preventing the formation of unnecessary OSPF adjacencies through the access layer.

To Verify OSPF Passive Interface:

```
DC201-5# sh ip ospf interface vlan 12
 Vlan12 is up, line protocol is up
    IP address 201.12.0.19/16, Process ID 2 VRF default, area 0.0.0.201
    Enabled by interface configuration
    State DR, Network type BROADCAST, cost 1000
    Index 9, Passive interface
```

3.2.3.1.2.8 OSPF Timers and Optimization

NVT has kept the OSPF hello/hold timers at their default values on DC2. This allows other resilience features such as SSO/NSF to provide high availability. BFD should be used for networks where fast peer failure detection is desired.

To Verify OSPF Timers and Optimization:

```
dc2-3# sh ip ospf

 Routing Process 2 with ID 40.2.0.15 VRF default
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart is configured
   Grace period: 60 state: Inactive
   Last graceful restart exit status: None
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 This router is an area border and autonomous system boundary.
 Redistributing External Routes from
   bgp-200
 Administrative distance 110
 Reference Bandwidth is 1000000 Mbps
 SPF throttling delay time of 100.000 msecs,
   SPF throttling hold time of 200.000 msecs,
   SPF throttling maximum wait time of 5000.000 msecs
 LSA throttling start time of 50.000 msecs,
   LSA throttling hold interval of 100.000 msecs,
   LSA throttling maximum wait time of 300.000 msecs
 Minimum LSA arrival 1000.000 msec
 LSA group pacing timer 10 secs
 Maximum paths to destination 8
 Number of external LSAs 74, checksum sum 0x235d07
 Number of opaque AS LSAs 0, checksum sum 0
 Number of areas is 8, 8 normal, 0 stub, 0 nssa
 Number of active areas is 8, 8 normal, 0 stub, 0 nssa
 Install discard route for summarized external routes.
 Install discard route for summarized internal routes.
 BFD is enabled
   Area BACKBONE(0.0.0.0) (Inactive)
        Area has existed for 2d19h
        Interfaces in this area: 3 Active interfaces: 3
        Passive interfaces: 0  Loopback interfaces: 2
        No authentication available
        SPF calculation has run 195 times
         Last SPF ran for 0.000168s
        Area ranges are
        Number of LSAs: 460, checksum sum 0xeb7312
   Area (0.0.0.201)
        Area has existed for 2d19h
        Interfaces in this area: 2 Active interfaces: 2
        Passive interfaces: 0  Loopback interfaces: 0
        No authentication available
        SPF calculation has run 195 times
         Last SPF ran for 0.001450s
        Area ranges are
        Number of LSAs: 790, checksum sum 0x165a7ba
```

```
dc2-3# sh ip ospf interface port-channel 15
 port-channel15 is up, line protocol is up
    IP address 40.201.1.15/24, Process ID 2 VRF default, area 0.0.0.201
    Enabled by interface configuration
    State P2P, Network type P2P, cost 50
    BFD is enabled
    Index 1, Transmit delay 1 sec
    1 Neighbors, flooding to 1, adjacent with 1
    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello timer due in 00:00:03
    Message-digest authentication, using key id 1
    Number of opaque link LSAs: 0, checksum sum 0
```

3.2.3.2 Unicast Forwarding Verification

On NX-OS platforms, routing is performed using hardware forwarding engines. The following sequence of commands illustrates verification of the programming of a host on a directly connected subnet on the Nexus 7000.

This Switch is the Authoritative Router for a Directly Connected Subnet on VLAN 11: 10.11.0.0/16:

```
DC201-5# show running-config interface vlan 11

interface Vlan11
  no ip redirects
  ip address 201.11.0.19/16
  ip address 201.111.0.19/16 secondary
  ipv6 address 2001:1:201:11::19/64
  no ipv6 redirects
  ip router ospf 2 area 0.0.0.201
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 201.11.0.1
  hsrp 2
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 201.111.0.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 2001:1:201:11::1
  ip dhcp relay address 94.253.253.2
  ip dhcp relay address 94.1.1.2
  no shutdown
```

The Host 201.11.7.1 has been Learned via ARP on this Subnet:

```
DC201-5# show ip arp 201.11.7.1

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
201.11.7.1      00:18:06  00c9.0b07.0100  Vlan11
```

On NX-OS, “show ip route” will also Show Directly Connected Hosts as /32 Routes:

```
DC201-5# show ip route 201.11.7.1

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

201.11.7.1/32, ubest/mbest: 1/0, attached
    *via 201.11.7.1, Vlan11, [250/0], 21:00:05, am
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
DC201-6# show ip adjacency 201.11.7.1

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
201.11.7.1      00c9.0b07.0100  50   arp        Vlan11
```

Find the PO Interface on which this MAC Address is Learnt:

```
DC201-6# sh mac address-table address 00c9.0b07.0100
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       00c9.0b07.0100    dynamic   0          F    F  Po7
```

Display PO7 Member Interface with Module Information:

```
DC201-6# sh port-channel summary | in Po7
7     Po7(SU)     Eth      LACP      Eth2/7(P)    Eth3/7(P)
```

Display Adjacency Index for this Route in Hardware Table:

```
DC201-6# sh system internal forwarding ip route 201.11.7.1 module 2

Routes for table default/base

----+---------------------+----------+----------+-----------
Dev | Prefix              | PfxIndex | AdjIndex | LIF
----+---------------------+----------+----------+-----------
  0   201.11.7.1/32          0x46e3       0xd7     0x16
  1   201.11.7.1/32          0x46e3       0xd7     0xc
  2   201.11.7.1/32          0xb2e3       0xd7     0xc
  3   201.11.7.1/32          0xb2e3       0xd7     0xc
```

Display DMAC Entry Programmed in Adjacency Table:

```
DC201-6# sh system internal forwarding adjacency module 2 entry 0xd7 detail
Device: 0   Index: 0xd7      DMAC: 00c9.0b07.0100 SMAC: f866.f207.2542
            LIF: 0x16 (Vlan11) DI: 0x0     ccc: 4   L2_FWD: NO  RDT: NO
            packets: 1137bytes: 109152zone enforce: 0
Device: 1   Index: 0xd7      DMAC: 00c9.0b07.0100 SMAC: f866.f207.2542
            LIF: 0xc (Vlan11) DI: 0x0     ccc: 4   L2_FWD: NO  RDT: NO
            packets: 0   bytes: 0   zone enforce: 0
Device: 2   Index: 0xd7      DMAC: 00c9.0b07.0100 SMAC: f866.f207.2542
```

```
              LIF: 0xc (Vlan11) DI: 0x0     ccc: 4   L2_FWD: NO  RDT: NO
              packets: 0   bytes: 0   zone enforce: 0
Device: 3     Index: 0xd7      DMAC: 00c9.0b07.0100 SMAC: f866.f207.2542
              LIF: 0xc (Vlan11) DI: 0x0     ccc: 4   L2_FWD: NO  RDT: NO
              packets: 0   bytes: 0   zone enforce: 0
```

Display Allocated Bridge Domain Matches in the Hardware Table:

```
DC201-6# sh vlan internal bd-info vlan-to-bd 11

VDC Id  Vlan Id  BD Id
------  -------  -------
2         11       4112
```

Display LTL Entry for this MAC Address Associated with the Bridge Domain:

```
DC201-6# sh hardware mac address-table 2 vlan 11 address 00c9.0b07.0100
FE | Valid| PI| BD  |      MAC       |  Index| Stat| SW  | Modi| Age| Tmr| GM| Sec| TR| NT| RM| RMA| Cap| Fld|Always| PV | RD| NN| UC|PI_E8| VIF | SWID| SSWID| LID
   |      |   |     |                |       | ic  |     | fied|Byte| Sel|   | ure| AP| FY|   |    |TURE|    | Learn|    |   |   |   |     |     |     |      |
---+------+---+-----+----------------+-------+-----+-----+-----+----+----+---+----+---+---+---+----+----+----+------|----|---|---|---|-----|-----|-----|------|-------
 0    1     0  4112  00c9.0b07.0100  0x0040a   0   0x089   1    236   1    0    0   0   0   0    0    0    0     0   0x00  0   0   0    0   0x000 0x000  0x000 0x0040a
 1    1     1  4112  00c9.0b07.0100  0x0040a   0   0x089   1    236   1    0    0   0   0   0    0    0    0     0   0x00  0   0   1    0   0x000 0x000  0x000 0x0040a
 2    1     0  4112  00c9.0b07.0100  0x0040a   0   0x009   0    109   1    0    0   0   0   0    0    0    0     0   0x00  1   0   0    0   0x000 0x000  0x000 0x0040a
 3    1     0  4112  00c9.0b07.0100  0x0040a   0   0x009   0    109   1    0    0   0   0   0    0    0    0     0   0x00  1   0   0    0   0x000 0x000  0x000 0x0040a
```

Display DMAC Sent to LTL Index for PO7:

```
DC201-6# sh system internal pixm info ltl 0x0040a

PC_TYPE    PORT    LTL      RES_ID       LTL_FLAG      CB_FLAG      MEMB_CNT
------------------------------------------------------------------------------
Normal     Po7     0x040a    0x16000006   0x00000000    0x00000002   2
```

3.2.3.3 Multicast Routing Design

Multicast routing has been enabled across the entire NVT network on DC2. On NX-OS, multicast routing is enabled by default, while it needs to be explicitly enabled on IOS.

DC2 Multicast Configuration:

```
feature pim
ip pim rp-address 40.2.50.1 group-list 230.2.0.0/16
ip pim rp-address 40.2.50.1 group-list 239.1.1.1/32
ip pim send-rp-announce loopback1 group-list 230.201.0.0/16
ip pim send-rp-discovery loopback1
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt

interface loopback1
  ip address 40.201.51.1/32
  ip router ospf 2 area 0.0.0.201
  ip pim sparse-mode
```

```
feature msdp
ip msdp originator-id loopback0
ip msdp peer 40.201.0.19 connect-source loopback0

interface loopback0
  ip address 40.201.0.21/32
  ip router ospf 2 area 0.0.0.201
```

```
  ip pim sparse-mode
```

3.2.3.3.1 PIM-ASM Rendezvous Point

The DC2 topology relies heavily on vPC and as such PIM Sparse Mode has been configured as the protocol of choice for multicast routing. NX-OS does not support PIM SSM and PIM Bidir operating over vPC.

3.2.3.3.1.1 Auto-RP

The DC2 testbed is designed to have an RP for each POD in the data center to support the groups sourced from that particular POD. Each RP is configured on the aggregation switches for a given POD. DC2 makes use of Auto-RP to automate distribution of RP information in the network.

To Verify PIM RP:

```
DC201-6# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.207.51.1, uptime: 1d20h, expires: 00:02:21
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.2.50.1, (0), uptime: 1d20h, expires: 00:02:21 (A),
  priority: 0, RP-source: 40.207.51.1 (A), (local), group ranges:
     239.1.1.1/32   230.2.0.0/16
RP: 40.201.51.1*, (0), uptime: 1d20h, expires: 00:02:21,
  priority: 0, RP-source: 40.207.51.1 (A), group ranges:
      230.201.0.0/16

DC201-6# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Action    Mode      RP-address       Shared-tree-only range
232.0.0.0/8        Accept    SSM       -                -
230.2.0.0/16       -         ASM       40.2.50.1        -
230.201.0.0/16     -         ASM       40.201.51.1      -
```

3.2.3.3.1.1.1 Auto-RP Forward Listen

DC2 has been enabled with Auto-RP listening and forwarding feature so that the Auto-RP mechanism can dynamically inform routers in the PIM domain of the group-to-RP mapping since PIM dense mode is not supported on NX-OS. By default, listening or forwarding of Auto-RP messages is not enabled on NX-OS.

3.2.3.3.1.2 Static RP

DC2 network is configured with a backup RP on the core routers for all groups in the network. This RP is statically configured on all routers in the network. Auto-RP takes precedence over static RP.
To Verify PIM RP:

```
DC201-6# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.207.51.1, uptime: 1d20h, expires: 00:02:32
```

```
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.2.50.1, (0), uptime: 1d20h, expires: 00:02:32 (A),
  priority: 0, RP-source: 40.207.51.1 (A), (local), group ranges:
      239.1.1.1/32   230.2.0.0/16
RP: 40.201.51.1*, (0), uptime: 1d20h, expires: 00:02:32,
  priority: 0, RP-source: 40.207.51.1 (A), group ranges:
      230.201.0.0/16

DC201-6# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.2.0.0/16       ASM       40.2.50.1        -
230.201.0.0/16     ASM       40.201.51.1      -
```

3.2.3.3.1.3 Anycast RP with MSDP

DC2 has been configured Anycast RP with MSDP within each POD at the aggregation layer and other switches.

Anycast RP and MSDP Configuration:

N7K aggregation 1:	N7K aggregation 2:

!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list 230.201.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  ip address 40.201.51.1/32
  ip router ospf 2 area 0.0.0.201
  ip pim sparse-mode
 	
! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.201.0.21 connect-source loopback0
interface loopback0
  ip address 40.201.0.19/32
  ip router ospf 2 area 0.0.0.201
  ip pim sparse-mode
	!Anycast RP configuration
	ip pim send-rp-announce loopback1 group-list 230.201.0.0/16
	ip pim send-rp-discovery loopback1
	interface loopback1
	  ip address 40.201.51.1/32
	  ip router ospf 2 area 0.0.0.201
	  ip pim sparse-mode
	
	! MSDP configuration
	ip msdp originator-id loopback0
	ip msdp peer 40.201.0.19 connect-source loopback0
	interface loopback0
	  ip address 40.201.0.21/32
	  ip router ospf 2 area 0.0.0.201
	  ip pim sparse-mode

To Verify MSDP Peer and SA_Cache:

```
DC201-5# sh ip msdp sa-cache
MSDP SA Route Cache for VRF "default" - 100 entries
Source          Group            RP               ASN         Uptime
201.11.7.1      230.201.0.1      40.201.0.21      0           22:58:41
201.11.7.2      230.201.0.1      40.201.0.21      0           22:58:41
201.11.7.3      230.201.0.1      40.201.0.21      0           22:58:41
201.11.7.4      230.201.0.1      40.201.0.21      0           22:58:41
201.11.7.5      230.201.0.1      40.201.0.21      0           22:58:41
201.11.7.6      230.201.0.1      40.201.0.21      0           22:58:41

DC201-5# sh ip msdp sum
MSDP Peer Status Summary for VRF "default"
Local ASN: 0, originator-id: 40.201.0.19

Number of configured peers:  1
Number of established peers: 1
Number of shutdown peers:    0
```

```
Peer            Peer        Connection      Uptime/   Last msg  (S,G)s
Address         ASN         State           Downtime  Received  Received
40.201.0.21         0           Established     23:00:10  00:00:58  100
```

3.2.3.3.2 PIM SPT-Threshold

DC2 testbed has been enabled *ip pim spt-threshold infinity* on the last hop non-vPC PIM routers to decrease the multicast entries hardware usage across the network. Nexus 7000 vPC does not support PIM spt-threshold configuration.

3.2.3.3.3 Multicast Multipath

Cisco NX-OS Multicast Multipath is enabled by default and the load sharing selection algorithm is based on the source and group addresses. On Cisco IOS, Multicast Multipath is disabled by default. When multipath is enabled on Cisco IOS, the default load sharing selection algorithm is source-based. The algorithm on IOS can be configured to match the behavior on NX-OS with the command "*ip multicast multipath s-g-hash basic".*

DC2 testbed been has enabled with multicast multipath across the whole network on all applicable platforms.

3.2.3.4 Multicast Forwarding Verification

The following sequence of commands illustrates the verification of the Cisco NX-OS multicast L2 and L3 forwarding.

Displays a Specific Multicast Route 230.202.0.1 with Incoming Interface Information:

```
DC201-6# show ip mroute 230.202.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.202.0.1/32), uptime: 2d08h, mrib pim ip igmp
  Incoming interface: port-channel3, RPF nbr: 40.201.2.15
  Outgoing interface list: (count: 20)
    Vlan2010, uptime: 03:50:20, igmp
    Vlan2009, uptime: 03:50:20, igmp
    Vlan2008, uptime: 03:50:20, igmp
    Vlan2007, uptime: 03:50:20, igmp
    Vlan2006, uptime: 03:50:20, igmp
    Vlan2005, uptime: 03:50:20, igmp
    Vlan2004, uptime: 03:50:20, igmp
    Vlan2003, uptime: 03:50:20, igmp
    Vlan2002, uptime: 03:50:20, igmp
    Vlan2001, uptime: 03:50:20, igmp
    Vlan20, uptime: 03:50:21, igmp
    Vlan19, uptime: 03:50:21, igmp
    Vlan18, uptime: 03:50:21, igmp
    Vlan17, uptime: 03:50:21, igmp
    Vlan16, uptime: 03:50:21, igmp
    Vlan15, uptime: 03:50:21, igmp
    Vlan14, uptime: 03:50:21, igmp
    Vlan13, uptime: 03:50:21, igmp
    Vlan12, uptime: 03:50:21, igmp
    Vlan11, uptime: 03:50:21, igmp

(202.11.17.1/32, 230.202.0.1/32), uptime: 04:22:33, ip mrib pim
  Incoming interface: port-channel3, RPF nbr: 40.201.2.15
  Outgoing interface list: (count: 20)
    Vlan2010, uptime: 03:50:20, mrib
```

```
    Vlan2009, uptime: 03:50:20, mrib
    Vlan2008, uptime: 03:50:20, mrib
    Vlan2007, uptime: 03:50:20, mrib
    Vlan2006, uptime: 03:50:20, mrib
    Vlan2005, uptime: 03:50:20, mrib
    Vlan2004, uptime: 03:50:20, mrib
    Vlan2003, uptime: 03:50:20, mrib
    Vlan2002, uptime: 03:50:20, mrib
    Vlan2001, uptime: 03:50:20, mrib
    Vlan20, uptime: 03:50:21, mrib
    Vlan19, uptime: 03:50:21, mrib
    Vlan18, uptime: 03:50:21, mrib
    Vlan17, uptime: 03:50:21, mrib
    Vlan16, uptime: 03:50:21, mrib
    Vlan15, uptime: 03:50:21, mrib
    Vlan14, uptime: 03:50:21, mrib
    Vlan13, uptime: 03:50:21, mrib
    Vlan12, uptime: 03:50:21, mrib
    Vlan11, uptime: 03:50:21, mrib
```

Display DR Information for Interface Vlan11:

```
DC201-6# sh ip pim interface brief
PIM Interface Status for VRF "default"
Interface            IP Address      PIM DR Address  Neighbor  Border
                                                     Count     Interface
Vlan11               201.11.0.21     201.11.0.21     1         no
port-channel3        40.201.2.21     40.201.2.21     1         no
port-channel4        40.201.4.21     40.201.4.21     1         no
port-channel10       40.201.10.21    40.201.10.21    0         no
loopback0            40.201.0.21     40.201.0.21     0         no
loopback1            40.201.51.1     40.201.51.1     0         no
```

Displays Mroute RPF Interface and Forwarding Counters in L3 Hardware Table:

```
DC201-6# sh forwarding multicast route group 230.202.0.1 source 202.11.17.1

slot  2
=======


  (202.11.17.1/32, 230.202.0.1/32), RPF Interface: port-channel4, flags:
    Received Packets: 859951 Bytes: 82555296
    Number of Outgoing Interfaces: 20
    Outgoing Interface List Index: 4
      Vlan11 Outgoing Packets:3469406846 Bytes:333063055872
      Vlan12 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan13 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan14 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan15 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan16 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan17 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan18 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan19 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan20 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2001 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2002 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2003 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2004 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2005 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2006 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2007 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2008 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2009 Outgoing Packets:3965275677 Bytes:380666463648
      Vlan2010 Outgoing Packets:3965275677 Bytes:380666463648
```

Displays the Multicast Routing Table with Packet Counts and Bit Rates for All Sources:

```
DC201-6# sh ip mroute 230.202.0.1 summary
IP Multicast Routing Table for VRF "default"

Total number of routes: 810
Total number of (*,G) routes: 9
Total number of (S,G) routes: 800
Total number of (*,G-prefix) routes: 1
Group count: 9, rough average sources per group: 88.8

Group: 230.202.0.1/32, Source count: 400
Source          packets     bytes           aps   pps       bit-rate     oifs
(*,G)           23525       1924338         81    0         0.000   bps  20
202.11.17.1     873974      71665868        82    7         4.964   kbps 20
202.11.17.2     874156      71680792        82    10        6.647   kbps 20
202.11.17.3     873668      71640776        82    7         4.964   kbps 20
202.11.17.4     874156      71680788        81    10        6.647   kbps 20
202.11.17.5     873668      71640776        82    7         4.964   kbps 20
202.11.17.6     874154      71680616        81    10        6.647   kbps 20
```

Display IGMP Snooping Groups Information:

```
DC201-6# sh ip igmp snooping groups 230.202.0.1 vlan 11
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
11    230.202.0.1        v2   D     Po7 Po8
```

Displays Detected Multicast Routers for VLAN:

```
DC201-6# sh ip igmp snooping mrouter vlan 11
Type: S - Static, D - Dynamic, V - vPC Peer Link
      I - Internal, F - FabricPath core port
      C - Co-learned, U - User Configured
      P - learnt by Peer
Vlan  Router-port   Type      Uptime      Expires
11    Po5           SVD       1d23h       00:04:34
11    Vlan11        I         1d23h       never
```

Displays IGMP Snooping Querier Information for VLAN:

```
DC201-6# sh ip igmp snooping querier vlan 11
Vlan  IP Address       Version    Expires      Port
11    201.11.0.19      v2         00:02:25     port-channel5
```

Display L2 MFDM Software Entries for Group/VLAN 11:

```
DC201-6# sh forwarding distribution ip igmp snooping vlan 11 group 230.202.0.1
Vlan: 11, Group: 230.202.0.1, Source: 0.0.0.0
  Outgoing Interface List Index: 3
  Reference Count: 320
  Platform Index: 0x7bce
  Number of Outgoing Interfaces: 3
    port-channel5
    port-channel7
    port-channel8

Vlan: 11, Aggregated Group: 230.202.0.1, Source: 0.0.0.0
  Outgoing Interface List Index: 3
  Reference Count: 320
  Platform Index: 0x7bce
  Number of Outgoing Interfaces: 3
```

```
    port-channel5
    port-channel7
    port-channel8
```

Display L2 Hardware Entry for Group/VLAN:

```
DC201-6# sh system internal ip igmp snooping vlan 11 group 230.202.0.1 module 2

Lookup Mode : IP

Vlan   Group           Source          Epoch   RID   DTL      sw-index
11    230.202.0.1                      0       3     0x7bce   0x2a       L

DC201-6# sh system internal ip igmp snooping vlan 11 group 230.202.0.1 module 3

Lookup Mode : IP

Vlan   Group           Source          Epoch   RID   DTL      sw-index
11    230.202.0.1                      0       3     0x7bce   0x2a       L
```

Display DTL Sent to LTL Index for PO7:

```
DC201-6# sh system internal pixm info ltl 0x7bce
MCAST LTLs allocated for VDC:2
============================================
LTL    IFIDX/RID   LTL_FLAG CB_FLAG
0x7bce 0x00000003 0x00     0x0002

mi | v5_f3_fpoe | v4_fpoe | v5_fpoe | clp_v4_l2 | clp_v5_l2 | clp20_v4_l3 | clp_cr_v4_l3 | flag |
proxy_if_index
0x10 | 0x8 | 0x0 | 0x88 | 0x0 | 0x88 | 0x48 | 0x48 | 0x0 | none

Member info
------------------
IFIDX           LTL
---------------------------------
Po8                0x040c
Po7                0x040a
Po5                0x0408
```

3.2.4 Layer-2/ Layer-3 Aggregation/Access Layer Network Design Overview

3.2.4.1 vPC

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco NX-OS switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

Figure 21 Creating a Single Logical Node through vPC (virtual PortChannel) Technology

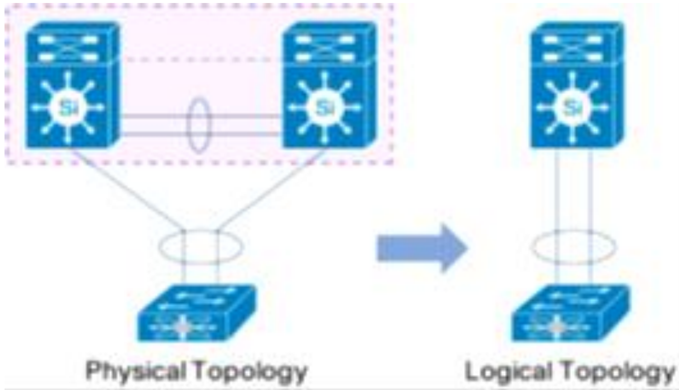

VPC peers configuration

```
N7K 1:
feature vpc

! vpc domain config
vpc domain 201
  peer-switch
  role priority 200
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  peer-gateway exclude-vlan 11
  track 10
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel6
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,10-20,2001-
2010,3000-3010
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet2/3
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,10-20,2001-
2010,3000-3010
  channel-group 6 mode active
  no shutdown

! vpc peer-keepalive config
interface Ethernet1/3
  vrf member vpc-keepalive
  ip address 1.1.1.1/24
  no shutdown

! vpc member port-channel config
interface port-channel7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  vpc 7

! vpc member port config
interface Ethernet2/7
```

```
N7K 2:
feature vpc

! vpc domain config
vpc domain 201
  peer-switch
  role priority 110
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  peer-gateway exclude-vlan 11
  track 10
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel5
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,10-20,2001-
2010,3000-3010
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,10-20,2001-
2010,3000-3010
  channel-group 5 mode active
  no shutdown

! vpc peer-keepalive config
interface Ethernet1/1
  vrf member vpc-keepalive
  ip address 1.1.1.2/24
  no shutdown

! vpc member port-channel config
interface port-channel7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  vpc 7

! vpc member port config
interface Ethernet8/1
```

<table>
<tr>
<td>

```
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  channel-group 7 mode active
  no shutdown

!vpc object tracking
!! uplinks
track 1 interface port-channel3 line-protocol
track 2 interface port-channel4 line-protocol
!!vpc peer-link
track 3 interface port-channel6 line-protocol
track 10 list boolean or
  object 1
  object 2
  object 3

! PIM prebuild SPT(only for non F2 mode)
ip pim pre-build-spt
```

</td>
<td>

```
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  channel-group 7 mode active
  no shutdown

! vpc object tracking
!! uplinks
track 1 interface port-channel3 line-protocol
track 2 interface port-channel4 line-protocol
!!vpc peer-link
track 3 interface port-channel5 line-protocol
track 10 list boolean or
  object 1
  object 2
  object 3

! PIM prebuild SPT(only for non F2 mode)
ip pim pre-build-spt
```

</td>
</tr>
</table>

Display vPC Status:

```
N7K-2# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 95
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 108
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po5    up     1-100,2001-2010,3001-3010,3951-3960

vPC status
----------------------------------------------------------------------
id   Port      Status Consistency Reason                 Active vlans
--   ----      ------ ----------- ------                 ------------
7    Po7       up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
8    Po8       up     success     success                   1,11-20,200
                                                            1-2010,3001
```

3.2.4.1.1 LACP

DC2 makes use of LACP mode active for all link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
N7K-2# show port-channel summary
```

```
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
3     Po3(RU)     Eth      LACP      Eth1/3(P)    Eth1/5(P)
4     Po4(RU)     Eth      LACP      Eth1/2(P)    Eth1/6(P)
5     Po5(SU)     Eth      LACP      Eth1/4(P)    Eth1/7(P)
7     Po7(SU)     Eth      LACP      Eth8/1(P)
8     Po8(SU)     Eth      LACP      Eth8/2(P)

DC201-6# show lacp interface Eth2/7
Interface Ethernet2/7 is up
  Channel group is 7 port channel is Po7
  PDUs sent: 432
  PDUs rcvd: 302
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(7f9b, 0-23-4-ee-be-c9, 8007, 8000, 207), (8000, 0-18-74-1e-e1-80, 6, 8000, 406)] ]
Operational as aggregated link since Mon Feb 17 12:07:06 2014

Local Port: Eth2/7   MAC Address= f8-66-f2-7-25-42
  System Identifier=0x8000,  Port Identifier=0x8000,0x207
  Operational key=32775
  LACP_Activity=passive
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=60
Actor Oper State=60
Neighbor: 0x406
  MAC Address= 0-18-74-1e-e1-80
  System Identifier=0x8000,  Port Identifier=0x8000,0x406
  Operational key=6
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=61
Partner Oper State=61
Aggregate or Individual(True=1)= 1
```

3.2.4.1.2 VLAN Trunking

DC2 testbed makes use of VLAN trunking in the aggregation-access blocks to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Display vlan information for Nexus 7000/7700:

```
N7K-2# show vlan internal usage
```

```
VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035,4048-4059     Online Diagnostic
4036-4039,4060-4087     ERSPAN
4042                    Satellite
4040                    Fabric scale
3968-4095               Current

N7K-2# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po5, Po7, Po8, Po17, Po27, Po71
                                                Po72, Po73, Po74, Po77, Po78
                                                Po201, Po221, Po401, Po421
                                                Po441, Po501, Po521, Eth1/4
                                                Eth1/7, Eth8/1, Eth8/2, Eth8/16
                                                Eth8/18, Eth8/29, Eth8/30
                                                Eth9/42, Eth10/31, Eth102/1/1
                                                Eth102/1/21, Eth102/1/41
                                                Eth104/1/25, Eth104/1/26
                                                Eth104/1/27, Eth104/1/28
                                                Eth104/1/29, Eth104/1/30
                                                Eth104/1/31, Eth104/1/32

VLAN Type         Vlan-mode
---- -----        ----------
11   enet         CE

Remote SPAN VLAN
----------------
Disabled

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------
```

Display vlan information for Nexus 5000/6000:

```
DC202-701# sh vlan internal usage

VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035               Online Diagnostic
4036-4039               ERSPAN
4042                    Satellite
3968-4047,4094          Current
```

Display vlan information for Nexus 3548:

```
DC204-47# sh vlan internal usage

VLAN         DESCRIPTION
---------    -------------------------------------------------------
3968-4031    Multicast
4032         Online diagnostics vlan1
4033         Online diagnostics vlan2
4034         Online diagnostics vlan3
4035         Online diagnostics vlan4
4036-4047    Reserved
4094         Reserved
```

3.2.4.1.3 Spanning Tree

vPC technology helps build a loop free topology by leveraging port-channels from access devices to the vPC domain. A port-channel is seen as a logical link from the spanning tree's standpoint, so a vPC domain with vPC-attached access devices forms a star topology at Layer 2 (there are no STP blocked ports in this type of topology). In this case, STP is used as a fail-safe mechanism to protect against any network loops.

DC2 makes use of Rapid-PVST which is the default spanning tree protocol for DC201. In DC202, MST is configured.

Display Spanning Tree Information:

```
N7K-DC201# sh spanning-tree vlan 11

VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    24587
             Address     0023.04ee.bec9
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24587  (priority 24576 sys-id-ext 11)
             Address     0023.04ee.bec9
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po6              Desg FWD 1         128.4101 (vPC peer-link) Network P2p
Po7              Desg FWD 1         128.4102 (vPC) P2p
Po8              Desg FWD 1         128.4103 (vPC) P2p
Po17             Desg FWD 1         128.4112 (vPC) P2p
Po27             Desg FWD 1         128.4122 (vPC) P2p
Po2011           Desg FWD 1         128.6106 (vPC) Edge P2p
Po2012           Desg FWD 1         128.6107 (vPC) Edge P2p

N7K-DC201# sh spanning-tree summary totals
Switch is in rapid-pvst mode
Root bridge for: VLAN0001, VLAN0010-VLAN0020, VLAN2001-VLAN2010
  VLAN3000-VLAN3010, VLAN3951-VLAN3960
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short
vPC peer switch                          is enabled (operational)
STP-Lite                                 is enabled

Name                     Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
43 vlans                     0         0        0        257        257

N7K-DC202#sh  spanning-tree  vlan  11
MST0000
  Spanning tree enabled protocol mstp
  Root ID    Priority    0
             Address     c84c.75fa.6000
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    0      (priority 0 sys-id-ext 0)
             Address     c84c.75fa.6000
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
```

```
---------------- ---- --- --------- -------- --------------------------------
Po17             Desg FWD  200       128.4112 (vPC) P2p
Po18             Desg FWD  200       128.4113 (vPC) P2p

N7K-DC202# sh spanning-tree mst

##### MST0    vlans mapped:   1-4094
Bridge        address c84c.75fa.6000  priority      0     (0 sysid 0)
Root          this switch for the CIST
Regional Root this switch
Operational   hello time 2 , forward delay 15, max age 20, txholdcount 6
Configured    hello time 2 , forward delay 15, max age 20, max hops    20

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po17             Desg FWD  200       128.4112 (vPC) P2p
Po18             Desg FWD  200       128.4113 (vPC) P2p
```

3.2.4.1.3.1 vPC Peer Switch Feature

The vPC Peer Switch feature allows a pair of vPC peer devices to appear as a single Spanning Tree Protocol root in the Layer 2 topology (they have the same bridge ID). vPC peer switch must be configured on both vPC peer devices to become operational.
This feature simplifies Spanning Tree Protocol configuration by configuring vPC VLANs on both peer devices with the same Spanning Tree Protocol priority. A vPC Peer Switch eliminates the need to map the Spanning Tree Protocol root to the vPC primary peer device.

Figure 22 vPC Peer-switch

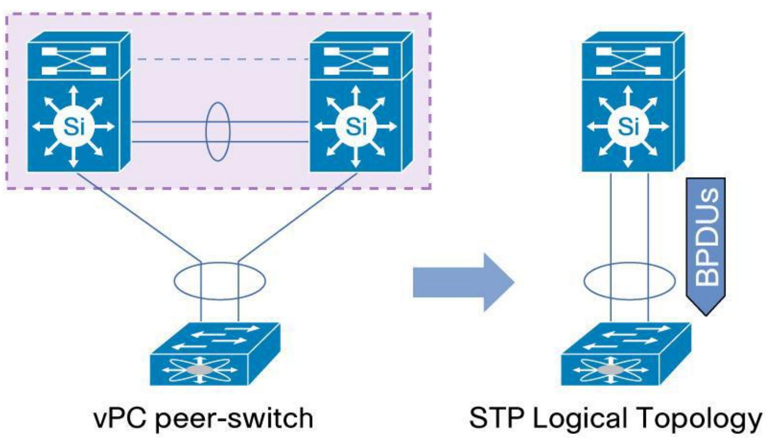

3.2.4.1.4 Configuration Parameters Consistency

After the vPC feature is enabled and the vPC peer-link on both peer devices is configured, Cisco Fabric Services messages provide a copy of the local vPC peer device configuration to the remote vPC peer device. The systems then determine whether any of the crucial configuration parameters differ on the two devices.
When a Type 1 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 1 consistency check failure, all vPC member ports are set to down state.

- For a vPC interface configuration Type 1 consistency check failure, the misconfigured vPC is set to down state

When a Type 2 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 2 consistency check failure, all vPC member ports remain in up state and vPC systems trigger protective actions.
- For a vPC interface configuration Type 2 consistency check failure, the misconfigured vPC remains in up state. However, depending on the discrepancy type, vPC systems will trigger protective actions. The most typical misconfiguration deals with the allowed VLANs in the vPC interface trunking configuration. In this case, vPC systems will disable the vPC interface VLANs that do not match on both sides.

Display vPC Consistency Parameters:

```
DC201-5# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
Interface-vlan admin up     2     1,10-20,2001-2010,3001 1,10-20,2001-2010,3001
                                  -3010                  -3010
Interface-vlan routing      2     1,10-20,2001-2010,3001 1,10-20,2001-2010,3001
capability                        -3010                  -3010
Allowed VLANs               -     1,10-20,2001-2010,3000 1,10-20,2001-2010,3000
                                  -3010                  -3010
Local error VLANs           -     -                      -

DC201-5# show vpc consistency-parameters interface port-channel 7

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Port Type               1     Default                Default
STP Port Guard              1     Default                Default
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-c9, 8007, 0-23-4-ee-be-c9, 8007,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-18-74-1e-e1-80, 6,   0-18-74-1e-e1-80, 6,
                                  0, 0)]                 0, 0)]
mode                        1     passive                passive
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
LACP Mode                   1     on                     on
Interface type              1     port-channel           port-channel
```

```
Admin port mode             1      trunk                  trunk
vPC card type               1      Clipper                Clipper
Allowed VLANs               -      1,11-20,2001-2010,3001 1,11-20,2001-2010,3001
                                   -3010                  -3010
Local error VLANs           -      -                      -
```

3.2.4.1.5 vPC Role Priority

There are two defined vPC roles: primary and secondary. The vPC role defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDUs) and responds to Address Resolution Protocol (ARP).

In case of a tie (same role priority value defined on both peer devices), the lowest system MAC will dictate the primary peer device.

Display vPC Role, System-MAC, System-Priority:

```
N7K-2# show vpc role

vPC Role status
----------------------------------------------------
vPC role                         : primary
Dual Active Detection Status     : 0
vPC system-mac                   : 00:23:04:ee:be:5f
vPC system-priority              : 32667
vPC local system-mac             : 00:23:ac:64:bb:c2
vPC local role-priority          : 110
```

3.2.4.1.6 vPC Peer-Link

The vPC peer-link is a standard 802.1Q trunk that performs the following actions:

- Carry vPC and non-vPC VLANs.
- Carry Cisco Fabric Services (CFS) messages that are tagged with CoS=4 for reliable communication CoS=4 for reliable communication.
- Carry flooded traffic between the vPC peer devices.
- Carry STP BPDUs, HSRP hello messages, and IGMP updates.

When the vPC peer-link fails and the vPC peer-keepalive link is still up, the vPC secondary peer device performs the following operations:

- Suspends its vPC member ports
- Shuts down the SVI associated to the vPC VLAN

Display vPC Peer-link Information:

```
DC201-5# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 201
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
```

```
Number of vPCs configured         : 32
Track object                      : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs       : 11
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po6    up     1,10-20,2001-2010,3000-3010

vPC status
----------------------------------------------------------------------
id   Port      Status Consistency Reason                Active vlans
--   ----      ------ ----------- ------                ------------
7    Po7       up     success     success                  1,11-20,200
                                                           1-2010,3001
                                                           -3010
8    Po8       up     success     success                  1,11-20,200
                                                           1-2010,3001
                                                           -3010
17   Po17      up     success     success                  1,11-20,200
                                                           1-2010,3001
                                                           -3010
27   Po27      up     success     success                  1,10-20,200
                                                           1-2010,3001
                                                           -3010
```

3.2.4.1.7 vPC Peer-Keepalive Link

The vPC peer-keepalive link is a Layer 3 link that joins one vPC peer device to the other vPC peer device and carries a periodic heartbeat between those devices. It is used at the boot up of the vPC systems to guarantee that both peer devices are up before forming the vPC domain. It is also used when the vPC peer-link fails, in which case, the vPC peer-keepalive link is leveraged to detect split brain scenario (both vPC peer devices are active-active).

Default Values for VPC Peer-Keepalive Links:

Timer	Default value
Keepalive interval	1 seconds
Keepalive hold timeout (on vPC peer-link loss)	3 seconds
Keepalive timeout	5 seconds

When building a vPC peer-keepalive link, use the following in descending order of preference:

1. Dedicated link(s) (1-Gigabit Ethernet port is enough) configured as L3. A port-channel with 2 X 1G port is preferred.
2. Mgmt0 interface (along with management traffic)
3. As a last resort, route the peer-keepalive link over the Layer 3 infrastructure.

DC2 makes use of the 1st option.

Display vPC Peer-Keepalive Information:

```
DC201-5# sh vpc peer-keepalive

vPC keep-alive status             : peer is alive
--Peer is alive for             : (31025) seconds, (523) msec
```

```
--Send status                   : Success
--Last send at                  : 2014.02.18 00:56:03 634 ms
--Sent on interface             : Eth1/3
--Receive status                : Success
--Last receive at               : 2014.02.18 00:56:03 359 ms
--Received on interface         : Eth1/3
--Last update from peer         : (0) seconds, (512) msec

vPC Keep-alive parameters
--Destination                   : 1.1.1.2
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc-keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

3.2.4.1.8 vPC Member Link

As suggested by the name, a vPC member port is a port-channel member of a vPC. A port-channel defined as a vPC member port always contains the keywords *vpc <vpc id>.*

A vPC only supports Layer 2 port-channels. The port-channel can be configured in access or trunk switchport mode. Any VLAN allowed on the vPC member port is by definition called a vPC VLAN. Whenever a vPC VLAN is defined on a vPC member port, it must also be defined on the vPC peer-link. Not defining a vPC VLAN on the vPC peer-link will cause the VLAN to be suspended.

The configuration of the vPC member port must match on both the vPC peer devices. If there is an inconsistency, a VLAN or the entire port channel may be suspended (depending on Type-1 or Type-2 consistency check for the vPC member port). For instance, a MTU mismatch will suspend the vPC member port.

Display vPC Member Port-channel Information:

```
DC201-5# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 201
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 32
Track object                      : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs       : 11
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po6    up     1,10-20,2001-2010,3000-3010

vPC status
```

```
-----------------------------------------------------------------------
id   Port      Status Consistency Reason                 Active vlans
--   ----      ------ ----------- ------                 ------------
7    Po7       up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
8    Po8       up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
17   Po17      up     success     success                   1,11-20,200
                                                            1-2010,3001
                                                            -3010
27   Po27      up     success     success                   1,10-20,200
                                                            1-2010,3001
                                                            -3010

DC201-5# show vpc consistency-parameters interface port-channel 7

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Port Type               1     Default                Default
STP Port Guard              1     Default                Default
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-c9, 8007, 0-23-4-ee-be-c9, 8007,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-18-74-1e-e1-80, 6,   0-18-74-1e-e1-80, 6,
                                  0, 0)]                 0, 0)]
mode                        1     passive                passive
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
LACP Mode                   1     on                     on
Interface type              1     port-channel           port-channel
Admin port mode             1     trunk                  trunk
vPC card type               1     Clipper                Clipper
Allowed VLANs               -     1,11-20,2001-2010,3001 1,11-20,2001-2010,3001
                                  -3010                  -3010
Local error VLANs           -     -                      -
```

3.2.4.1.9 vPC ARP Synchronization

The vPC ARP Synchronization feature improves the convergence time for Layer 3 flows (North to South traffic). When the vPC peer-link fails and subsequently recovers, vPC ARP Synchronization performs ARP bulk synchronization over Cisco Fabric Services (CFS) from the vPC primary peer device to the vPC secondary peer device.

Displays vPC ARP Synchronization Information:

```
DC201-5# show ip arp sync-entries

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Address         Age       MAC Address     Interface
201.210.7.1     00:00:06  00c9.d207.0100  Vlan2010
201.210.7.2     00:00:06  00c9.d207.0101  Vlan2010
201.210.7.3     00:00:06  00c9.d207.0102  Vlan2010
201.210.7.4     00:00:06  00c9.d207.0103  Vlan2010
201.210.7.5     00:00:06  00c9.d207.0104  Vlan2010
```

```
201.210.7.6      00:00:06  00c9.d207.0105  Vlan2010
```

3.2.4.1.10 vPC Delay Restore

After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from the access to the core until the Layer 3 connectivity is reestablished.

The vPC Delay Restore feature delays the vPCs leg bringup on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on the vPC leg. The result provides a graceful restoration along with zero packet loss during the recovery phase (traffic still gets diverted to the alive vPC peer device).

This feature is enabled by default with a vPC restoration default timer of 30 seconds, which is used on DC2.

3.2.4.1.11 vPC Object-Tracking

A vPC deployment where the L3 core uplinks and vPC peer-link interfaces are localized on the same module, is vulnerable to access layer isolation if that module fails on the primary vPC (vPC member ports are defined on both 1-Gbps line cards and on 10-Gbps line card).

Figure 23 vPC Object Tracking Feature – Behavior when vPC peer-link Fails – add M2 Module Patch.

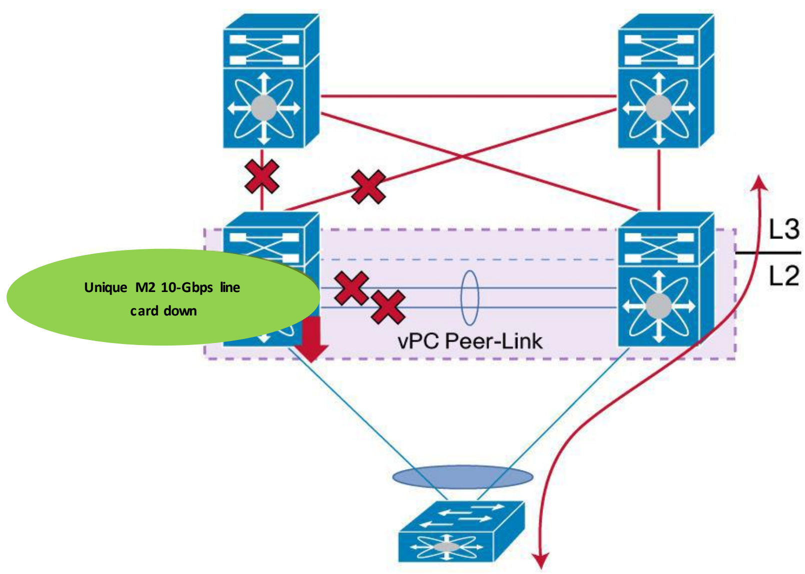

The vPC Object Tracking feature suspends the vPCs on the impaired device so that traffic can be diverted over the remaining vPC peer.

To use vPC object tracking, track both Peer-link interfaces and L3 core interfaces as a list of Boolean objects. Note that the Boolean AND operation is not supported with vPC object tracking. The vPC object tracking configuration must be applied on both vPC peer devices.

```
! Track the vpc peer link
track 1 interface port-channel5 line-protocol
! Track the uplinks to the core
track 2 interface port-channel3 line-protocol
track 3 interface port-channel4 line-protocol
! Combine all tracked objects into one.
! "OR" means if ALL objects are down, this object will go down
! ==> lost all connectivity to the L3 core and the peer link
track 10 list boolean OR
  object 1
  object 2
  object 3
! If object 10 goes down on the primary vPC peer,
! system will switch over to other vPC peer and disable all local vPCs
vpc domain 201
  track 10
```

Display Tracked Object Status:

```
DC201-5# show track 10
Track 10
  List  Boolean or
  Boolean or is UP
  2 changes, last change 17:21:21
  Track List Members:
    object 3 UP
    object 2 UP
    object 1 UP
  Tracked by:
    vPCM Domain 201
    ISCM iscm configuration
```

3.2.4.1.12 vPC Auto-Recovery

vPC auto-recovery feature was designed to address 2 enhancements to vPC.

- To provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure (vPC auto-recovery feature).
- To handle a specific case where both vPC peer devices reload but only one comes back to life (vPC auto-recovery reload-delay feature).

The switch which unsuspends its vPC role with vPC auto-recovery continues to remain primary even after peer-link is on. The other peer takes the role of secondary and suspends its own vPC until a consistency check is complete. Therefore, to avoid this situation from occurring erroneously, auto-recovery reload-delay-timer should be configured to be long enough for the system to fully complete its bootup sequence.

Helpful Commands for vPC Object Tracking:

Show vpc brief	Displays Auto-recovery status

Configuration Check:

```
DC201-5# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                   : 201
Peer status                     : peer adjacency formed ok
```

```
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 32
Track object                      : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs       : 11
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po6    up     1,10-20,2001-2010,3000-3010
```

3.2.4.1.13 HSRP Active/Active with vPC

HSRP in the context of vPC has been improved from a functional and implementation standpoint to take full benefits of the L2 dual-active peer devices nature offered by vPC technology. HSRP operates in active-active mode from a data plane standpoint, as opposed to classical active/standby implementation with a STP based network. No additional configuration is required. As soon as a vPC domain is configured and interface VLAN with an associated HSRP group is activated, HSRP will behave by default in active/active mode (on the data plane side).

From a control plane standpoint, active-standby mode still applies for HSRP in context of vPC; the active HSRP instance responds to ARP request. ARP response will contain the HSRP vMAC which is the same on both vPC peer devices. The standby HSRP vPC peer device just relays the ARP request to active HSRP/VRRP peer device through the vPC peer-link.

Sample Configuration:

```
! N7K-1:
interface Vlan11
  no ip redirects
  ip address 201.11.0.19/16
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 201.11.0.1
  no shutdown

! N7K-2:
interface Vlan11
  no ip redirects
  ip address 201.11.0.21/16
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    ip 201.11.0.1
  no shutdown
```

Helpful Commands for HSRP Active/Active with vPC:

Show hsrp brief	Displays hsrp status
Show mac address-table vlan <vlan id>	Displays mac addresses including HSRP vMAC; check for G-flag on vMAC for active/active HSRP

Configuration Check:

```
DC201-6# show hsrp brief
*:IPv6 group   #:group belongs to a bundle
                     P indicates configured to preempt.
                     |
 Interface   Grp  Prio P State    Active addr      Standby addr     Group addr
  Vlan1       1    100  P Active   local            201.0.1.19       201.0.1.1       (conf)
  Vlan10      1    100  P Active   local            201.10.0.19      201.10.0.1      (conf)
  Vlan11      1    200  P Active   local            201.11.0.19      201.11.0.1      (conf)
  Vlan11      2    210  P Active   local            201.11.0.19      201.111.0.1     (conf)

DC201-6# show mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G 11       0000.0c9f.f001    static       -       F    F  sup-eth1(R)
G 11       0000.0c9f.f002    static       -       F    F  sup-eth1(R)
```

3.2.4.1.14 PIM Pre-Build-SPT with vPC

PIM Pre-build SPT on non-forwarder attracts multicast traffic by triggering upstream PIM J/Ps (Join/Prune) without setting any interface in the OIF (Outgoing Interface) list. Multicast traffic is then always pulled to the non-active forwarder and finally dropped due to no OIFs.

In the vPC implementation in F2-mode, because of a hardware limitation, the PIM dual DR mode is disabled. In this case (with F2 mode), even if the **ip pim pre-build-spt** command is configured, there is no value added because the corresponding (S,G) route is not created in the first place.

3.2.4.1.15 Double-Sided vPC Topology

A double-sided vPC topology superposes two layers of vPC domain and the bundle between vPC domain 1 and vPC domain 2 is by itself a vPC. The vPC domain at the bottom is used for active/active connectivity from end-point devices to the network access layer. The vPC domain at the top is used for active/active FHRP in the L2/L3 boundary aggregation layer.

Figure 24 Double-Sided vPC Topology

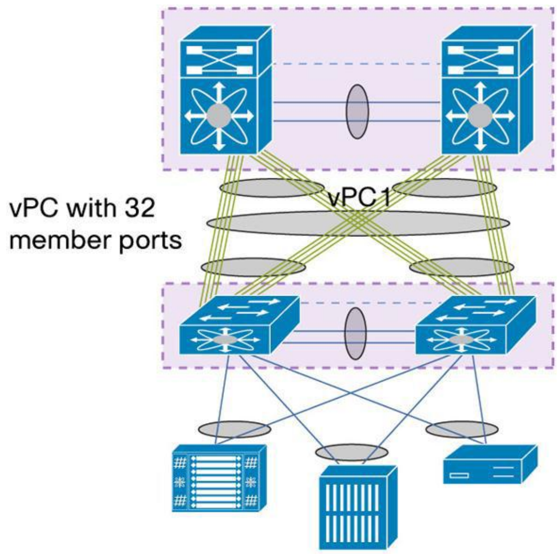

Benefits of double-sided vPC over single-sided vPC topology are listed below:

- Enables a larger Layer 2 domain.
- Provides a highly resilient architecture. In double-sided vPC, two access switches are connected to two aggregation switches whereas in single-sided vPC, one access switch is connected to two aggregation switches.
- Provides more bandwidth from the access to aggregation layer. Using a Cisco Nexus F2 Series modules line card for vPC and Cisco Nexus 5000 Series Switches with Release 4.1(3)N1(1a) or later, a vPC with 32 active member ports (that is, 320 Gbps) can be instantiated.

3.2.4.2 FabricPath

DC2 FabricPath topology is designed to have four FabricPath spines using Nexus 7000 at the aggregation layer. Two N7000 switches are used on DC2 which are configured to have two FabricPath VDC's each to simulate four spines. There are six Nexus 5000 leaf switches and two Nexus 6000 leaf switches on access layer that are connected to all four spines. We have 2 Nexus 7700 leaf connected to only two spines. The FabricPath feature is only supported on the F Series modules on the Nexus 7000.

Additional FabricPath links have been configured between two of the spines to simulate up to 14 FabricPath leaf nodes.

The FabricPath spines are deployed using sup2-e and F2e modules on Nexus 7000. Because of the multiple forwarding engines (FEs) on the F Series modules, the port pairs and port sets in the table below must be configured to be in the same VDC.

Nexus 7000 F Series Modules Port Pairs and Port Sets	
Port Pairs for F1 Modules	**Port Sets for F2 Modules**
Ports 1 and 2	Ports 1, 2, 3, 4

Ports 3 and 4	Ports 5, 6, 7, 8
Ports 5 and 6	Ports 9, 10, 11, 12
Ports 7 and 8	Ports 13, 14, 15, 16
Ports 9 and 10	Ports 17, 18, 19, 20
Ports 11 and 12	Ports 21, 22, 23, 24
Ports 13 and 14	Ports 25, 26, 27, 28
Ports 15 and 16	Ports 29, 30, 31, 32
Ports 17 and 18	Ports 33, 34, 35, 36
Ports 19 and 20	Ports 37, 38, 39, 40
Ports 21 and 22	Ports 41, 42, 43, 44
Ports 23 and 24	Ports 45, 46, 47, 48
Ports 25 and 26	
Ports 27 and 28	
Ports 29 and 30	
Ports 31 and 32	

DC2 FabricPath Configuration on Spines DC202-51/52 is as the Following:

```
feature-set FabricPath

logging level FabricPath isis 5

vlan 1,11-3010
  mode FabricPath
FabricPath switch-id 251
logging level FabricPath switch-id 5
vpc domain 211
  FabricPath switch-id 300
  FabricPath multicast load-balance

interface port-channel52
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel71
  switchport mode FabricPath

interface port-channel72
  switchport mode FabricPath

interface port-channel73
  switchport mode FabricPath

interface port-channel74
  switchport mode FabricPath

interface port-channel75
  switchport mode FabricPath

interface port-channel76
  switchport mode FabricPath

interface port-channel83
```

```
  switchport mode FabricPath

interface port-channel84
  switchport mode FabricPath

interface port-channel85
  switchport mode FabricPath

interface port-channel86
  switchport mode FabricPath

interface port-channel91
  switchport mode FabricPath

interface port-channel92
  switchport mode FabricPath

interface port-channel93
  switchport mode FabricPath

interface port-channel94
  switchport mode FabricPath

interface port-channel95
  switchport mode FabricPath

interface port-channel96
  switchport mode FabricPath

interface port-channel701
  switchport mode FabricPath

interface port-channel702
  switchport mode FabricPath

interface port-channel703
  switchport mode FabricPath

interface port-channel704
  switchport mode FabricPath

interface port-channel705
  switchport mode FabricPath

interface port-channel706
  switchport mode FabricPath

interface port-channel707
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel708
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel709
  switchport mode FabricPath

interface port-channel710
  switchport mode FabricPath

FabricPath domain default
  root-priority 109
FabricPath load-balance unicast include-vlan
FabricPath load-balance multicast rotate-amount 0x3 include-vlan
```

DC2 FabricPath Configuration on Spines DC202-53/54 is as the Following:

```
feature-set FabricPath

logging level FabricPath isis 5

vlan 1,11-3010
  mode FabricPath
FabricPath switch-id 253
logging level FabricPath switch-id 5
vpc domain 212
  FabricPath switch-id 400
  FabricPath multicast load-balance

interface port-channel54
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel701
  switchport mode FabricPath

interface port-channel702
  switchport mode FabricPath

interface port-channel703
  switchport mode FabricPath

interface port-channel704
  switchport mode FabricPath

interface port-channel705
  switchport mode FabricPath

interface port-channel706
  switchport mode FabricPath

interface port-channel707
  switchport mode FabricPath
  FabricPath isis metric 200

interface port-channel708
  switchport mode FabricPath
  FabricPath isis metric 200

FabricPath domain default
  root-priority 110
FabricPath load-balance unicast include-vlan
FabricPath load-balance multicast include-vlan
```

3.2.4.2.1 FabricPath Switch-IDs

Cisco FabricPath can assign switch IDs to all the devices in the network automatically; however, it is convenient to use a meaningful numbering scheme. During network troubleshooting, having a distinct numbering scheme allows for faster and easier switch role identification.
DC202 has been assigned with switch IDs using the following scheme in the FabricPath domain network:

- The devices in the spine layer have been assigned an ID related to spine VDC naming: 251 to 254
- The devices in the leaf layer have been assigned an ID related to leaf device naming: 701 to 710
- The virtual switch for the domain has ID's: 300 and 400

Figure 25 DC202 FabricPath POD Logical Topology

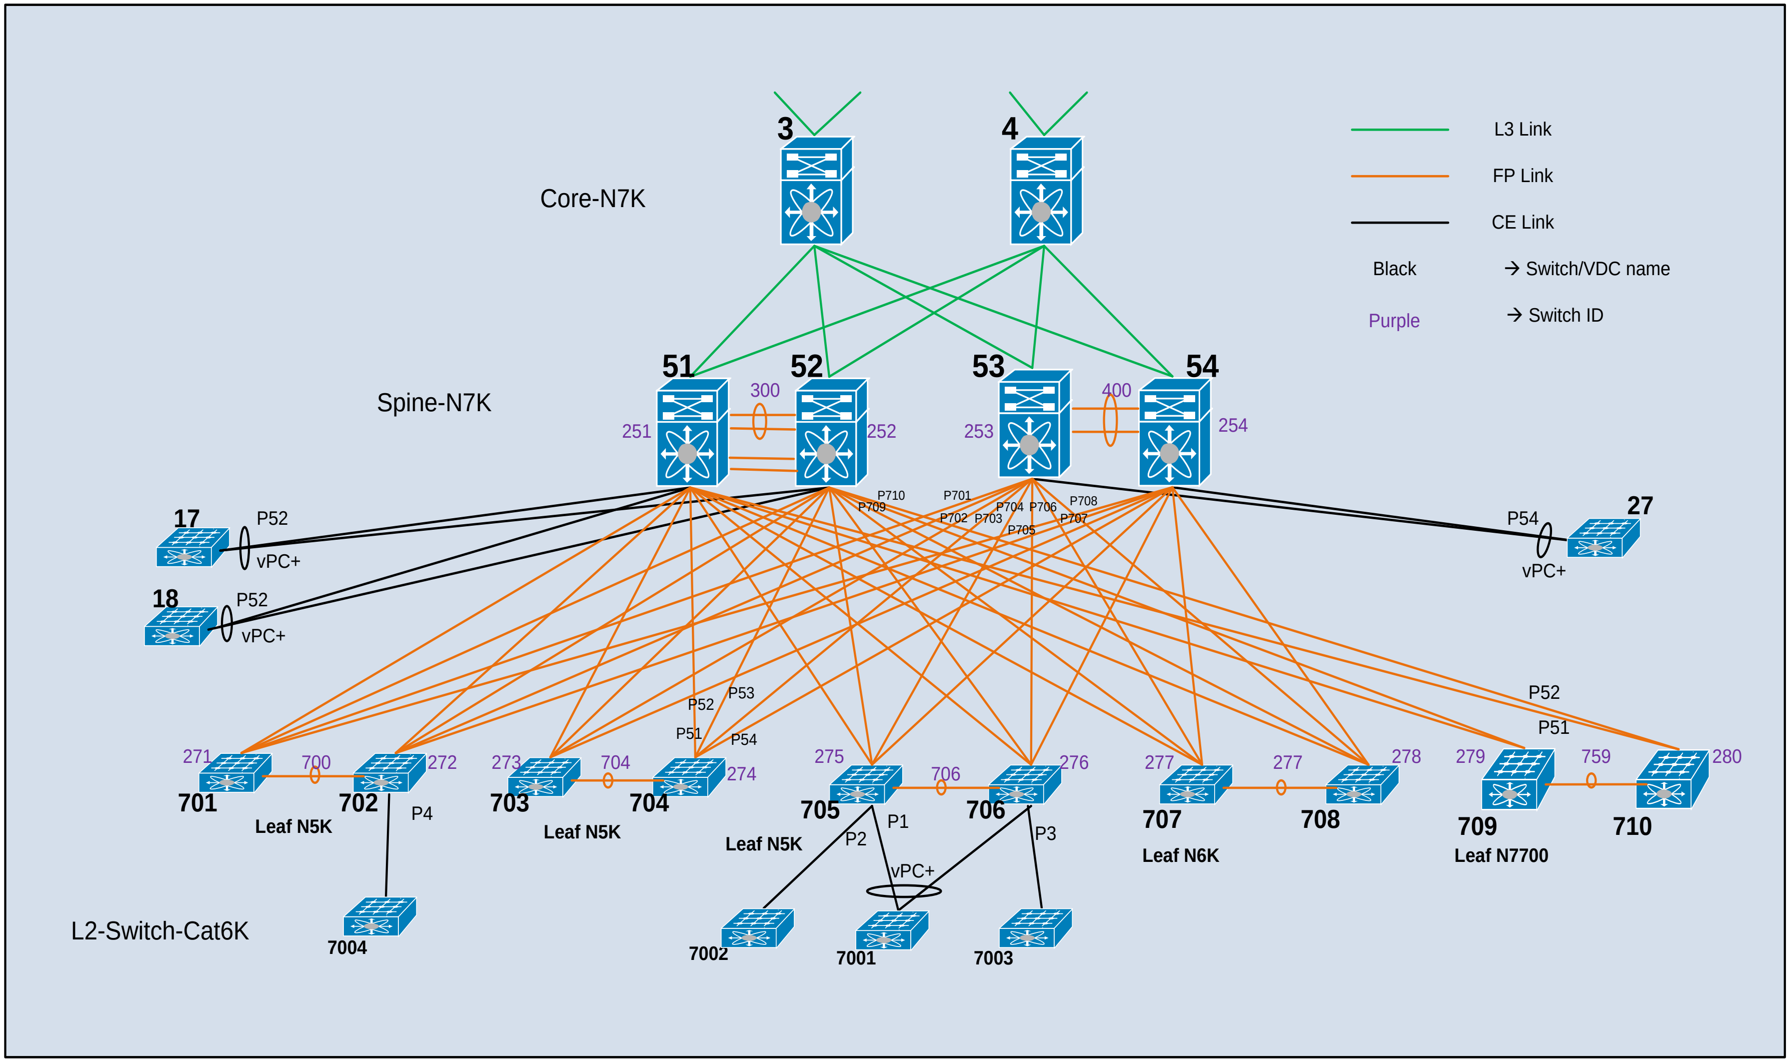

To Verify the FabricPath Switch ID:

```
DC202-51# sh FabricPath switch-id local
Switch-Id: 251
System-Id: 0026.980c.c0c3

DC202-51# sh FabricPath switch-id
                        FABRICPATH SWITCH-ID TABLE
Legend: '*' - this system
        '[E]' - local Emulated Switch-id
        '[A]' - local Anycast Switch-id
Total Switch-ids: 26
=============================================================================
    SWITCH-ID      SYSTEM-ID       FLAGS         STATE    STATIC  EMULATED/
                                                                  ANYCAST
--------------+----------------+------------+-----------+--------------------
*   251         0026.980c.c0c3    Primary     Confirmed Yes      No
    252         f866.f207.2543    Primary     Confirmed Yes      No
    253         0026.980c.c0c4    Primary     Confirmed Yes      No
    254         f866.f207.2544    Primary     Confirmed Yes      No
    271         547f.eef7.dafc    Primary     Confirmed Yes      No
    272         547f.eef7.d97c    Primary     Confirmed Yes      No
    273         547f.eef7.e3fc    Primary     Confirmed Yes      No
    274         547f.eebb.bd3c    Primary     Confirmed Yes      No
    275         547f.eef7.debc    Primary     Confirmed Yes      No
    276         547f.eede.927c    Primary     Confirmed Yes      No
    277         002a.6a3f.dac1    Primary     Confirmed No       Yes
    277         002a.6a3f.e301    Primary     Confirmed No       Yes
[E] 300         0026.980c.c0c3    Primary     Confirmed No       Yes
    300         f866.f207.2543    Primary     Confirmed No       Yes
    400         f866.f207.2544    Primary     Confirmed No       Yes
    400         0026.980c.c0c4    Primary     Confirmed No       Yes
    700         547f.eef7.d97c    Primary     Confirmed No       Yes
    700         547f.eef7.dafc    Primary     Confirmed No       Yes
    706         547f.eede.927c    Primary     Confirmed No       Yes
    706         547f.eef7.debc    Primary     Confirmed No       Yes
    707         002a.6a3f.e301    Primary     Confirmed Yes      No
    708         002a.6a3f.dac1    Primary     Confirmed Yes      No
    709         002a.6a5b.f7c2    Primary     Confirmed Yes      No
    710         547f.eede.c0c2    Primary     Confirmed Yes      No
    759         547f.eede.c0c2    Primary     Confirmed No       Yes
```

3.2.4.2.2 FabricPath VLANs

Cisco FabricPath VLANs should be consistently defined on all the Cisco FabricPath switches in a particular FabricPath topology. DC202 has 2000 VLANs configured on the spine layer switches.

To Verify the FabricPath VLANs:

```
DC202-54# sh FabricPath isis vlan-range
FabricPath IS-IS domain: default
MT-0
Vlans configured:
1, 11-3010, 4040
```

3.2.4.2.3 FabricPath Core Port

The configuration of a FabricPath core port is performed with the command *switchport mode FabricPath*. The FabricPath core port exchanges topology info through L2 ISIS adjacency and forwarding based on the Switch ID Table.

To Verify the FabricPath Interface:

```
DC202-51# sh FabricPath isis interface port-channel 701
FabricPath IS-IS domain: default
Interface: port-channel701
  Status: protocol-up/link-up/admin-up
  Index: 0x0003, Local Circuit ID: 0x01, Circuit Type: L1
  No authentication type/keychain configured
  Authentication check specified
  Extended Local Circuit ID: 0x160002BC, P2P Circuit ID: 0000.0000.0000.00
  Retx interval: 5, Retx throttle interval: 66 ms
  LSP interval: 33 ms, MTU: 1500
  P2P Adjs: 1, AdjsUp: 1, Priority 64
  Hello Interval: 10, Multi: 3, Next IIH: 00:00:04
  Level   Adjs   AdjsUp  Metric   CSNP  Next CSNP  Last LSP ID
  1          1        1      20     60  Inactive   ffff.ffff.ffff.ff-ff
  Topologies enabled:
    Level Topology Metric  MetricConfig Forwarding
    0     0        4000    no           UP
    1     0        20      no           UP

DC202-51# sh FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel52 P2P   4     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel71 P2P   8     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel72 P2P   9     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel73 P2P   10    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel74 P2P   11    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel75 P2P   12    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel76 P2P   13    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel83 P2P   16    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel84 P2P   19    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel85 P2P   20    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel91 P2P   21    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel92 P2P   24    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel93 P2P   22    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel94 P2P   23    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel95 P2P   25    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel701 P2P   3     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel702 P2P   6     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel703 P2P   5     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel704 P2P   7     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel705 P2P   14    Up/Ready   0x01/L1   1500 20      64          1/1
port-channel706 P2P   15    Up/Ready   0x01/L1   1500 20      64          1/1
port-channel707 P2P   2     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel708 P2P   1     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel709 P2P   17    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel710 P2P   18    Up/Ready   0x01/L1   1500 200     64          1/1



DC202-54# sh FabricPath isis interface port-channel 701
FabricPath IS-IS domain: default
Interface: port-channel701
  Status: protocol-up/link-up/admin-up
  Index: 0x0001, Local Circuit ID: 0x01, Circuit Type: L1
  No authentication type/keychain configured
  Authentication check specified
  Extended Local Circuit ID: 0x160002BC, P2P Circuit ID: 0000.0000.0000.00
  Retx interval: 5, Retx throttle interval: 66 ms
```

```
  LSP interval: 33 ms, MTU: 1500
  P2P Adjs: 1, AdjsUp: 1, Priority 64
  Hello Interval: 10, Multi: 3, Next IIH: 00:00:05
  Level   Adjs   AdjsUp  Metric   CSNP  Next CSNP  Last LSP ID
  1          1        1      20     60  Inactive   ffff.ffff.ffff.ff-ff
  Topologies enabled:
    Level Topology Metric  MetricConfig Forwarding
    0     0        4000    no           UP
    1     0        20      no           UP

DC202-54# sh FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel53 P2P    9     Up/Ready    0x01/L1   1500 200      64          1/1
port-channel701 P2P   1     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel702 P2P   3     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel703 P2P   2     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel704 P2P   4     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel705 P2P   7     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel706 P2P   8     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel707 P2P   5     Up/Ready    0x01/L1   1500 200      64          1/1
port-channel708 P2P   6     Up/Ready    0x01/L1   1500 200      64          1/1
```

3.2.4.2.4 FabricPath Metric

Cisco FabricPath ISIS calculates the preferred path to any switch-id based on the metric to any given destination. The metric is as follows:

- 1-Gbps Ethernet links have a cost of 400
- 10-Gigabit Ethernet links have a cost of 40
- 20-Gbps have a cost of 20

For FabricPath on DC2, NVT has set a higher ISIS metric on vPC peer links between the spine switches and on FabricPath links between the spines to prevent traffic from flowing through the vPC peer links.

To Verify the FabricPath ISIS Metric, use the Following Commands:

```
DC202-51# sh FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel52 P2P   4     Up/Ready    0x01/L1   1500 200      64          1/1
port-channel71 P2P   8     Up/Ready    0x01/L1   1500 200      64          1/1
port-channel72 P2P   9     Up/Ready    0x01/L1   1500 200      64          1/1
port-channel73 P2P   10    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel74 P2P   11    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel75 P2P   12    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel76 P2P   13    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel83 P2P   16    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel84 P2P   19    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel85 P2P   20    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel91 P2P   21    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel92 P2P   24    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel93 P2P   22    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel94 P2P   23    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel95 P2P   25    Up/Ready    0x01/L1   1500 200      64          1/1
port-channel701 P2P   3     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel702 P2P   6     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel703 P2P   5     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel704 P2P   7     Up/Ready    0x01/L1   1500 20       64          1/1
port-channel705 P2P   14    Up/Ready    0x01/L1   1500 20       64          1/1
port-channel706 P2P   15    Up/Ready    0x01/L1   1500 20       64          1/1
port-channel707 P2P   2     Up/Ready    0x01/L1   1500 200      64          1/1
```

```
port-channel708 P2P   1     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel709 P2P   17    Up/Ready   0x01/L1   1500 200     64          1/1
port-channel710 P2P   18    Up/Ready   0x01/L1   1500 200     64          1/1


DC202-54# sh FabricPath isis interface brief
FabricPath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel53 P2P   1     Up/Ready   0x01/L1   1500 200     64          1/1
port-channel701 P2P   2     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel702 P2P   3     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel703 P2P   4     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel704 P2P   5     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel705 P2P   6     Up/Ready   0x01/L1   1500 20      64          1/1
port-channel706 P2P   7     Up/Ready   0x01/L1   1500 20      64          1/1
```

3.2.4.2.5 Root for FabricPath Multi-Destination Trees

In FabricPath, multicast, broadcast and flooded traffic are forwarded along a multi-destination tree. FabricPath allows for multiple multi-destination trees in order to achieve traffic load balancing for multi-destination frames.
Two multi-destination trees are defined in Cisco FabricPath network by default, and multi-destination traffic is mapped to either of those trees for load-balancing purposes. The root of those multi-destination trees in the network should be explicitly set so as to provide an optimal topology.

Cisco FabricPath Intermediate Switch-to-Intermediate Switch (IS-IS) Protocol elects the switch with the highest configured root priority as the root for multi-destination tree 1. The switch with the second-highest root priority becomes the root for multi-destination tree 2. If there is no root priority configured, the other two parameters will be compared, system ID and switch ID, with higher values being better in all cases.

NVT has set the roots of the two multi-destination trees at two spine switches, one from each pair of vPC+ switches. If either of those switches fails, a replacement root would be elected out of all the FabricPath domain switches. This backup root should be configured in advance so that the system falls back to a predetermined topology in a failure scenario.

The Figure 26 shows the DC202 FabricPath Root design for the multi-destination trees. Spine 54 has highest root priority and is selected as root of FTag 1 and Spine 52 has second highest root priority and is selected as root of FTag 2.

Figure 26 FabricPath Root design for the multi-destination trees on DC202

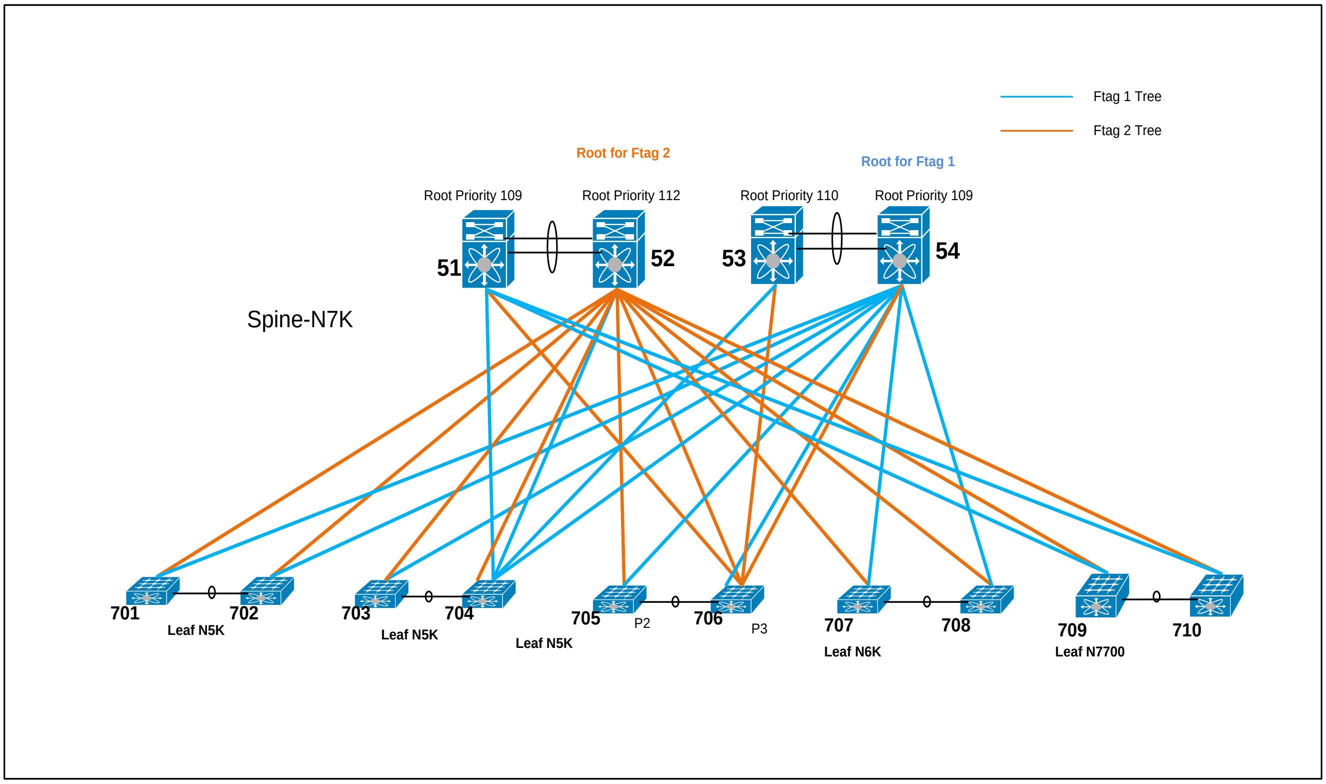

FTag trees are used as follows:

- FTag1 tree is used for unknown unicast, broadcast, and multicast.
- FTag2 tree is used only for multicast traffic.

To Verify FabricPath Multi-destination Tree Root:

```
DC202-54# sh FabricPath isis topology summary
FabricPath IS-IS Topology Summary
FabricPath IS-IS domain: default
MT-0
  Configured interfaces:  port-channel53  port-channel701  port-channel702
  port-channel703  port-channel704  port-channel705  port-channel706
 Max number of trees: 2  Number of trees supported: 2
    Tree id: 1, ftag: 1, root system: f866.f207.2544, 254
    Tree id: 2, ftag: 2 [transit-traffic-only], root system: f866.f207.2543, 252
Ftag Proxy Root: f866.f207.2544
```

To Verify which Multicast FTag Tree is Used in N7K:

```
DC202-54# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 202.11.27.1 dst-ip
230.202.0.1 vlan 12 module 2
128b Hash Key generated : 48 2e 00 00 00 00 00 00 32 82 c6 c0 79 b2 80 00
0x7f
```

```
    FTAG SELECTED IS : 2 (HASH 127)

DC202-53# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 202.11.27.1 dst-ip
230.202.0.1 vlan 11 module 2
128b Hash Key generated : 00 32 82 c6 c0 40 00 00 00 08 31 00 00 00 00 00
0x2
    FTAG SELECTED IS : 1 (HASH 2)
```

To Verify which Multicast FTag Tree is Used in N5K:

```
DC202-702# sh FabricPath load-balance multicast ftag-selected vlan 12 macg 0100.5e4d.0002

 If the traffic is received on a non-vPC port:
 Ftag selected : 2

 If the traffic is received on a vPC port:
 Ftag selected : 1

======================================

 Vlan : 12 (int_vlan : 3956)
 Macg : 0100.5e4d.0002

 Hash-key : 0x0f740000 00000000
 Hash-val : 237
 Num_trees : 2

======================================
```

3.2.4.2.6 vPC+ for FabricPath

NVT DC2 testbed is designed to have 2 pairs of vPC+ peers on the FabricPath spine and 3 pairs of vPC+ peers on the FabricPath leaf. The vPC+ peer-link must be configured as a FabricPath core link. FabricPath vPC+ configuration is as the following:

N7K aggregation VDC 5:

```
!vPC+ configuration
feature vpc
vpc domain 211
  peer-switch
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  delay restore 120
  dual-active exclude interface-vlan 1,11-20,2001-
2010
  track 10
  auto-recovery
  FabricPath switch-id 300
  FabricPath multicast load-balance
  ip arp synchronize


!vPC+ member configuration
interface port-channel17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  medium p2p
  vpc 17

!vPC+ peer link configuration
```

N7K aggregation VDC 6:

```
!vPC+ configuration
feature vpc
vpc domain 211
  peer-switch
  role priority 110
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  delay restore 120
  dual-active exclude interface-vlan 1,11-20,2001-
2010
  track 10
  auto-recovery
  FabricPath switch-id 300
  FabricPath multicast load-balance
  ip arp synchronize

!vPC+ member configuration
interface port-channel17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,11-20,2001-
2010,3001-3010
  medium p2p
  vpc 17

!vPC+ peer link configuration
```

interface port-channel52
  switchport
  switchport mode FabricPath
  spanning-tree port type network
  medium p2p
  vpc peer-link
  FabricPath isis metric 200
 	
!vPC+ peer keepalive configuration
interface Ethernet1/19
  vrf member vpc-keepalive
  ip address 1.1.1.1/24
  no shutdown
	interface port-channel51
	  switchport
	  switchport mode FabricPath
	  spanning-tree port type network
	  medium p2p
	  vpc peer-link
	  FabricPath isis metric 200
	
	!vPC+ peer keepalive configuration
	interface Ethernet1/19
	  vrf member vpc-keepalive
	  ip address 1.1.1.2/24
	  no shutdown

To Verify the vPC+:

```
DC202-51# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 211
vPC+ switch id                    : 300
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
vPC FabricPath status             : peer is reachable through FabricPath
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : 1,11-20,2001-2010
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)
FabricPath load balancing         : Enabled
Port Channel Limit                : limit to 244

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po52   up     1,11-20,2001-2010,3001-3010

vPC status
-------------------------------------------------------------------------------
id   Port      Status Consistency Reason           Active vlans  vPC+ Attribute
--   ----      ------ ----------- ------           ------------  --------------
17   Po17      up     success     success          1,11-20,2001- DF: Partial,
                                                   2010,3001-301  FP MAC:
                                                   1              300.11.65535
```

3.2.4.2.6.1 HSRP Active/Active with vPC+

Figure 27 HSRP Active/Active with vPC+

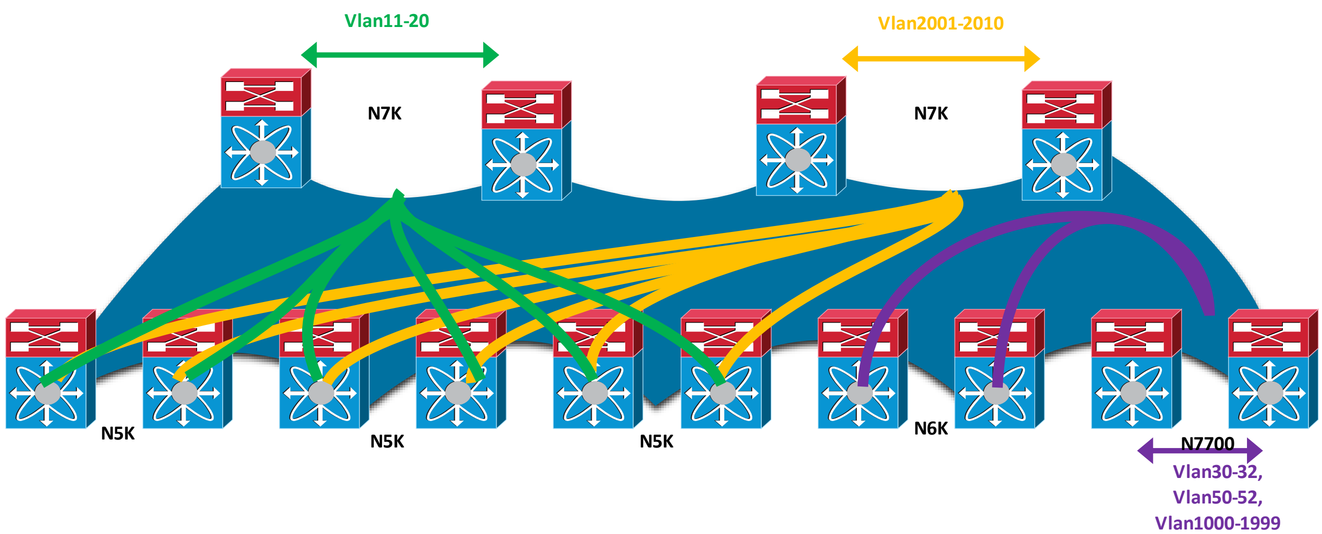

DC202 has split HSRP for VLANs among four spines with ten VLANs running HSRP between the first pair of spines only and another ten VLANs running HSRP between the other pair of spines only. In addition, the N7700 leaf pair is configured with HSRP for another 1000 VLANs.

DC202 HSRP configuration is as the below. Two HSRP groups with authentication and priority are configured for each VLAN.

```
interface Vlan11
  no shutdown
  no ip redirects
  ip address 202.11.0.51/16
  ip address 202.111.0.51/16 secondary
  ipv6 address 2001:1:202:11::51/64
  ip router ospf 2 area 0.0.0.202
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 202.11.0.1
  hsrp 2
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 200
    ip 202.111.0.1
```

To Verify HSRP peers and virtual mac address on N7000 spine/N7700 leaf:

```
DC202-51# sh hsrp interface vlan 11
Vlan11 - Group 1 (HSRP-V2) (IPv4)
  Local state is Active, priority 200 (Cfged 200), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 200
  Preemption Delay (Seconds) Minimum:120
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 2.047000 sec(s)
  Virtual IP address is 202.11.0.1 (Cfged)
  Active router is local
  Standby router is unknown
  Authentication MD5, key-string "cisco"
  Virtual mac address is 0000.0c9f.f001 (Default MAC)
  2 state changes, last state change 01:25:08
  IP redundancy name is hsrp-Vlan11-1 (default)

Vlan11 - Group 2 (HSRP-V2) (IPv4)
  Local state is Active, priority 200 (Cfged 200), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 200
  Preemption Delay (Seconds) Minimum:120
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 1.885000 sec(s)
  Virtual IP address is 202.111.0.1 (Cfged)
  Active router is local
  Standby router is unknown
  Authentication MD5, key-string "cisco"
  Virtual mac address is 0000.0c9f.f002 (Default MAC)
  2 state changes, last state change 01:25:08
  IP redundancy name is hsrp-Vlan11-2 (default)

DC202-51# sh hsrp brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr      Standby addr     Group addr
Vlan11      1   200  P Active   local            202.11.0.52      202.11.0.1      (conf)
Vlan11      2   200  P Active   local            202.11.0.52      202.111.0.1     (conf)

DC202-51# sh mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G 11       0000.0c9f.f001    static       -       F    F  sup-eth1(R)
G 11       0000.0c9f.f002    static       -       F    F  sup-eth1(R)
```

To Verify HSRP virtual mac learnt on Nexus 5000, N6000 and N7700 edge switches mac table:

```
N5k-705# sh mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0000.0c9f.f001    dynamic   0          F    F  300.0.65535
* 11       0000.0c9f.f002    dynamic   0          F    F  300.0.65535
```

3.2.4.2.6.2 vPC+ Dual-Active Exclude

As a result of declaring the link that connects the spines as a vPC peer-link, the default behavior of vPC applies, whereby, if the peer-link goes down, the SVIs on the vPC secondary device are shut down.
In the context of FabricPath designs, this behavior is not beneficial, because the FabricPath links are still available, and there is no good reason to shut down the SVIs on the secondary. It is thus recommended to configure *dual-active exclude* for all the vPC+ vlans.

To Verify Dual-Active Exclude VLAN:

```
DC202-51# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 211
vPC+ switch id                    : 300
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
vPC FabricPath status             : peer is reachable through FabricPath
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Track object                      : 10
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : 1,11-20,2001-2010
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)
FabricPath load balancing         : Enabled
Port Channel Limit                : limit to 244

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po52   up     1,11-3010

vPC status
-------------------------------------------------------------------------------
id   Port      Status Consistency Reason           Active vlans  vPC+ Attribute
--   ----      ------ ----------- ------           ------------  --------------
17   Po17      up     success     success          1,11-20,2001- DF: Partial,
                                                   2010,3001-301  FP MAC:
                                                   0              300.11.65535
```

3.2.4.2.7 Routed Multicast in FabricPath vPC+

PIM is enabled on the four Nexus 7000 spine VDCs with FabricPath VLANs configured under SVIs. It follows the same rules as all other non-FabricPath PODs. DC202 has defined all four spines as an auto-RP with Anycast RP/MSDP configured. From an operational perspective, it is advisable to align the PIM designated router (DR) priority with the HSRP primary.

3.2.4.3 FabricPath Load-Balancing and Verification

In Nexus 7000 F2 VDC, to modify default unicast/multicast load-balancing, the port-channel load-balancing has to be changed first, followed by the unicast/multicast load-balancing change.

3.2.4.3.2 FabricPath Unicast Load-Balancing and Verification

Cisco NX-OS FabricPath unicast Layer 2 ISIS ECMP is on by default.
The default FabricPath unicast load balancing mechanism on the Nexus 7000 with F2 line cards uses Layer 3/Layer 4 source and destination addresses without the VLAN included, Nexus 5000 uses Layer 2/Layer 3/Layer 4 source and destination addresses and VLAN with symmetric hashing. Nexus 6000 the default load-balancing scheme for ECMP is a mixed mode (Layer 2, Layer 3 and Layer 4 ports), it uses source and destination addresses with VLAN. To avoid hash polarization, each Cisco FabricPath switch automatically rotates the hash string by a number of bytes based on the system MAC address.

DC2 testbed has changed Nexus 7000 spine FabricPath unicast load-balancing mechanism using the following command and kept the Nexus 5000, Nexus 6000 and Nexus 7700 FabricPath unicast load-balance as default.

```
F2 VDC
! Change port-channel load-balance on main VDC

DC5-sup2(config)# FabricPath load-balance source-destination
```

```
! Change FabricPath load-balance unicast on spine F2 VDC

DC202-51(config)# FabricPath load-balance unicast rotate-amount 0xb include-vlan
```

```
! Verify Nexus 7000 spine FabricPath load-balance after modify

DC202-51(config)# sh FabricPath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 11 bytes
Use VLAN: TRUE


Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

In DC202 FabricPath network topology there are four equal cost paths from one leaf switch to any other leaf switch, except its vPC+ peer.

To Verify the FabricPath unicast ECMP path and load-balancing in leaf switchs - Nexus 5000, Nexus 6000 and Nexus 7700, use the following commands.

Display Information About All FabricPath Topology Interfaces:

```
N5K-705# sh FabricPath topology interface
Interface          Topo-Description               Topo-ID    Topo-IF-State
```

```
------------------ ------------------------------ ---------- ------------
port-channel51      0                               0              Up
port-channel52      0                               0              Up
port-channel53      0                               0              Up
port-channel54      0                               0              Up
port-channel706     0                               0              Up
```

Display All FabricPath IS-IS Adjacency Information:

```
N5K-705# sh FabricPath isis adjacency
FabricPath IS-IS domain: default FabricPath IS-IS adjacency database:
System ID       SNPA            Level  State  Hold Time   Interface
DC202-51        N/A             1      UP     00:00:30    port-channel51
DC202-52        N/A             1      UP     00:00:32    port-channel52
DC202-53        N/A             1      UP     00:00:23    port-channel53
DC202-54        N/A             1      UP     00:00:23    port-channel54
DC202-706       N/A             1      UP     00:00:24    port-channel706
```

Display the FabricPath Layer 2 IS-IS Routing Table for Unicast Routes:

```
N5K-705# sh FabricPath isis route
FabricPath IS-IS domain: default MT-0
Topology 0, Tree 0, Swid routing table
251, L1
 via port-channel51, metric 20
252, L1
 via port-channel52, metric 20
253, L1
 via port-channel53, metric 20
254, L1
 via port-channel54, metric 20
271, L1
 via port-channel51, metric 40
 via port-channel53, metric 40
 via port-channel52, metric 40
 via port-channel54, metric 40
272, L1
 via port-channel51, metric 40
 via port-channel53, metric 40
 via port-channel52, metric 40
 via port-channel54, metric 40
273, L1
 via port-channel51, metric 40
 via port-channel53, metric 40
 via port-channel52, metric 40
 via port-channel54, metric 40
274, L1
 via port-channel51, metric 40
 via port-channel53, metric 40
 via port-channel52, metric 40
 via port-channel54, metric 40
276, L1
 via port-channel706, metric 20
277, L1
 via port-channel51, metric 220
 via port-channel53, metric 220
 via port-channel52, metric 220
 via port-channel54, metric 220
300, L1
 via port-channel51, metric 20
 via port-channel52, metric 20
400, L1
 via port-channel53, metric 20
 via port-channel54, metric 20
```

```
700, L1
 via port-channel51, metric 40
 via port-channel53, metric 40
 via port-channel52, metric 40
 via port-channel54, metric 40
706, L1
 via port-channel706, metric 20
707, L1
 via port-channel51, metric 220
 via port-channel53, metric 220
 via port-channel52, metric 220
 via port-channel54, metric 220
708, L1
 via port-channel51, metric 220
 via port-channel53, metric 220
 via port-channel52, metric 220
 via port-channel54, metric 220
709, L1
 via port-channel51, metric 25
 via port-channel52, metric 25
710, L1
 via port-channel51, metric 25
 via port-channel52, metric 25
759, L1
 via port-channel51, metric 25
 via port-channel52, metric 25
```

Display Unicast Routes to Switch-ID 271:

```
N5k-705# sh l2 route switchid 271
FabricPath Unicast Route Table
'a/b/c' denotes ftag/switch-id/subswitch-id
'[x/y]' denotes [admin distance/metric]
ftag 0 is local ftag
subswitch-id 0 is default subswitch-id


FabricPath Unicast Route Table for Topology-Default

1/271/0, number of next-hops: 4
        via Po51, [115/40], 0 day/s 12:18:25, isis_FabricPath-default
        via Po52, [115/40], 0 day/s 04:00:31, isis_FabricPath-default
        via Po53, [115/40], 0 day/s 10:54:47, isis_FabricPath-default
        via Po54, [115/40], 0 day/s 03:59:12, isis_FabricPath-default
```

Display FabricPath Unicast Ftag Information:

```
N5K-705# sh FabricPath topology ftag unicast
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1
```

Display which Path the FabricPath Unicast Load-balancing Utilizes for a Given Flow:

```
N5K-705# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 dst-ip 201.11.7.1
Missing params will be substituted by 0's.

crc8_hash: 28
This flow selects interface Po51

N5K-705# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 dst-ip 201.11.7.2

Missing params will be substituted by 0's.

crc8_hash: 42
```

```
This flow selects interface Po53

N6K-707# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 dst-ip 201.11.7.1
Missing params will be substituted by 0's.

hash select: CRC10d (12)
crc8_hash  : 62
This flow selects interface Po53

N6K-707# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 dst-ip 201.11.7.2
Missing params will be substituted by 0's.

hash select: CRC10d (12)
crc8_hash  : 32
This flow selects interface Po51

N7700-709# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 flow-type l3 src-ip
201.11.27.1 dst-ip 230.201.0.1 vlan 11 module 1
This flow selects interface Po52

N7700-709# sh FabricPath load-balance unicast forwarding-path ftag 1 switchid 271 flow-type l3 src-ip
201.11.27.2 dst-ip 230.201.0.1 vlan 11 module 1
This flow selects interface Po51
```

3.2.4.3.3 FabricPath Multicast Load-Balancing and Verification

In the DC202 FabricPath topology excerpt shown in Figure 28, the multicast traffic source is located on the L2 switch, 27, and the receiver is located on the L2 switch, 7003. Multicast traffic that reaches the spine 53, selects FTag 1 and uses tree 1 to forward the multicast data to the receiver which is attached to the leaf switch, 706. Note that the multicast traffic is also forwarded to all other spines because of PIM neighborship.

Figure 28 FabricPath Ftag 1 Multi-Destination Tree

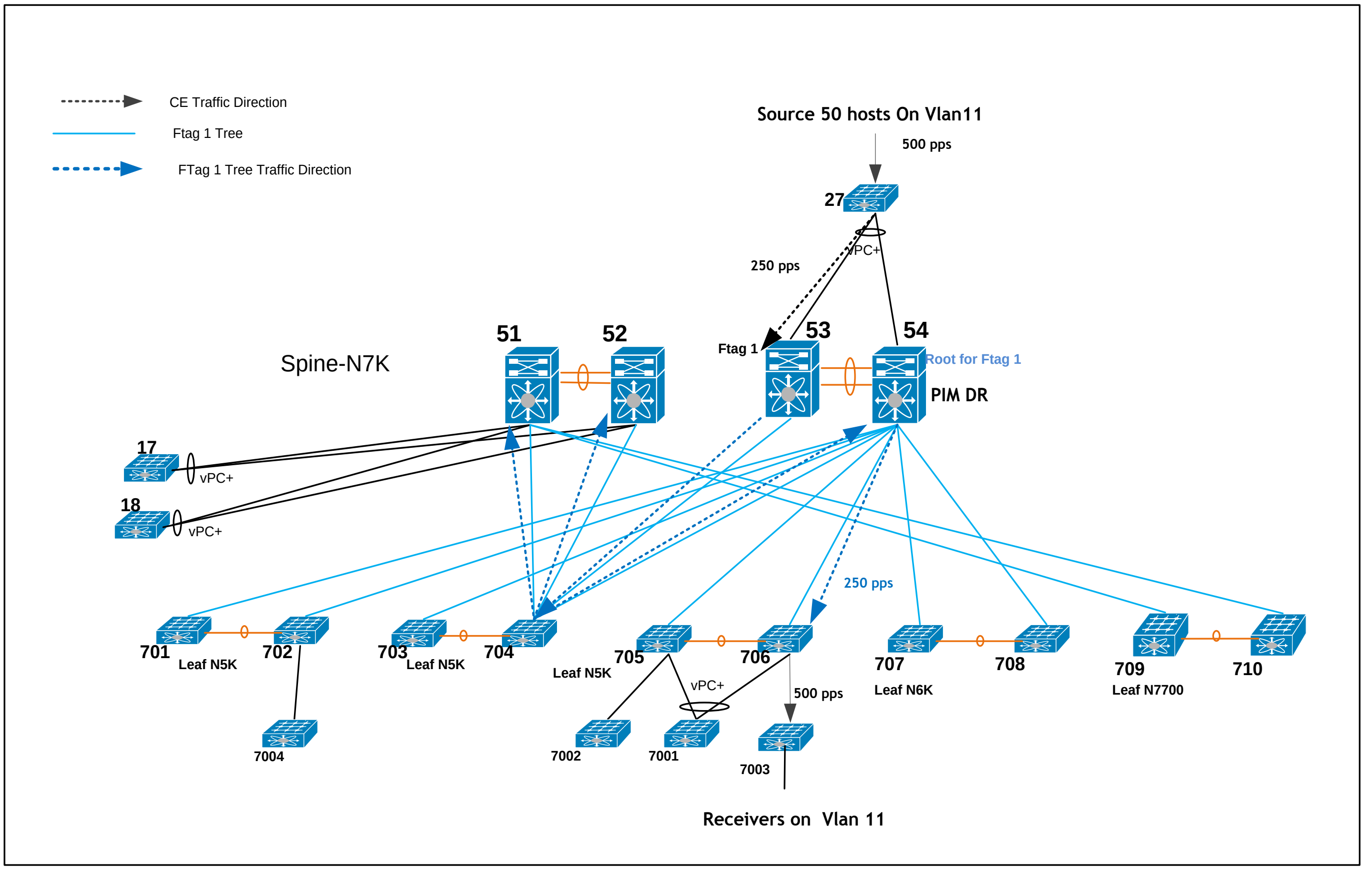

The Multicast traffic that reaches the spine 54, selects FTag 2 and uses tree 2 to forward the multicast data to the receiver which is attached to the leaf switch, 706. Note that this multicast traffic is also forwarded to all other spines because of PIM neighborship.

Figure 29 FabricPath Ftag 2 Multi-Destination Tree

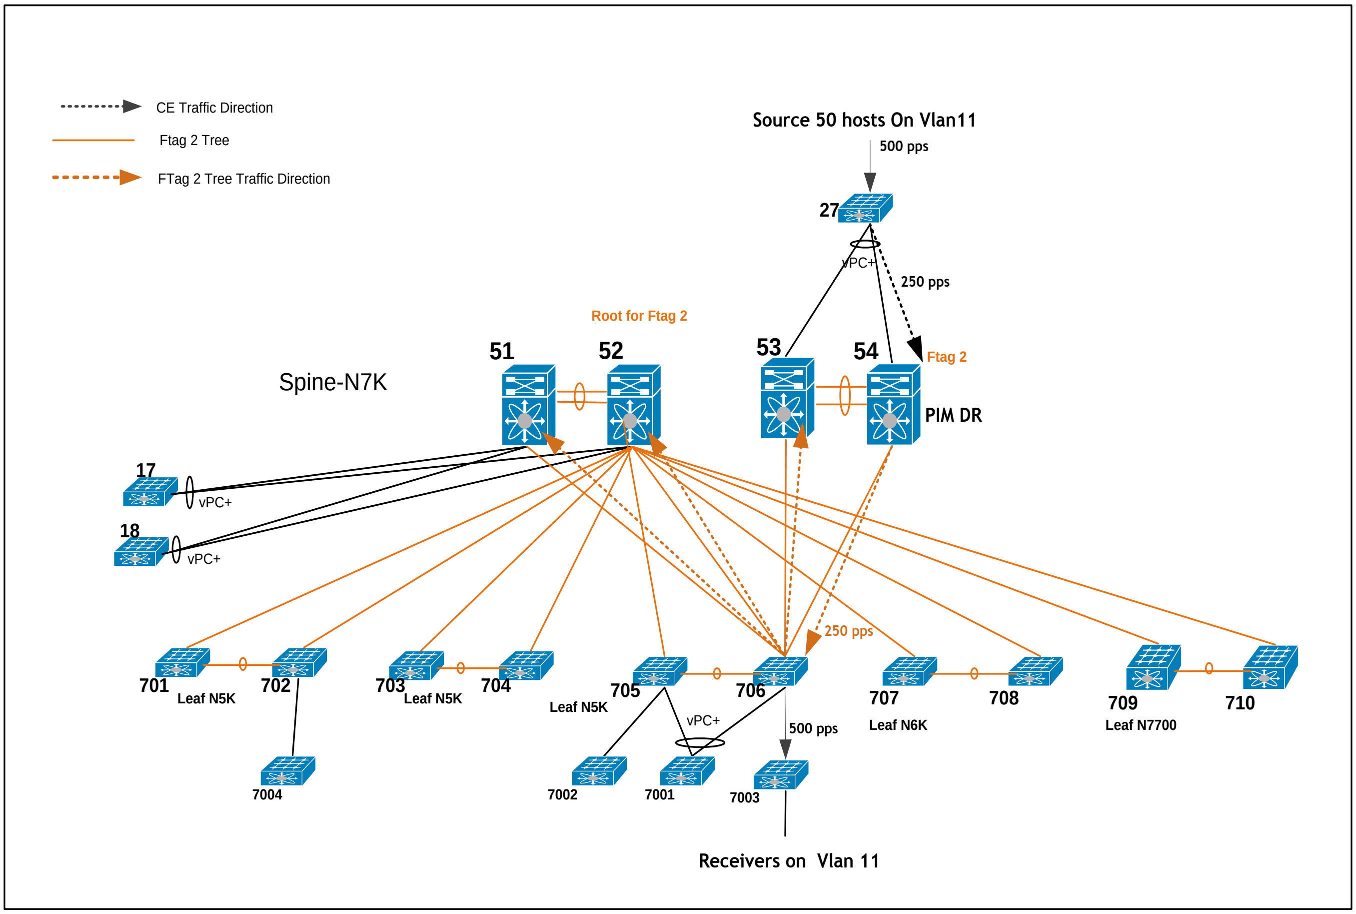

The hashing to either multi-destination tree is platform-dependent and the hash function is per flow. The default multicast load balancing mechanism for Nexus 7000 F1 VDC uses a symmetric hash input combining both Layer 3 (source and destination IP addresses) and Layer 4 (source and destination TCP and UDP port numbers, if present) information, as well as the VLAN ID; while in Nexus 7000 F2 VDC it does not include the VLAN ID. The default multicast load balancing mechanism for the Nexus 5000 uses symmetric hash with Layer 2/Layer 3/Layer 4 source and destination addresses, as well as VLAN ID.

DC202 has changed Nexus 7000 F2 VDC multicast load balancing mechanism to include Vlan while leaving Nexus 7000 F1/M1 VDC and Nexus 5000 as default.

```
DC202-51(config)#  FabricPath load-balance multicast rotate-amount 0x3 include-vlan

N7k-51(config)# sh run FabricPath  | in "multicast"
FabricPath load-balance multicast rotate-amount 0x3 include-vlan

N7k-51(config)# sh FabricPath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
```

```
Rotate amount: 11 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

To Verify the FabricPath multicast load-balancing path for a given multicast group in Nexus 7000, use the following commands.

Display the IP Multicast Routes for VLAN 11, Group 230.202.0.1

```
DC202-54# sh FabricPath isis ip mroute vlan 11 group 230.202.0.1
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.202.0.1)
  Outgoing interface list: (count: 6)
    SWID: 0xfb (251)
    SWID: 0xfc (252)
    SWID: 0xfd (253)
    SWID: 0x110 (272)
    SWID: 0x113 (275)
    SWID: 0x114 (276)
```

Display FabricPath Multicast Routes for VLAN 11:

```
DC202-54# sh FabricPath mroute vlan 11

(vlan/11, 0.0.0.0, 224.0.1.39), uptime: 19:21:48, isis igmp
 Outgoing interface list: (count: 4)
  Interface Vlan11, [SVI] uptime: 19:20:25, igmp
  Switch-id 251, uptime: 16:16:00, isis
  Switch-id 252, uptime: 19:20:23, isis
  Switch-id 253, uptime: 16:15:06, isis

(vlan/11, 0.0.0.0, 224.0.1.40), uptime: 19:21:48, isis igmp
 Outgoing interface list: (count: 4)
  Interface Vlan11, [SVI] uptime: 19:20:25, igmp
  Switch-id 251, uptime: 16:16:07, isis
  Switch-id 252, uptime: 19:20:23, isis
  Switch-id 253, uptime: 16:15:06, isis

(vlan/11, 0.0.0.0, 230.202.0.1), uptime: 16:00:05, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 16:00:04, igmp
  Switch-id 251, uptime: 16:00:04, isis
  Switch-id 252, uptime: 16:00:04, isis
  Switch-id 253, uptime: 16:00:04, isis
  Switch-id 272, uptime: 16:00:04, isis
  Switch-id 275, uptime: 16:00:05, isis
  Switch-id 276, uptime: 16:00:05, isis

 (vlan/11, *, *), Flood, uptime: 19:21:48, isis
 Outgoing interface list: (count: 9)
  Switch-id 251, uptime: 16:16:16, isis
  Switch-id 252, uptime: 19:21:48, isis
  Switch-id 253, uptime: 16:15:16, isis
  Switch-id 271, uptime: 19:21:48, isis
```

```
  Switch-id 272, uptime: 19:21:48, isis
  Switch-id 273, uptime: 19:21:48, isis
  Switch-id 274, uptime: 19:21:48, isis
  Switch-id 275, uptime: 19:21:48, isis
  Switch-id 276, uptime: 19:21:48, isis

(vlan/11, *, *), Router ports (OMF), uptime: 19:21:48, isis igmp
 Outgoing interface list: (count: 5)
  Interface Vlan11, [SVI] uptime: 19:21:48, igmp
  Interface port-channel53, uptime: 19:21:48, igmp
  Switch-id 251, uptime: 16:16:14, isis
  Switch-id 252, uptime: 19:21:48, isis
  Switch-id 253, uptime: 16:15:16, isis

Found total 5 route(s)
```

Display FabricPath Topology FTag Information:

```
DC202-54# sh FabricPath topology ftag multicast
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1
0                         0          2         2
DC202-54# sh FabricPath topology ftag active
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          2         2
```

Display FabricPath Multicast Load-balancing Information:

```
DC202-51# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 202.11.27.1 dst-ip
230.202.0.1 vlan 11 module 2
128b Hash Key generated : 00 00 00 32 82 c6 c0 40 00 00 00 00 32 00 00 00
0xf2
    FTAG SELECTED IS : 1 (HASH 242)
```

Display FabricPath Multicast Route for VLAN 11, Ftag 2:

```
DC202-54# sh FabricPath mroute vlan 11 ftag 2

(ftag/2, vlan/11, 0.0.0.0, 224.0.1.39), uptime: 19:25:55, isis igmp
 Outgoing interface list: (count: 4)
  Interface Vlan11, [SVI] uptime: 19:24:32, igmp
  Interface port-channel706,   Switch-id 251, uptime: 16:20:24, isis
  Interface port-channel706,   Switch-id 252, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 253, uptime: 16:19:23, isis

(ftag/2, vlan/11, 0.0.0.0, 224.0.1.40), uptime: 19:25:55, isis igmp
 Outgoing interface list: (count: 4)
  Interface Vlan11, [SVI] uptime: 19:24:32, igmp
  Interface port-channel706,   Switch-id 251, uptime: 16:20:24, isis
  Interface port-channel706,   Switch-id 252, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 253, uptime: 16:19:23, isis

 (ftag/2, vlan/11, 0.0.0.0, 230.202.0.1), uptime: 16:04:12, isis igmp
 Outgoing interface list: (count: 7)
  Interface port-channel27, uptime: 16:04:11, igmp
  Interface port-channel706,   Switch-id 251, uptime: 16:20:24, isis
```

```
  Interface port-channel706,   Switch-id 252, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 253, uptime: 16:19:23, isis
  Interface port-channel706,   Switch-id 272, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 275, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 276, uptime: 19:25:55, isis

 (ftag/2, vlan/11, *, *), Flood, uptime: 19:25:55, isis
 Outgoing interface list: (count: 9)
  Interface port-channel706,   Switch-id 251, uptime: 16:20:24, isis
  Interface port-channel706,   Switch-id 252, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 253, uptime: 16:19:23, isis
  Interface port-channel706,   Switch-id 271, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 272, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 273, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 274, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 275, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 276, uptime: 19:25:55, isis

(ftag/2, vlan/11, *, *), Router ports (OMF), uptime: 19:25:55, isis igmp
 Outgoing interface list: (count: 5)
  Interface Vlan11, [SVI] uptime: 19:25:55, igmp
  Interface port-channel53, uptime: 19:25:55, igmp
  Interface port-channel706,   Switch-id 251, uptime: 16:20:24, isis
  Interface port-channel706,   Switch-id 252, uptime: 19:25:55, isis
  Interface port-channel706,   Switch-id 253, uptime: 16:19:23, isis

Found total 5 route(s)
```

To Verify the traffic path for a given multicast group in Nexus 5000/6000/7700 leaf switch, use the following commands.

Display the IP Multicast Routes for VLAN 11, Group 230.202.0.1

```
N5K-706# sh FabricPath isis ip mroute vlan 11 group 230.202.0.1
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.202.0.1)
  Outgoing interface list: (count: 6)
    SWID: 0xfb (251)
    SWID: 0xfc (252)
    SWID: 0xfd (253)
    SWID: 0xfe (254)
    SWID: 0x110 (272)
    SWID: 0x113 (275)

N6K-708# sh FabricPath isis ip mroute vlan 11 group 230.202.0.1
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.202.0.1)
  Outgoing interface list: (count: 7)
    SWID: 0xfb (251)
    SWID: 0xfc (252)
    SWID: 0xfd (253)
    SWID: 0xfe (254)
    SWID: 0x110 (272)
    SWID: 0x113 (275)
    SWID: 0x114 (276)

N7700-709# sh FabricPath isis ip mroute vlan 11 group 230.202.0.1
```

```
FabricPath IS-IS domain: default
FabricPath IS-IS IPv4 Multicast Group database
VLAN 11: (*, 230.202.0.1)
  Outgoing interface list: (count: 7)
    SWID: 0xfb (251)
    SWID: 0xfc (252)
    SWID: 0xfd (253)
    SWID: 0xfe (254)
    SWID: 0x110 (272)
    SWID: 0x113 (275)
    SWID: 0x114 (276)
```

Display FabricPath Multicast Routes for VLAN 11 on the Leaf’s (Nexus 5000/6000/7700):

```
N5K-706# sh FabricPath mroute vlan 11

(vlan/11, 0.0.0.0, 224.0.1.39), uptime: 6d21h, isis
 Outgoing interface list: (count: 4)
  Switch-id 251, uptime: 16:54:03, isis
  Switch-id 252, uptime: 19:58:27, isis
  Switch-id 253, uptime: 16:53:09, isis
  Switch-id 254, uptime: 19:58:27, isis

(vlan/11, 0.0.0.0, 224.0.1.40), uptime: 6d21h, isis
 Outgoing interface list: (count: 4)
  Switch-id 251, uptime: 16:54:10, isis
  Switch-id 252, uptime: 19:58:27, isis
  Switch-id 253, uptime: 16:53:09, isis
  Switch-id 254, uptime: 19:58:27, isis

(vlan/11, 0.0.0.0, 230.202.0.1), uptime: 16:38:09, isis igmp
 Outgoing interface list: (count: 8)
  Switch-id 251, uptime: 16:38:07, isis
  Switch-id 252, uptime: 16:38:07, isis
  Switch-id 253, uptime: 16:38:07, isis
  Switch-id 254, uptime: 16:38:07, isis
  Switch-id 272, uptime: 16:38:07, isis
  Switch-id 275, uptime: 16:38:08, isis
  Interface port-channel1, uptime: 16:38:09, igmp
  Interface port-channel3, uptime: 16:38:08, igmp

 (vlan/11, *, *), Flood, uptime: 7w6d, isis
 Outgoing interface list: (count: 9)
  Switch-id 251, uptime: 16:54:19, isis
  Switch-id 252, uptime: 21:14:55, isis
  Switch-id 253, uptime: 16:53:20, isis
  Switch-id 254, uptime: 21:13:45, isis
  Switch-id 271, uptime: 4w2d, isis
  Switch-id 272, uptime: 7w6d, isis
  Switch-id 273, uptime: 7w6d, isis
  Switch-id 274, uptime: 7w6d, isis
  Switch-id 275, uptime: 7w6d, isis

(vlan/11, *, *), Router ports (OMF), uptime: 7w6d, isis
 Outgoing interface list: (count: 4)
  Switch-id 251, uptime: 16:54:17, isis
  Switch-id 252, uptime: 21:14:55, isis
  Switch-id 253, uptime: 16:53:20, isis
  Switch-id 254, uptime: 21:13:43, isis

Found total 5 route(s)
```

Display FabricPath Topology FTag Information:

```
N5K-706# sh FabricPath topology ftag multicast
```

```
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1
0                         0          2         2

N5K-706# sh FabricPath topology ftag active
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1

N7700-709# sh FabricPath topology ftag multicast
Topo-Description          Topo-ID    Graph-ID  Ftag
------------------------- ---------- --------- ---------
0                         0          1         1
0                         0          2         2
```

Display FabricPath Multicast Load-balancing Information:

```
N5K-706# sh FabricPath load-balance multicast ftag-selected vlan 11 macg 0100.5e4d.0001

 If the traffic is received on a non-vPC port:
 Ftag selected : 1

 If the traffic is received on a vPC port:
 Ftag selected : 1

======================================

 Vlan : 11 (int_vlan : 15)
 Macg : 0100.5e4d.0001

 Hash-key : 0x000f0000 00000000
 Hash-val : 34
 Num_trees : 2

======================================

N6K-708# sh FabricPath load-balance multicast ftag-selected vlan 11 macg 0100.5e4d.0001

 If the traffic is received on a non-vPC port:
 Ftag selected : 2

 If the traffic is received on a vPC port:
 Ftag selected : 1

======================================

 Vlan : 11 (int_vlan : 3957)
 Macg : 0100.5e4d.0001

 Hash-key : 0x0f750000 00000000
 Hash-val : 819
 Num_trees : 2

 Hash-val 6-bits: 51
 Offset : 1

======================================

N7700-709# sh FabricPath load-balance multicast ftag-selected flow-type l3 src-ip 202.11.27.1 dst-ip
230.202.0.1 vlan 11 module 1
128b Hash Key generated : 00 32 82 c6 c0 79 b2 80 00 48 88 00 00 00 00 00
0x2
    FTAG SELECTED IS : 2 (HASH 2)
```

Display FabricPath Multicast Route for VLAN 11, Ftag 1 on Leaf's Nexus 5000/6000/7700:

```
N5K-706# sh FabricPath mroute vlan 11 ftag 1


(ftag/1, vlan/11, 0.0.0.0, 224.0.1.39), uptime: 6d21h, isis
 Outgoing interface list: (count: 4)
  Interface port-channel54, uptime: 17:02:31, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 17:01:30, isis
  Interface port-channel54, uptime: 21:21:33, isis

(ftag/1, vlan/11, 0.0.0.0, 224.0.1.40), uptime: 6d21h, isis
 Outgoing interface list: (count: 4)
  Interface port-channel54, uptime: 17:02:31, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 17:01:30, isis
  Interface port-channel54, uptime: 21:21:33, isis

(ftag/1, vlan/11, 0.0.0.0, 230.202.0.1), uptime: 16:46:22, isis igmp
 Outgoing interface list: (count: 8)
  Interface port-channel54, uptime: 17:02:31, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 17:01:30, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel1, uptime: 16:46:22, igmp
  Interface port-channel3, uptime: 16:46:20, igmp

 (ftag/1, vlan/11, *, *), Flood, uptime: 7w6d, isis
 Outgoing interface list: (count: 9)
  Interface port-channel54, uptime: 17:02:31, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 17:01:30, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 21:21:33, isis

(ftag/1, vlan/11, *, *), Router ports (OMF), uptime: 7w6d, isis
 Outgoing interface list: (count: 4)
  Interface port-channel54, uptime: 17:02:31, isis
  Interface port-channel54, uptime: 21:21:33, isis
  Interface port-channel54, uptime: 17:01:30, isis
  Interface port-channel54, uptime: 21:21:33, isis

Found total 5 route(s)
```

3.2.5 Fabric Extenders (FEX)

The Fabric Extender integrates with its parent switch, which is a Cisco Nexus Series device, to allow automatic provisioning and configuration taken from the settings on the parent device.

The Fabric Interface is an uplink port that is designated for connection from the Fabric Extender to its parent switch. A fabric interface cannot be used for any other purpose. It must be directly connected to the parent switch. Multiple fabric interfaces can be combined together to form a port-channel fabric interface.

On DC2, FEX is attached to Nexus 5000, Nexus 6000 and Nexus 7000. For Nexus 7000, beginning with Cisco NX-OS Release 6.1(3), a minimum number of links for the FEX fabric port channel can be configured so that when a certain number of FEX fabric port-channel member ports go down, the host-facing interfaces of the FEX are suspended.

The host interfaces are Ethernet host interfaces for connection to a server or host system.

```
feature-set fex

fex 101
  pinning max-links 1
  description FEX0101

! Port-channel fabric interface
interface port-channel101
  switchport mode fex-fabric
  fex associate 101

interface Ethernet1/41
  switchport mode fex-fabric
  fex associate 101
  channel-group 101
  no shutdown

! Port-channel host interface
interface port-channel1001
  switchport mode trunk
  switchport trunk allowed vlan 1,1000-1999
  spanning-tree port type edge trunk
  vpc 1001

interface Ethernet101/1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,1000-1999
  spanning-tree port type edge trunk
  channel-group 1001 mode active
```

Display the Fabric Extenders Attached to the System:

```
N5K-704# sh fex
  FEX         FEX           FEX               FEX              Fex
Number    Description      State             Model            Serial
------------------------------------------------------------------------
101        FEX0101                Online   N2K-C2248TP-E-1GE   SSI163704DZ
102        FEX0102                Online     N2K-C2224TP-1GE   SSI15480FC7
103        FEX0103                Online   N2K-C2248TP-E-1GE   SSI16370GLQ
104        FEX0104                Online   N2K-C2248TP-E-1GE   SSI16370GLM

N7K# sh fex
  FEX         FEX           FEX                          FEX
Number    Description      State             Model            Serial
------------------------------------------------------------------------
111        FEX0111                Online   N2K-C2248TP-E-1GE   SSI171608PZ
112        FEX0112                Online   N2K-C2248TP-E-1GE   FOX1724GZHW
113        FEX0113                Online   N2K-C2248TP-E-1GE   FOX1724G9DJ
114        FEX0114                Online   N2K-C2248TP-E-1GE   FOX1724G9DM
115        FEX0115                Online   N2K-C2248TP-E-1GE   SSI17160D9M
116        FEX0116                Online   N2K-C2248TP-E-1GE   SSI171608R9
117        FEX0117                Online   N2K-C2248TP-E-1GE   FOX1724G9G7
118        FEX0118                Online   N2K-C2248TP-E-1GE   FOX1725G4D4
119        FEX0119                Online   N2K-C2248TP-E-1GE   FOX1724G9BG
120        FEX0120                Online   N2K-C2248TP-E-1GE   FOX1724GZL5
```

```
121        FEX0121                 Online   N2K-C2248TP-E-1GE   FOX1724G9E5
122        FEX0122                 Online   N2K-C2248TP-E-1GE   SSI171708HV
123        FEX0123                 Online   N2K-C2248TP-E-1GE   FOX1724GXGF
124        FEX0124                 Online   N2K-C2248TP-E-1GE   FOX1724G9DP
```

Since the FEX host interfaces are supposed to be connected directly to hosts, certain defaults should be noted as shown below. Also, **cdp** is **not** supported on the Fabric Extenders connected to a Nexus 5000/6000/7000 parent switch.

```
N5K-704# sh run interface ethernet 101/1/1 all
interface Ethernet101/1/1
  no description
  lacp port-priority 32768
  lacp rate normal
  priority-flow-control mode auto
  lldp transmit
  lldp receive
  no switchport block unicast
  no switchport block multicast
  hardware multicast hw-hash
  no hardware vethernet mac filtering per-vlan
  cdp enable
  switchport
  switchport mode trunk
  no switchport dot1q ethertype
  no switchport priority extend
  switchport trunk allowed vlan 1,1000-1250
  spanning-tree port-priority 128
  spanning-tree cost auto
  spanning-tree link-type auto
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  no spanning-tree bpdufilter
  speed auto
  duplex auto
  flowcontrol receive off
  flowcontrol send on
  no link debounce
  no beacon
  delay 1
  snmp trap link-status
  logging event port link-status default
  logging event port trunk-status default
  mdix auto
  storm-control broadcast level 100.00
  storm-control multicast level 100.00
  storm-control unicast level 100.00
  no shutdown lan
  load-interval counter 1 30
  load-interval counter 2 300
  no load-interval counter 3
  medium broadcast
  channel-group 1001 mode active
  no shutdown
```

3.3 DC3 Core Network and Configuration

DC3 core is composed of three Nexus 7000 devices.

3.3.1 Configuration of Platform Specific Features On DC3 Core

3.3.1.1 Licensing

License Usage in DC3 core:

```
DC3-3# sh license usage
Feature                       Ins  Lic   Status Expiry Date Comments
                                  Count
--------------------------------------------------------------------------------
MPLS_PKG                       Yes   -   Unused Never       -
STORAGE-ENT                    No    -   Unused             -
VDC_LICENSES                   Yes   4   Unused Never       -
ENTERPRISE_PKG                 No    -   Unused             -
FCOE-N7K-F132XP                No    0   Unused             -
FCOE-N7K-F248XP                No    0   Unused             -
ENHANCED_LAYER2_PKG            Yes   -   Unused Never       -
SCALABLE_SERVICES_PKG          Yes   -   Unused Never       -
TRANSPORT_SERVICES_PKG         Yes   -   Unused Never       -
LAN_ADVANCED_SERVICES_PKG      Yes   -   Unused Never       -
LAN_ENTERPRISE_SERVICES_PKG    Yes   -   In use Never       -
```

3.3.1.2 Out-of-Band Management Network

DC3 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic.

Configuration:

```
interface mgmt0
  vrf member management
  ip address 10.2.3.15/16
```

3.3.1.3 Common Configurations

3.3.1.3.1 SSH and TACACS+

SSH is enabled in DC3 to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration and Verification:

```
feature tacacs+


ip tacacs source-interface mgmt0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
use-vrf management

DC3-3# sh ssh server
ssh version 2 is enabled

DC3-3# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
```

```
interop      pts/0        Feb 18 09:57 06:44        13662 (172.28.92.47) session=ssh
interop      pts/1        Feb 18 17:09    .         19327 (172.28.92.47) session=ssh *
```

3.3.1.3.2 CDP and LLDP

CDP is pervasively used on the DC3 core test bed for inter-device discovery.

CDP Configuration and Verification:

```
DC3-3# sh run cdp all

!Command: show running-config cdp all
!Time: Tue Feb 18 17:12:41 2014

version 6.2(6)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
no cdp format device-id system-name


interface Ethernet7/1
  cdp enable

interface Ethernet7/2
  cdp enable

DC3-3# sh cdp neighbors interface mgmt 0
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform      Port ID
mgmt-sw1.interop.cisco.com
                    mgmt0          170    R S I      WS-C6509-E    Gig2/44
```

3.3.1.3.3 Syslog

Syslog is used to record all network events on the DC3 core test bed. Whenever possible, NVT uses a separate management VRF for syslog.

Configuration and Verification:

```
logging server syslog.interop.cisco.com 7 use-vrf management facility local6

DC3-3# sh logging server
Logging server:                 enabled
{syslog.interop.cisco.com}
        server severity:        debugging
        server facility:        local6
        server VRF:             management
```

3.3.1.3.4 SNMP

SNMP is used for system monitoring in NVT DC3 core. Scripts are used to poll the systems asynchronously during the course of all NVT test execution.

Configuration:

```
version 6.2(6)
snmp-server source-interface trap mgmt0
snmp-server user admin vdc-admin auth md5 0x1ba9d057b6e00be1e3ca7f59269ddaa1 priv
0x1ba9d057b6e00be1e3ca7f59269ddaa1 localizedkey
snmp-server host 172.28.92.62 traps version 2c public
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server enable traps callhome event-notify
snmp-server enable traps callhome smtp-send-fail
snmp-server enable traps cfs state-change-notif
snmp-server enable traps cfs merge-failure
snmp-server enable traps aaa server-state-change
snmp-server enable traps feature-control FeatureOpStatusChange
snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended
snmp-server enable traps config ccmCLIRunningConfigChanged
snmp-server enable traps snmp authentication
snmp-server enable traps link cisco-xcvr-mon-status-chg
snmp-server enable traps vtp notifs
snmp-server enable traps vtp vlancreate
snmp-server enable traps vtp vlandelete
snmp-server enable traps bridge newroot
snmp-server enable traps bridge topologychange
snmp-server enable traps stpx inconsistency
snmp-server enable traps stpx root-inconsistency
snmp-server enable traps stpx loop-inconsistency
snmp-server enable traps system Clock-change-notification
snmp-server enable traps feature-control ciscoFeatOpStatusChange
snmp-server community private group vdc-admin
snmp-server community public group vdc-operator
snmp-server community cisco group vdc-operator
```

3.3.1.3.5 NTP

NTP is used to synchronize the clocks on all DC3 core devices to provide consistent timestamps on all network logs and events.

Configuration and Verification:

```
ntp distribute
ntp server 172.28.92.1 use-vrf management
ntp commit

DC3-3a# sh ntp status
Distribution : Enabled
Last operational state: No session

DC3-3a# sh ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                 8   64     377   0.00107 management
```

3.3.1.3.6 SPAN

SPAN has been enabled on NVT switches to allow packet captures to assist in network debugging.

Configuration and Verification:

```
monitor session 1
  source interface port-channel41 both
  destination interface Ethernet7/14
  no shut

interface Ethernet7/14
  switchport
  switchport monitor
  no shutdown

DC3-4# sh monitor session 1
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Po41
    tx            : Po41
    both          : Po41
source VLANs      :
    rx            :
    tx            :
    both          :
source exception  :
filter VLANs      : filter not specified
destination ports : Eth7/14

Feature       Enabled   Value   Modules Supported       Modules Not-Supported
-----------------------------------------------------------------------------
MTU-Trunc     No
rate-limit-rx No
rate-limit-tx No
Sampling      No
MCBE          No
L3-TX         -         -       5                       7  8  10
RB span       No


Legend:
  MCBE  = Multicast Best Effort
  L3-TX = L3 Multicast Egress SPAN
  ExSP-X = Exception Span for type X (L3, FP, or misc)
```

3.3.1.3.7 DNS

DNS has been enabled to provide name lookup in NVT network.

Configuration and Verification:

```
vrf context management
  ip domain-lookup
  ip domain-name interop.cisco.com
  ip domain-list cisco.com
  ip domain-list interop.cisco.com
  ip name-server 172.28.92.9 172.28.92.10

DC3-3# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=1.73 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=1.428 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=1.432 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=1.344 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=1.328 ms
```

```
--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.328/1.452/1.73 ms
```

3.3.1.3.8 UDLD

UDLD is used to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a device detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. In aggressive mode, if the link state of the port is determined to be bi-directional and the UDLD information times out while the link on the port is still up, UDLD tries to re-establish the state of the port.

Configuration:

```
version 6.2(6)
feature udld

udld aggressive


DC3-3# sh udld neighbors
Port             Device Name    Device ID   Port ID         Neighbor State
---------------------------------------------------------------------------
Ethernet7/13     TBM16492630    1           Ethernet7/13    bidirectional
Ethernet7/17     FOC1712R0D6    1           Ethernet1/49    bidirectional
Ethernet7/18     FOC1711R1LP    1           Ethernet1/49    bidirectional
Ethernet7/19     FOC1711R1U9    1           Ethernet1/49    bidirectional
```

3.3.1.3.9 MTU

System MTU is configured to jumbo MTU on DC3 core test bed.

Configuration and Verification:

```
interface Ethernet8/18
  mtu 9216
  channel-group 62 mode active
  no shutdown

DC3-3# sh int e8/18
Ethernet8/18 is down (SFP not inserted)
admin state is up, Dedicated Interface
  Belongs to Po62
  Hardware: 1000/10000 Ethernet, address: 64a0.e73f.a4c2 (bia 6c9c.ed47.fd81)
  MTU 9216 bytes, BW 10000000 Kbit, DLY 10 usec
  reliability 255/255, txload 1/255, rxload 1/255
```

3.3.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

When the switch comes up for the first time, there are multiple CoPP configuration templates that are presented: *strict, moderate, lenient and dense*. NVT has chosen the *lenient* template on DC3 core.

Configuration on Nexus 7000 for release 6.2.x:

```
copp profile lenient
```

Default Lenient CoPP on Nexus 7000 for Software Release 6.2.x as Used in NVT DC3

```
policy-map type control-plane copp-system-p-policy-lenient
  class copp-system-p-class-critical
    set cos 7
    police cir 36000 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-important
    set cos 6
    police cir 1400 kbps bc 1500 ms conform transmit violate drop
  class copp-system-p-class-multicast-router
    set cos 6
    police cir 2600 kbps bc 1000 ms conform transmit violate drop
  class copp-system-p-class-management
    set cos 2
    police cir 10000 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-multicast-host
    set cos 1
    police cir 1000 kbps bc 1000 ms conform transmit violate drop
  class copp-system-p-class-normal
    set cos 1
    police cir 680 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-ndp
    set cos 6
    police cir 680 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-normal-dhcp
    set cos 1
    police cir 1500 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-normal-dhcp-relay-response
    set cos 1
    police cir 1800 kbps bc 750 ms conform transmit violate drop
  class copp-system-p-class-redirect
    set cos 1
    police cir 280 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-exception
    set cos 1
    police cir 360 kbps bc 375 ms conform transmit violate drop
  class copp-system-p-class-monitoring
    set cos 1
    police cir 130 kbps bc 1500 ms conform transmit violate drop
  class copp-system-p-class-l2-unpoliced
    police cir 8 gbps bc 5 mbytes conform transmit violate transmit
  class copp-system-p-class-undesirable
    set cos 0
    police cir 32 kbps bc 375 ms conform drop violate drop
  class copp-system-p-class-fcoe
    set cos 6
    police cir 1060 kbps bc 1500 ms conform transmit violate drop
  class copp-system-p-class-l2-default
    police cir 100 kbps bc 375 ms conform transmit violate drop
  class class-default
    set cos 0
    police cir 100 kbps bc 250 ms conform transmit violate drop
```

3.3.1.5 Rate Limiters

Rate limiters are an additional set of features on Nexus 7000 to prevent undesirable packets from overwhelming the CPU on the supervisor module.

Default Values:

```
Dc3-3# show hardware rate-limiter

Units for Config: packets per second
```

```
Allowed, Dropped & Total: aggregated since last clear counters


Module: 3
  R-L Class           Config           Allowed         Dropped            Total
 +------------------+--------+---------------+---------------+-----------------+
  L3 mtu                 500              436               0               436
  L3 ttl                 500           171234        14981787          15153021
  L3 control           10000                0               0                 0
  L3 glean               100              823            6036              6859
  L3 mcast dirconn    Disable
  L3 mcast loc-grp      3000                0               0                 0
  L3 mcast rpf-leak      500              165               0               165
  L2 storm-ctrl       Disable
  access-list-log        100                0               0                 0
  copy                 30000         16351350               0          16351350
  receive              30000          9922819               0           9922819
  L2 port-sec            500                0               0                 0
  L2 mcast-snoop       10000                0               0                 0
  L2 vpc-low            4000                0               0                 0
  L2 l2pt                500                0               0                 0
  f1 rl-1               4500                                0
  f1 rl-2               1000                                0
  f1 rl-3               1000                                0
  f1 rl-4                100                                0
  f1 rl-5               1500                                0
  L2 vpc-peer-gw        5000                0               0                 0
  L2 lisp-map-cache     5000                0               0                 0
  L2 dpss                100                0               0                 0
  L3 glean-fast          100                0               0                 0
```

3.3.1.6 VDCs and Resource Allocation

VDCs on the Nexus 7000 are used in the NVT testbed to partition a single physical device into multiple logical devices that provide fault isolation, management isolation, address allocation isolation, service differentiation domains, and adaptive resource management.

```
DC3-3a# sh vdc

Switchwide mode is m1 f1 m1xl f2 m2xl f2e f3

vdc_id  vdc_name                    state           mac                type       lc
------  --------                    -----           ----------         ---------  ------
1       DC3-3a                      active          64:a0:e7:3f:a4:c1  Admin      None
2       DC3-3                       active          64:a0:e7:3f:a4:c2  Ethernet   f2 f2e f3
3       DC3-1                       active          64:a0:e7:3f:a4:c3  Ethernet   m1 m1xl m2xl f2e
```

Resource allocation for VDC’s is done from the main VDC based on the requirements. The configuration used in the NVT testbed is as shown below.

The Following Command Can Be Used to Help Estimate the VDC Resource Allocation:

```
DC3-3a# show routing memory estimate routes 30000 next-hops 2
Shared memory estimates:
  Current max      8 MB;   2652 routes with 32 nhs
          in-use   1 MB;      9 routes with  1 nhs (average)
  Configured max   8 MB;   2652 routes with 32 nhs
  Estimate memory with fixed overhead:  11 MB;  30000 routes with  2 nhs
  Estimate with variable overhead included:
  - With MVPN enabled VRF:  12 MB
  - With OSPF route (PE-CE protocol):  15 MB
```

```
  - With EIGRP route (PE-CE protocol):  16 MB
```

Configuration:

```
vdc DC3-3a id 1
  cpu-share 5
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
vdc DC3-3 id 2
  limit-resource module-type f2 f2e f3
  cpu-share 5
  allocate interface Ethernet7/1-48
  allocate interface Ethernet8/1-48
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
vdc DC3-1 id 3
  limit-resource module-type m1 m1xl m2xl f2e
  cpu-share 5
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12
```

3.3.2 Routing Design Overview

3.3.2.1 Unicast Routing Design

3.3.2.1.1 BGP Routing Design

BGP has been chosen as the routing protocol for the NVT DC3 core, spine and leaf layers as shown in Figure 30. The layers are logically connected to each other through eBGP. The N7K core layer in BGP AS 3 is shared between other DC3 networks (DC31, DC32, DC33, and DC36).

Figure 30 BGP Core Logical Design

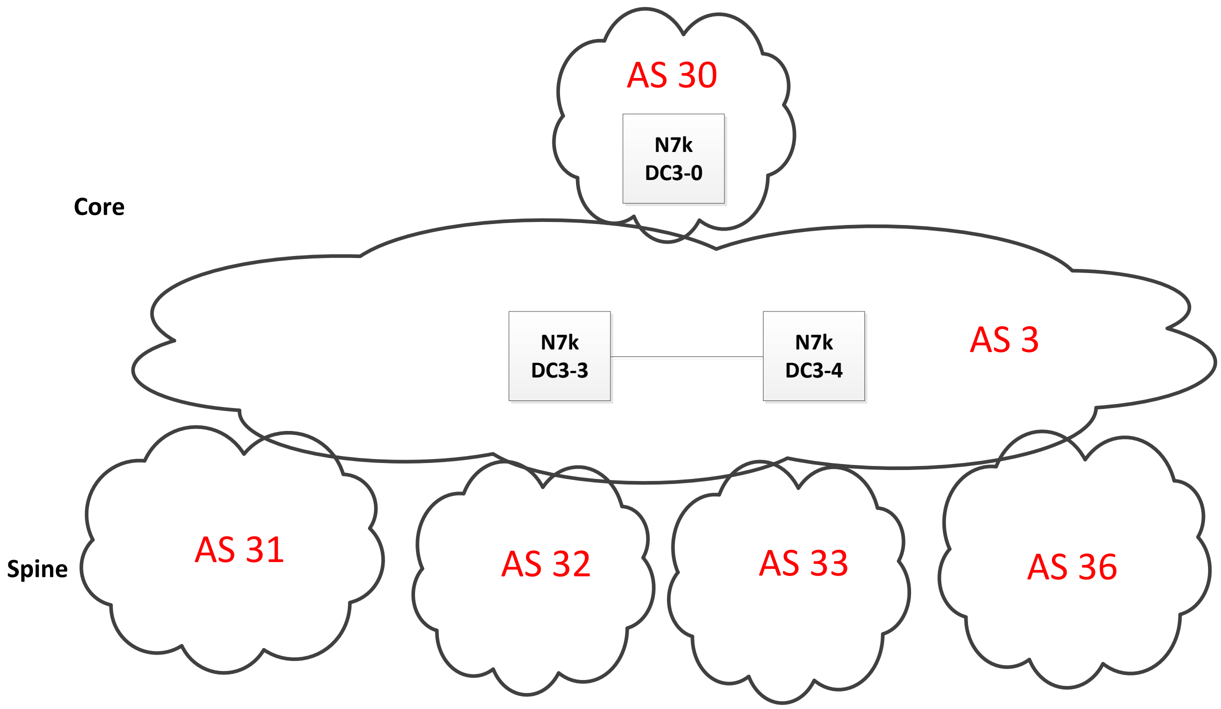

BGP Core Configuration:

```
router bgp 3
  router-id 40.33.0.15
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
    network 40.3.0.15/32
    network 40.3.4.0/24
    network 40.3.254.1/32
    network 40.30.31.0/24
    network 40.32.21.0/24
    network 40.32.22.0/24
    network 40.32.23.0/24
    network 40.32.24.0/24
    network 40.33.31.0/24
    network 40.33.32.0/24
    network 40.33.33.0/24
```

```
    network 40.33.34.0/24
    network 40.34.11.0/24
    network 40.34.12.0/24
    network 40.36.31.0/24
    network 40.36.32.0/24
    network 40.36.33.0/24
    network 40.36.34.0/24
    network 40.36.35.0/24
    network 40.36.36.0/24
    maximum-paths 32
  address-family ipv6 unicast
    network 2001:1:40:33:31::/80
    network 2001:1:40:33:32::/80
    network 2001:1:40:33:33::/80
    network 2001:1:40:33:34::/80
    network 2001:1:40:34:11::/80
    network 2001:1:40:34:12::/80
    network 2001:1:40:36:31::/80
    network 2001:1:40:36:32::/80
    network 2001:1:40:36:33::/80
    network 2001:1:40:36:34::/80
    network 2001:1:40:36:35::/80
    network 2001:1:40:36:36::/80
    network 2001:1:40:3::15:0:15/128
    network 2001:1:40:3:4::/80
    network 2001:40:30:31::/64
    maximum-paths 32
  template peer DC31
    remote-as 31
    address-family ipv4 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer DC32
    remote-as 32
    address-family ipv4 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer DC33
    remote-as 33
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
```

```
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      soft-reconfiguration inbound
  template peer DC34
    remote-as 34
    address-family ipv4 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer DC36
    remote-as 36
    address-family ipv4 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map NO-DEFAULT in
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  neighbor 40.3.4.17 remote-as 3
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  neighbor 40.30.31.10 remote-as 30
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  neighbor 40.31.11.1
    inherit peer DC31
  neighbor 40.31.12.2
    inherit peer DC31
  neighbor 40.32.21.1
    inherit peer DC32
  neighbor 40.32.22.2
    inherit peer DC32
  neighbor 40.32.23.3
    inherit peer DC32
  neighbor 40.32.24.4
```

```
    inherit peer DC32
  neighbor 40.33.31.1
    inherit peer DC33
  neighbor 40.33.32.2
    inherit peer DC33
  neighbor 40.33.33.3
    inherit peer DC33
  neighbor 40.33.34.4
    inherit peer DC33
  neighbor 40.34.11.1
    inherit peer DC34
  neighbor 40.34.12.2
    inherit peer DC34
  neighbor 40.36.31.1
    inherit peer DC36
  neighbor 40.36.32.2
    inherit peer DC36
  neighbor 40.36.33.3
    inherit peer DC36
  neighbor 40.36.34.4
    inherit peer DC36
  neighbor 40.36.35.5
    inherit peer DC36
  neighbor 40.36.36.6
    inherit peer DC36
```

BGP DC3-0 Configuration:

```
feature bgp

router bgp 30
  address-family ipv4 unicast
    network 40.3.0.10/32
    network 40.30.1.0/24
    network 40.30.2.0/24
    network 40.30.3.0/24
    network 40.30.4.0/24
    network 40.30.5.0/24
    network 40.30.6.0/24
    network 40.30.7.0/24
    network 40.30.8.0/24
    network 40.30.31.0/24
    network 40.30.41.0/24
    maximum-paths 16
  address-family ipv6 unicast
    network 2001:1:40:3::10:0:10/128
    network 2001:40:30:1::/64
    network 2001:40:30:2::/64
    network 2001:40:30:31::/64
    network 2001:40:30:3::/64
    network 2001:40:30:41::/64
    network 2001:40:30:4::/64
    network 2001:40:30:5::/64
    network 2001:40:30:6::/64
    network 2001:40:30:7::/64
    network 2001:40:30:8::/64
    maximum-paths 16
  neighbor 40.30.31.15 remote-as 3
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  neighbor 40.30.41.17 remote-as 3
    address-family ipv4 unicast
```

```
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
```

3.3.2.2 Multicast Routing Design

Multicast routing has been enabled across DC3 network.

DC3 core Multicast Configuration:

```
version 6.2(6)
feature pim

ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen

interface port-channel11
  ip pim sparse-mode
  ip pim border

interface loopback1
  description dc3-RP
  ip address 40.3.254.1/32
  ip pim sparse-mode
```

```
version 6.2(6)
feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.3.0.17 connect-source loopback0 remote-as 3

interface loopback0
  ip address 40.3.0.15/32
  ipv6 address 2001:1:40:3:0:15:0:15/128
  ip pim sparse-mode
```

3.3.2.2.1 PIM-ASM Rendezvous Point

PIM Sparse Mode has been configured as the protocol of choice for multicast routing on DC3 core.

3.3.2.2.1.1 Static RP

For the groups with a Rendezvous Point on the core, the RP is statically configured across DC3 network.

To Verify PIM RP:

```
DC3-3# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: unknown
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1*, (0), uptime: 5d14h, expires: never,
  priority: 0, RP-source: (local), group ranges:
```

```
      230.3.0.0/16
DC3-3# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Action    Mode      RP-address       Shared-tree-only range
232.0.0.0/8        Accept    SSM       -                -
230.3.0.0/16       -         ASM       40.3.254.1       --
```

3.3.2.2.1.2 Anycast RP with MSDP

NVT has configured Anycast RP with MSDP on the two core routers.

NVT Anycast RP and MSDP Configuration:

N7K core 1:	N7K core 2:

!Anycast RP configuration
ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  description dc3-RP
  ip address 40.3.254.1/32
  ip pim sparse-mode
 	
! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.3.0.17 connect-source loopback0
interface loopback0
  ip address 40.3.0.15/32
  ipv6 address 2001:1:40:3:0:15:0:15/128
  ip pim sparse-mode
	!Anycast RP configuration
	ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
	ip pim send-rp-discovery loopback1
	interface loopback1
	  description dc3-RP
	  ip address 40.3.254.1/32
	  ip pim sparse-mode
	
	! MSDP configuration
	ip msdp originator-id loopback0
	ip msdp peer 40.3.0.15 connect-source loopback0
	interface loopback0
	  ip address 40.3.0.17/32
	  ipv6 address 2001:1:40:3:0:17:0:17/128
	  ip pim sparse-mode

To Verify MSDP peer:

```
DC3-4# sh ip msdp sum
MSDP Peer Status Summary for VRF "default"
Local ASN: 3, originator-id: 40.3.0.17

Number of configured peers:  3
Number of established peers: 2
Number of shutdown peers:    0

Peer            Peer        Connection      Uptime/   Last msg  (S,G)s
Address         ASN         State           Downtime  Received  Received
40.3.0.15       3           Established     5d19h     00:00:01  0
```

3.3.2.2.2 PIM Border

On DC3 core network, PIM border is configured on all interfaces that connect to spine layer to prevent candidate-RP and Auto-RP messages from being sent or received to each POD.

Configuration:

```
DC3-4# sh run int p41

!Command: show running-config interface port-channel41
!Time: Fri Feb 21 15:17:25 2014

version 6.2(6)

interface port-channel41
```

```
  description DC33-1 - Po41 - Po4
  mtu 9216
  ip address 40.33.41.17/24
  ipv6 address 2001:1:40:33:41:17:0:17/80
  ip pim sparse-mode
  ip pim border

DC3-4# sh ip pim interface  po41
PIM Interface Status for VRF "default"
port-channel41, Interface status: protocol-up/link-up/admin-up
  IP address: 40.33.41.17, IP subnet: 40.33.41.0/24
  PIM DR: 40.33.41.17, DR's priority: 1
  PIM neighbor count: 1
  PIM hello interval: 30 secs, next hello sent in: 00:00:11
  PIM neighbor holdtime: 105 secs
  PIM configured DR priority: 1
  PIM border interface: yes
  PIM GenID sent in Hellos: 0x1a0f2019
  PIM Hello MD5-AH Authentication: disabled
  PIM Neighbor policy: none configured
  PIM Join-Prune inbound policy: none configured
  PIM Join-Prune outbound policy: none configured
  PIM Join-Prune interval: 1 minutes
  PIM Join-Prune next sending: 1 minutes
  PIM BFD enabled: no
  PIM passive interface: no
  PIM VPC SVI: no
  PIM Interface Statistics, last reset: never
    General (sent/received):
      Hellos: 17729/17736 (early: 0), JPs: 2/13, Asserts: 0/0
      Grafts: 0/0, Graft-Acks: 0/0
      DF-Offers: 0/0, DF-Winners: 0/0, DF-Backoffs: 0/0, DF-Passes: 0/0
    Errors:
      Checksum errors: 0, Invalid packet types/DF subtypes: 0/0
      Authentication failed: 0
      Packet length errors: 0, Bad version packets: 0, Packets from self: 0
      Packets from non-neighbors: 0
          Packets received on passiveinterface: 0
      JPs received on RPF-interface: 0
      (*,G) Joins received with no/wrong RP: 0/0
      (*,G)/(S,G) JPs received for SSM/Bidir groups: 0/0
      JPs filtered by inbound policy: 0
      JPs filtered by outbound policy: 0
```

3.4 DC31

3.4.1 Configuration of Platform Specific Features On DC31

3.4.1.1 Licensing

License Usage for Nexus 6000 in DC31:

```
dc31-1# show license usage
Feature                      Ins  Lic    Status Expiry Date Comments
                                  Count
--------------------------------------------------------------------------------
FCOE_NPV_PKG                  No    -   Unused             -
FM_SERVER_PKG                 No    -   Unused             -
ENTERPRISE_PKG                No    -   Unused             -
FC_FEATURES_PKG               No    -   Unused             -
VMFEX_FEATURE_PKG             No    -   Unused             -
ENHANCED_LAYER2_PKG           No    -   Unused             -
LAN_BASE_SERVICES_PKG         Yes   -   In use Never       -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   In use Never       -
--------------------------------------------------------------------------------
```

3.4.1.2 Out-of-Band Management Network

DC31 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic.

Configuration:

```
interface mgmt0
  description mgmt0==Gig2/18 mgmt-sw1
  vrf member management
  ip address 10.2.31.1/16
```

3.4.1.3 Common Configurations

3.4.1.3.1 SSH and TACACS+

SSH is enabled in DC31 to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration and Verification:

```
feature tacacs+


ip tacacs source-interface mgmt0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
    use-vrf management


dc31-1# show ssh server
ssh version 2 is enabled
dc31-1# show users
NAME     LINE         TIME         IDLE          PID COMMENT
admin    ttyS0        Feb 10 11:37  old         8287
interop  pts/0        Feb 10 12:18 00:02       12264 (taro.interop.cisco.com) session=ssh
interop  pts/1        Feb 13 16:16   .         29164 (taro.interop.cisco.com) session=ssh *
```

3.4.1.3.2 CDP and LLDP

CDP and LLDP are pervasively used on the DC31 testbed for inter-device discovery.

Configuration and Verification:

```
dc31-1# sh run cdp all

!Command: show running-config cdp all
!Time: Tue Feb 25 16:20:34 2014

version 6.0(2)N2(4)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
cdp format device-id system-name


interface mgmt0
  cdp enable

interface Ethernet1/1
  cdp enable

interface Ethernet1/2
  cdp enable

interface Ethernet1/3
  cdp enable
dc31-1# sh run lldp

!Command: show running-config lldp
!Time: Tue Feb 25 16:20:48 2014

version 6.0(2)N2(4)
feature lldp


dc31-1# sh cdp ne
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute


Device-ID             Local Intrfce Hldtme Capability  Platform      Port ID
mgmt-sw1.interop.cisco.com
                    mgmt0          167    R S I     WS-C6509-E    Gig2/18
dc31-1# sh lldp ne
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID             Local Intf      Hold-time  Capability  Port ID
dc31-103.interop.cisco.com Eth1/4/1        120        BR          Eth1/1
```

3.4.1.3.3 Syslog

Syslog is used to record all network events on the DC31 test bed. Whenever possible, DC31 makes use of a separate management VRF for syslog.

Configuration and Verification:

```
logging server syslog.interop.cisco.com 5 use-vrf management facility local6

dc31-1# sh logg ser
Logging server:                enabled
{syslog.interop.cisco.com}
```

```
        server severity:        notifications
        server facility:        local6
        server VRF:             management
```

3.4.1.3.4 SNMP

SNMP is used for system monitoring in DC31. Scripts are used to poll the systems asynchronously during the course of all DC31 test execution.

Configuration:

```
snmp-server source-interface trap mgmt0
snmp-server user admin network-admin auth md5 0x624c5ee49e857e9665cfb4ec04d3920a priv
0x624c5ee49e857e9665cfb4ec04d3920a localiz
edkey
snmp-server host 172.28.92.62 traps version 2c public
snmp-server enable traps callhome event-notify
snmp-server enable traps callhome smtp-send-fail
snmp-server enable traps cfs state-change-notif
snmp-server enable traps lldp lldpRemTablesChange
snmp-server enable traps cfs merge-failure
snmp-server enable traps aaa server-state-change
snmp-server enable traps upgrade UpgradeOpNotifyOnCompletion
snmp-server enable traps upgrade UpgradeJobStatusNotify
snmp-server enable traps feature-control FeatureOpStatusChange
snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended
snmp-server enable traps config ccmCLIRunningConfigChanged
snmp-server enable traps snmp authentication
snmp-server enable traps link cisco-xcvr-mon-status-chg
snmp-server enable traps vtp notifs
snmp-server enable traps vtp vlancreate
snmp-server enable traps vtp vlandelete
snmp-server enable traps bridge newroot
snmp-server enable traps bridge topologychange
snmp-server enable traps stpx inconsistency
snmp-server enable traps stpx root-inconsistency
snmp-server enable traps stpx loop-inconsistency
snmp-server enable traps poe portonoff
snmp-server enable traps poe pwrusageon
snmp-server enable traps poe pwrusageoff
snmp-server enable traps poe police
snmp-server community public group network-operator
snmp-server community private group network-admin
snmp-server community cisco group network-operator
```

3.4.1.3.5 NTP

NTP is used to synchronize the clocks on all DC31 devices to provide consistent timestamps on all network logs and events.

Configuration and Verification:

```
ntp distribute
ntp server 172.28.92.1 use-vrf management
ntp commit


dc31-1# show ntp status
Distribution : Enabled
Last operational state: No session
dc31-1# show ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
```

```
    remote               local                 st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                 8   64      377   0.00069 management*172.28.92.1
0.0.0.0                8   64      377   0.00092 management
```

3.4.1.3.6 SPAN

SPAN has been enabled on DC31 switches to provide packet captures to assist in network debugging.

Configuration and Verification:

```
monitor session 1
  source interface port-channel11 both
  destination interface Ethernet1/50
  no shut

dc31-1# sh monitor session 1
   session 1
---------------
type              : local
state             : up
acl-name          : acl-name not specified
source intf       :
    rx            : Po11
    tx            : Po11
    both          : Po11
source VLANs      :
    rx            :
destination ports : Eth1/50

Legend: f = forwarding enabled, l = learning enabled
```

3.4.1.3.7 DNS

DNS has been enabled to provide name lookup in DC31 network.

Configuration and Verification:

```
vrf context management
  ip domain-name interop.cisco.com
  ip domain-list interop.cisco.com
  ip domain-list cisco.com
  ip name-server 172.28.92.9 172.28.92.10

dc31-1# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=0.621 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=0.529 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=0.545 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=0.527 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=0.499 ms

--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.499/0.544/0.621 ms
```

3.4.1.3.8 MTU

System MTU is configured as jumbo MTU across the DC31 test bed. When the system MTU is set to greater than 9192, the message *"1%KERN-3-SYSTEM_MSG: packet sendmsg: packet size 9250 > MTU*

9230" could be seen. The internal header on the Nexus 6000 is 24 bytes hence why the MTU can become greater than 9216. In DC31, MTU size of 9000 is configured.

Configuration:

```
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9000
system qos
  service-policy type network-qos jumbo

interface Ethernet1/1
  description DC31-2
  no switchport
  no negotiate auto
  mtu 9000
  channel-group 2 mode active

dc31-1# sh queuing interface ethernet 1/1
Ethernet1/1 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0       WRR            100

  RX Queuing
    qos-group 0
    q-size: 100160, HW MTU: 9000 (9000 configured)
    drop-type: drop, xon: 0, xoff: 0
    Statistics:
        Pkts received over the port             : 269
        Ucast pkts sent to the cross-bar        : 264
        Mcast pkts sent to the cross-bar        : 5
        Ucast pkts received from the cross-bar  : 0
        Pkts sent to the port                   : 0
        Pkts discarded on ingress               : 0
        Per-priority-pause status               : Rx (Inactive), Tx (Inactive)
```

3.4.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU. The DC31 testbed uses the default CoPP.

Configuration:

```
dc31-1# sh copp status
Last Config Operation: None
Last Config Operation Timestamp: None
Last Config Operation Status: None
Policy-map attached to the control-plane: copp-system-policy-default

dc31-1# show policy-map type control-plane name copp-system-policy-default

  policy-map type control-plane copp-system-policy-default
    class copp-system-class-igmp
      police cir 1024 kbps bc 65535 bytes
    class copp-system-class-pim-hello
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bridging
      police cir 20000 kbps bc 4800000 bytes
    class copp-system-class-arp
      police cir 1024 kbps bc 3600000 bytes
    class copp-system-class-dhcp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-mgmt
```

```
      police cir 12000 kbps bc 4800000 bytes
    class copp-system-class-lacp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-lldp
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-udld
      police cir 2048 kbps bc 4800000 bytes
    class copp-system-class-isis
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-msdp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-cdp
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-fip
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-bgp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-eigrp
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-exception
      police cir 64 kbps bc 4800000 bytes
    class copp-system-class-glean
      police cir 1024 kbps bc 4800000 bytes
    class copp-system-class-hsrp-vrrp
      police cir 1024 kbps bc 256000 bytes
    class copp-system-class-icmp-echo
      police cir 64 kbps bc 3600000 bytes
    class copp-system-class-ospf
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-bfd
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-pim-register
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-rip
      police cir 9600 kbps bc 4800000 bytes
    class copp-system-class-l3dest-miss
      police cir 64 kbps bc 16000 bytes
    class copp-system-class-mcast-miss
      police cir 256 kbps bc 3200000 bytes
    class copp-system-class-excp-ip-frag
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-same-if
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-excp-ttl
      police cir 64 kbps bc 3200000 bytes
    class copp-system-class-default
      police cir 512 kbps bc 6400000 bytes
    class copp-system-class-rpf-fail
      police cir 512 kbps bc 3200000 bytes
    class copp-system-class-mcast-last-hop
      police cir 512 kbps bc 3200000 bytes
```

3.4.2 Image Upgrade and Downgrade

DC31 makes use of “install all” to upgrade/downgrade software images whenever possible, but the upgrade will be disruptive as Layer 3 features have been configured on the Nexus 6000.

```
dc31-102# show install all status
This is the log of last installation.


Verifying image bootflash:/n6000-uk9-kickstart.6.0.2.N2.3.52.bin for boot variable "kickstart".
SUCCESS
```

```
Verifying image bootflash:/n6000-uk9.6.0.2.N2.3.52.bin for boot variable "system".
SUCCESS

Verifying image type.
SUCCESS

Extracting "system" version from image bootflash:/n6000-uk9.6.0.2.N2.3.52.bin.
SUCCESS

Extracting "kickstart" version from image bootflash:/n6000-uk9-kickstart.6.0.2.N2.3.52.bin.
SUCCESS

Extracting "bios" version from image bootflash:/n6000-uk9.6.0.2.N2.3.52.bin.
SUCCESS

Performing module support checks.
SUCCESS

Notifying services about system upgrade.
SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes      disruptive         reset  Non-disruptive install not supported if L3 was enabled
     2       yes      disruptive         reset  Non-disruptive install not supported if L3 was enabled



Images will be upgraded according to following table:
Module             Image         Running-Version             New-Version  Upg-Required
------  ----------------  ----------------------  ----------------------  ------------
     1            system             6.0(2)N2(3)             6.0(2)N2(4)           yes
     1         kickstart             6.0(2)N2(3)             6.0(2)N2(4)           yes
     1              bios      v1.5.0(12/29/2012)      v1.5.0(12/29/2012)            no
     1         power-seq                    v4.0                    v4.0            no
     1  fabric-power-seq                    v4.0                    v4.0            no
     2         power-seq                    v1.0                    v4.0           yes
     1   microcontroller                v1.2.0.5                v1.2.0.5            no


Switch will be reloaded for disruptive upgrade.

Install is in progress, please wait.

Performing runtime checks.
SUCCESS

Setting boot variables.
SUCCESS

Performing configuration copy.
SUCCESS

Module 2: Refreshing compact flash and upgrading bios/loader/bootrom/power-seq.
Warning: please do not remove or power off the module at this time.
Note: Power-seq upgrade needs a power-cycle to take into effect.
On success of power-seq upgrade, SWITCH OFF THE POWER to the system and then, power it up.
Note: Micro-controller upgrade needs a power-cycle to take into effect.
On success of micro-controller upgrade, SWITCH OFF THE POWER to the system and then, power it up.
SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

3.4.3 Routing Design Overview

3.4.3.1 Unicast Routing Design

3.4.3.1.1 BGP Routing Design

The network is split into three layers: core, spine, and leaf. The layers are logically connected to each other through eBGP, as shown in Figure 31. The N7K core layer in BGP AS 3 is shared with other DC3 networks (DC32, DC33, and DC36). The spine layer runs OSPF to provide inter-switch connectivity to support iBGP sessions. The leaf layer is divided into multiple BGP ASes. This BGP logical design is easier to configure, maintain and debug than full mesh ibgp, route reflector, or confe derations; the core can consolidate these as private ASes if there is a need to advertise to other BGP exchanges.

The spine layer is eBGP connected to the ASes configured at the Leaf layer over both IPv4 and IPv6 address families (eBGP dual stack). The spine routers also inject the default route down to the leaf ASes for both IPv4 and IPv6 address families (default-originate). ECMP is enabled on both IPv4 and IPv6 address families (maximum-path 64) across the DC31 network.

The leaf layer represents different top of rack topologies that can be deployed. AS 31101 employs two N6001 in a vPC topology, using HSRP for gateway redundancy for nodes. AS 31103 employs a routed top of rack with N6001. AS 31104 employs a routed Nexus 3548 ToR. AS 31105 employs a r outed Nexus 3048 ToR. AS 31106 is used as a test tool rather than network under test. The Nexus 7000 is divided into multiple VRFs, with each VRF representing an extra ToR in the network. The goal is to test increasing number of ToR supported by the spine layer.

Figure 31DC31 BGP Logical Design

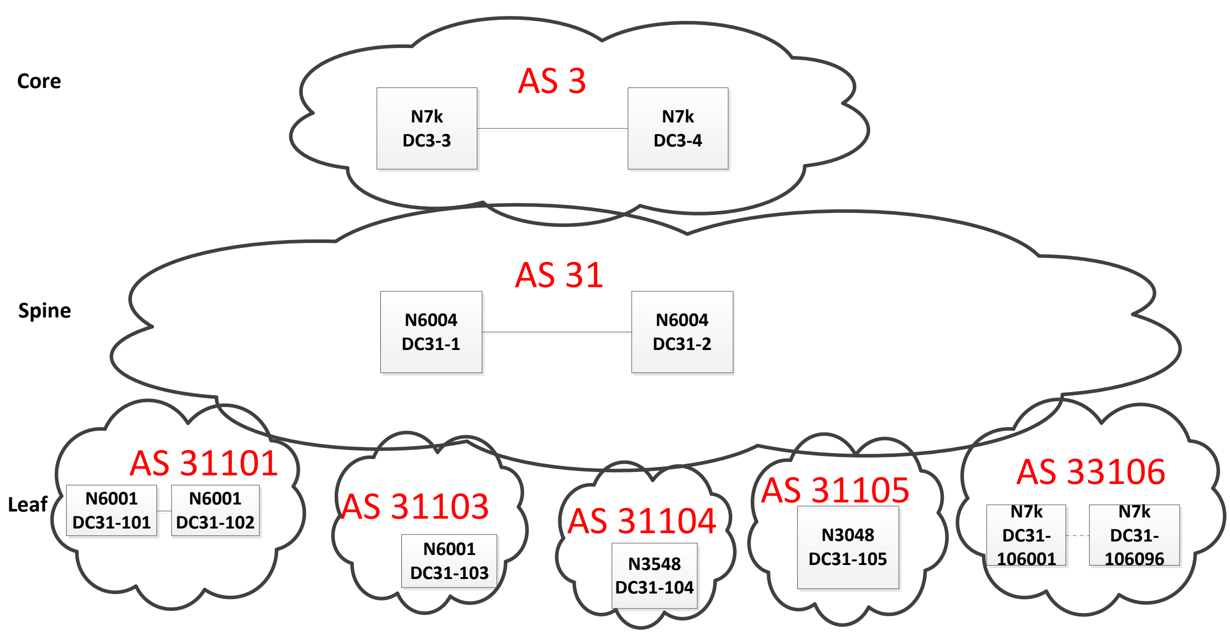

DC31 BGP configuration:

```
feature bgp

router bgp 31
  router-id 40.31.0.1
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
    network 31.101.11.0/24
…
    network 40.31.254.1/32
    maximum-paths 64
  address-family ipv6 unicast
    network 2001:1:40:31::1:0:1/128
…
    network 2001:31:106:196::/64
    maximum-paths 64
  neighbor 31.101.11.101 remote-as 31101
    inherit peer BGPLEAF
…
  neighbor 31.106.196.106 remote-as 31106
    inherit peer BGPLEAF106
    no shutdown
  neighbor 40.31.0.2 remote-as 31
    inherit peer BGPSPINE
  neighbor 40.31.11.15 remote-as 3
    inherit peer BGPCORE
  neighbor 40.31.13.17 remote-as 3
    inherit peer BGPCORE
  template peer BGPCORE
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  template peer BGPLEAF
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer BGPLEAF106
    address-family ipv4 unicast
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      route-map DEFAULT-ONLY out
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer BGPSPINE
    update-source loopback0
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
```

3.4.3.1.1.1 BGP Router-Id

To establish BGP sessions between peers, BGP must have a router ID, which is sent to BGP peers in the OPEN message when a BGP session is established. On DC31, NVT has configured a loopback interface IP address as the BGP router-ID. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represen t the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

To Verify the BGP Router-ID:

```
dc31-1# sh ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 37967, local router ID is 40.31.0.1
```

3.4.3.1.1.2 BGP Address Family

BGP address family for IPv4 and Ipv6 have been configured to achieve BGP peering, load -balancing, default route injection.

To Verify the BGP Address Family:

```
dc31-1# show ip bgp all summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 40.31.0.1, local AS number 31
BGP table version is 37967, IPv4 Unicast config peers 117, capable peers 117
1435 network entries and 8433 paths using 513844 bytes of memory
BGP attribute entries [14/1792], BGP AS path entries [6/36]
BGP community entries [0/0], BGP clusterlist entries [0/0]
8314 received paths for inbound soft reconfiguration
8314 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
31.101.11.101   4 31101    4707    4874    37967    0    0    3d05h 411
31.101.12.101   4 31101    4714    4854    37967    0    0    3d05h 411
31.102.11.102   4 31101    4750    6052    37967    0    0    2d23h 411
31.102.12.102   4 31101    4756    6033    37967    0    0    2d23h 411
31.103.101.103  4 31103    4665    4909    37967    0    0    3d05h 423
31.103.102.103  4 31103    4667    4909    37967    0    0    3d05h 423


BGP summary information for VRF default, address family IPv6 Unicast
BGP router identifier 40.31.0.1, local AS number 31
BGP table version is 34505, IPv6 Unicast config peers 117, capable peers 117
1136 network entries and 7538 paths using 461792 bytes of memory
BGP attribute entries [10/1280], BGP AS path entries [4/24]
BGP community entries [0/0], BGP clusterlist entries [0/0]
7422 received paths for inbound soft reconfiguration
7422 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
31.101.11.101   4 31101    4707    4874    34505    0    0    3d05h 411
31.101.12.101   4 31101    4714    4854    34505    0    0    3d05h 411
31.102.11.102   4 31101    4750    6052    34505    0    0    2d23h 411
31.102.12.102   4 31101    4756    6033    34505    0    0    2d23h 411
```

3.4.3.1.1.3 BGP Load Sharing and ECMP

DC31 has configured the maximum-paths that BGP adds to the route table for equal-cost multipath load balancing as 64 for both spine and leaf peers for IPv4/IPv6 address families.

3.4.3.1.1.4 BGP Authentication

DC31 has configured MD5 Authentication for BGP sessions.

To Verify the BGP Authentication:

```
dc31-1# sh ip bgp neighbors 31.101.11.101
BGP neighbor is 31.101.11.101,  remote AS 31101, ebgp link,  Peer index 4
  Inherits peer configuration from peer-template BGPLEAF
  BGP version 4, remote router ID 31.0.0.101
  BGP state = Established, up for 00:00:02
  Peer is directly attached, interface Ethernet1/2
  TCP MD5 authentication is enabled
```

3.4.3.1.1.5 BGP Update-Source

DC31 has configured BGP update-source to establish a BGP multi-hop sessions. DC31 has multi-hop sessions only on the iBGP peering between the spine switches.

To Verify the BGP Update-Source:

```
dc31-1# sh ip bgp neighbors 40.31.0.2
BGP neighbor is 40.31.0.2,  remote AS 31, ibgp link,  Peer index 1
  Inherits peer configuration from peer-template BGPSPINE
  BGP version 4, remote router ID 40.31.0.2
  BGP state = Established, up for 3d05h
  Using loopback0 as update source for this peer
```

3.4.3.1.1.6 BGP Default Route

BGP default route is advertised from the spine peers to the leaf peers for both Ipv4 and Ipv6 address families.

To Verify the BGP Default Route:

```
dc31-1# sh ip bgp neighbors 31.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 38782, neighbor version 38782
  411 accepted paths consume 19728 bytes of memory
  1032 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 31.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 3d05h after start

  For address family: IPv6 Unicast
  BGP table version 35320, neighbor version 35320
  411 accepted paths consume 19728 bytes of memory
  733 sent paths
  Inbound soft reconfiguration allowed
  Default information originate, default sent
  Last End-of-RIB received 3d05h after start
```

3.4.3.1.1.7 BGP Next-Hop-Self

BGP next-hop-self is configured for BGP sessions between the spine switches for both IPv4 and IPv6 address families. However, a cosmetic issue prevents the next-hop-self from showing up in the *show ip bgp neighbors* output (CSCun31570).

To Verify the BGP Next-Hop-Self:

```
dc31-1# sh ip bgp neighbors 31.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 38782, neighbor version 38782
  411 accepted paths consume 19728 bytes of memory
  1032 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 31.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 3d05h after start
```

3.4.3.1.1.8 BGP Soft-Reconfiguration

BGP Soft reset is recommended because it allows routing tables to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis.

```
dc31-1# sh ip bgp neighbors 31.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 38782, neighbor version 38782
  411 accepted paths consume 19728 bytes of memory
  1032 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 31.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 3d05h after start

  For address family: IPv6 Unicast
  BGP table version 35320, neighbor version 35320
  411 accepted paths consume 19728 bytes of memory
  733 sent paths
  Inbound soft reconfiguration allowed
  Default information originate, default sent
  Last End-of-RIB received 3d05h after start
```

3.4.3.1.2 OSPF Routing Design

OSPF/OSPFv3 is used as the IGP to provide reachability for establishing iBGP peering at the spine layer only. The OSPF/OSPFv3 process is enabled only on directly connected interfaces and the Loopback interface. All the OSPF enabled interfaces are in Area 0.0.0.0. Each OSPF network type is set to point-to-point to decrease OSPF neighbor setup latency. In order to improve OSPF convergence, SPF and LSA timers are throttled to (100 200 5000 and 50 100 300) respectively.

DC31 OSPF/OSPFv3 Configuration:

```
feature ospf
router ospf 1
```

```
  router-id 40.31.0.2
  log-adjacency-changes
  timers throttle spf 100 200 5000
  timers throttle lsa 50 100 300
  auto-cost reference-bandwidth 100000

interface loopback0
  ip router ospf 1 area 0.0.0.0

interface loopback1
  ip router ospf 1 area 0.0.0.0

interface port-channel1
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0



feature ospfv3
router ospfv3 31
  router-id 40.31.0.2
  auto-cost reference-bandwidth 100000

interface loopback0
  ipv6 router ospfv3 31 area 0.0.0.0

interface port-channel1
  ospfv3 network point-to-point
  ipv6 router ospfv3 31 area 0.0.0.0
```

3.4.3.1.3 Unicast Forwarding Verification

This Switch is the Authoritative Router for a Directly Connected Subnet on VLAN 11 131.10.11.0/24:

```
dc31-101# show running-config interface vlan 11

!Command: show running-config interface Vlan11
!Time: Fri Feb 14 14:11:27 2014

version 6.0(2)N2(3)

interface Vlan11
  no shutdown
  mtu 9000
  ip address 131.10.11.2/24
  ipv6 address 2001:131:10:11::2/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 50
    ip 131.10.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 50
    ip 2001:131:10:11::1
```

The Host 131.10.11.51 has been Learned via ARP on this Subnet:

```
dc31-101# sh ip arp 131.10.11.51

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
```

```
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
131.10.11.51    00:03:53  0083.0a0b.3300  Vlan11
```

On NX-OS, "show ip route" will also Show Directly Connected Hosts as /32 Routes:

```
dc31-101# show ip route 131.10.11.51
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

131.10.11.51/32, ubest/mbest: 1/0, attached
    *via 131.10.11.51, Vlan11, [250/0], 4w4d, am
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
dc31-101# show ip adjacency 131.10.11.51

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
131.10.11.51    0083.0a0b.3300  50   arp        Vlan11
```

Find the PO Interface on which this MAC Address is Learnt:

```
dc31-101# sh mac address-table address 0083.0a0b.3300
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0083.0a0b.3300    dynamic   20          F    F  Po11
```

Display PO11 Member Interface with Module Information:

```
dc31-101# sh port-channel summary | in Po11
11    Po11(SU)    Eth      LACP      Eth1/1(P)
```

Display Adjacency Index for this Route in Hardware Table:

```
dc31-101# sh system internal forwarding ip route 131.10.11.51

Routes for table default/base

----+----------------+--------------------------+---------------+-----------
Dev | Prefix         | UC/MC Handle (Index)     |AdjIdx(nhcount)|     LIF
----+----------------+--------------------------+---------------+-----------
1 131.10.11.51/32    0x23de2/0xdeadbeef             0x10000(0x1)  0x901000b
```

Display DMAC Entry Programmed in Adjacency Table:

```
dc31-101# sh system internal forwarding adjacency entry  0x10000 detail
Index 0x10000 MAC 002A.6A35.A8C1  BD 400
Prefix =     131.10.11.3/32 HANDLE = 0x20274
Index 0x10000 MAC 0083.0A0B.0B00  BD 400
Prefix =    131.10.11.11/32 HANDLE = 0x2356c
Index 0x10000 MAC 0083.0A0B.0B01  BD 400
Prefix =    131.10.11.12/32 HANDLE = 0x225c5
Index 0x10000 MAC 0083.0A0B.0B02  BD 400
Prefix =    131.10.11.13/32 HANDLE = 0x22326
Index 0x10000 MAC 0083.0A0B.0B03  BD 400
Prefix =    131.10.11.14/32 HANDLE = 0x22803
Index 0x10000 MAC 0083.0A0B.0B04  BD 400
Prefix =    131.10.11.15/32 HANDLE = 0x22ee0
Index 0x10000 MAC 0083.0A0B.1500  BD 400
Prefix =    131.10.11.21/32 HANDLE = 0x23ced
Index 0x10000 MAC 0083.0A0B.1501  BD 400
Prefix =    131.10.11.22/32 HANDLE = 0x237c8
Index 0x10000 MAC 0083.0A0B.1502  BD 400
Prefix =    131.10.11.23/32 HANDLE = 0x2312b
Index 0x10000 MAC 0083.0A0B.1503  BD 400
Prefix =    131.10.11.24/32 HANDLE = 0x2169a
Index 0x10000 MAC 0083.0A0B.1504  BD 400
Prefix =    131.10.11.25/32 HANDLE = 0x21079
Index 0x10000 MAC 0083.0A0B.3300  BD 400
Prefix =    131.10.11.51/32 HANDLE = 0x23de2
```

Display if Packets are Getting Dropped:

```
dc31-101# sh platform fwm info asic-errors all

Printing non zero Carmel error registers - 48 bits:
BIG_DROP_INGRESS_FW_PARSING_ERROR: res0 = 2 res1 = 0 [4]
BIG_DROP_INGRESBIG_BIG_DROP_S_INVALID_IF: res0 = 4 res1 = 0 [5]
BIG_DROP_INGRESS_UC_PC_DROP: res0 = 2 res1 = 0 [19]
BIG_DROP_CDCE_SW_TBL_RPF_MISS: res0 = 852571 res1 = 0 [49]
BIG_DROP_HIT_DROP_PORT_MAP_IDX: res0 = 2 res1 = 0 [53]
BIG_DROP_SRC_VLAN_MBR: res0 = 4 res1 = 0 [59]
BIG_DROP_EGRESS_ACL: res0 = 256752553 res1 = 0 [76]

Printing non zero Carmel error registers - 32 bits:
```

3.4.3.2 Multicast Routing Design

Multicast routing has been enabled across the entire DC31 network.

DC31 Multicast Configuration:

```
feature pim

ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
ip pim send-rp-announce loopback1 group-list 230.31.0.0/16
ip pim send-rp-discovery loopback1
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen

interface loopback1
  description dc31-RP
  ip address 40.31.254.1/32
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
```

```
feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.31.0.2 connect-source loopback0

interface loopback0
  ip address 40.31.0.1/32
  ipv6 address 2001:1:40:31:0:1:0:1/128
  ip router ospf 1 area 0.0.0.0
  ipv6 router ospfv3 31 area 0.0.0.0
  ip pim sparse-mode
```

3.4.3.2.1 PIM-ASM Rendezvous Point

PIM Sparse Mode has been configured as the protocol of choice for multicast routing. NX-OS does not support PIM SSM and PIM Bidir operating over vPC.

3.4.3.2.1.1 Auto-RP

The DC31 testbed is designed to have the RP on the spine to support the groups sourced from that particular POD. DC31 makes use of Auto-RP to automate distribution of RP information in the network.

To Verify PIM RP:

```
dc31-1# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.31.254.1*, next Discovery message in: 00:00:05
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1, (0), uptime: 3d23h, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.31.254.1*, (0), uptime: 3d23h, expires: 00:02:26,
  priority: 0, RP-source: 40.31.254.1 (A), group ranges:
      230.31.0.0/16
dc31-1# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.31.0.0/16      ASM       40.31.254.1      -
```

3.4.3.2.1.1.1 Auto-RP Forward Listen

DC31 has enabled the Auto-RP listening and forwarding feature so that the Auto-RP mechanism can dynamically inform routers in the PIM domain of the group-to-RP mapping since PIM dense mode is not supported on NX-OS. By default, listening or forwarding of Auto-RP messages is not enabled on NX-OS.

3.4.3.2.1.2 Static RP

For the groups with a Rendezvous Point on the core, the RP is statically configured on all routers in the DC31 network.

To Verify PIM RP:

```
dc31-101# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.31.254.1, uptime: 1d17h, expires: 00:02:48
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1, (0), uptime: 1d17h, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.31.254.1, (0), uptime: 1d17h, expires: 00:02:48,
  priority: 0, RP-source: 40.31.254.1 (A), group ranges:
      230.31.0.0/16

dc31-101# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.31.0.0/16      ASM       40.31.254.1      -
```

3.4.3.2.1.3 Anycast RP with MSDP

DC31 has configured Anycast RP with MSDP within each POD at the spine layer. DC31 has also configured Anycast RP with MSDP among the core switches.

DC31 Anycast RP and MSDP Configuration:

```
N6K spine 1:

!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list
230.31.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  description dc31-RP
  ip address 40.31.254.1/32
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.31.0.2 connect-source loopback0
interface loopback0
  ip address 40.31.0.1/32
  ipv6 address 2001:1:40:31:0:1:0:1/128
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
```

```
N6K spine 2:

!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list
230.31.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  description dc31-RP
  ip address 40.31.254.1/32
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.31.0.1 connect-source loopback0
interface loopback0
  ip address 40.31.0.2/32
  ipv6 address 2001:1:40:31:0:1:0:2/128
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
```

To Verify MSDP Peer and SA_Cache:

```
dc31-1# sh ip msdp sa-cache
MSDP SA Route Cache for VRF "default" - 1767 entries
Source          Group           RP              ASN         Uptime
131.10.11.12    230.31.0.1      40.31.0.2       0           00:05:57
131.10.11.21    230.31.0.1      40.31.0.2       0           00:05:57
131.10.11.35    230.31.0.1      40.31.0.2       0           00:07:00
131.10.12.34    230.31.0.1      40.31.0.2       0           00:07:00
131.10.13.15    230.31.0.1      40.31.0.2       0           00:05:57
131.10.14.15    230.31.0.1      40.31.0.2       0           00:05:57
```

```
…

dc31-1# sh ip msdp summary
MSDP Peer Status Summary for VRF "default"
Local ASN: 31, originator-id: 40.31.0.1

Number of configured peers:  1
Number of established peers: 1
Number of shutdown peers:    0

Peer            Peer       Connection      Uptime/   Last msg  (S,G)s
Address         ASN        State           Downtime  Received  Received
40.31.0.2       0          Established     1d19h     00:00:53  1767
```

3.4.3.2.2 PIM SPT-Threshold

DC31 has enabled *ip pim spt-threshold infinity* on the last hop non-vPC PIM routers to decrease the multicast entries hardware usage across the network. However, on the Nexus 6000 it was found that the spt-threshold infinity config caused the cessation of (S,G) state creation on the first-hop router thereby leading to the dropping of all incoming traffic (CSCul56319). Following this finding the spt-threshold config was removed from all Nexus 6000 switches.

3.4.3.2.3 Multicast Multipath

Cisco NX-OS Multicast Multipath is enabled by default and the load sharing selection algorithm is based on the source and group addresses.

3.4.3.2.4 Multicast Forwarding Verification

The following sequence of commands illustrates the verification of the Cisco NX-OS multicast L2 and L3 forwarding.

Displays a Specific Multicast Route 230.31.0.1 with Incoming Interface Information:

```
dc31-102# sh ip mroute 230.31.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.31.0.1/32), uptime: 00:31:21, igmp pim ip
  Incoming interface: Ethernet2/3, RPF nbr: 31.102.12.1
  Outgoing interface list: (count: 10)
    Vlan20, uptime: 00:31:21, igmp
    Vlan17, uptime: 00:31:21, igmp
    Vlan14, uptime: 00:31:21, igmp
    Vlan19, uptime: 00:31:21, igmp
    Vlan16, uptime: 00:31:21, igmp
    Vlan13, uptime: 00:31:21, igmp
    Vlan18, uptime: 00:31:21, igmp
    Vlan15, uptime: 00:31:21, igmp
    Vlan12, uptime: 00:31:21, igmp
    Vlan11, uptime: 00:31:21, igmp


(131.30.11.11/32, 230.31.0.1/32), uptime: 00:25:51, ip pim mrib
  Incoming interface: Ethernet2/1, RPF nbr: 31.102.11.1
  Outgoing interface list: (count: 10)
    Vlan11, uptime: 00:25:49, mrib
    Vlan12, uptime: 00:25:49, mrib
    Vlan13, uptime: 00:25:49, mrib
    Vlan14, uptime: 00:25:49, mrib
```

```
    Vlan15, uptime: 00:25:49, mrib
    Vlan16, uptime: 00:25:49, mrib
    Vlan17, uptime: 00:25:49, mrib
    Vlan18, uptime: 00:25:49, mrib
    Vlan19, uptime: 00:25:49, mrib
    Vlan20, uptime: 00:25:49, mrib
```

Display DR Information for Interface Vlan11:

```
dc31-102# sh ip pim interface brief
PIM Interface Status for VRF "default"
Interface            IP Address      PIM DR Address  Neighbor  Border
                                                     Count     Interface
Vlan20               131.10.20.3     131.10.20.3     1         no
Vlan19               131.10.19.3     131.10.19.3     1         no
Vlan18               131.10.18.3     131.10.18.3     1         no
Vlan17               131.10.17.3     131.10.17.3     1         no
Vlan16               131.10.16.3     131.10.16.3     1         no
Vlan15               131.10.15.3     131.10.15.3     1         no
Vlan14               131.10.14.3     131.10.14.3     1         no
Vlan13               131.10.13.3     131.10.13.3     1         no
Vlan12               131.10.12.3     131.10.12.3     1         no
Vlan11               131.10.11.3     131.10.11.3     1         no
Vlan1                131.10.1.3      131.10.1.3      1         no
Ethernet2/1          31.102.11.102   31.102.11.102   1         no
Ethernet2/2          31.102.21.102   31.102.21.102   1         no
Ethernet2/3          31.102.12.102   31.102.12.102   1         no
Ethernet2/4          31.102.22.102   31.102.22.102   1         no
```

Displays Mroute RPF Interface and Forwarding Counters in L3 Hardware Table:

```
dc31-102# sh forwarding multicast route group 230.31.0.1 source 131.30.11.11

  (131.30.11.11/32, 230.31.0.1/32), RPF Interface: Ethernet2/1, flags:
    Received Packets: 134 Bytes: 8710
    Number of Outgoing Interfaces: 10
    Outgoing Interface List Index: 10
      Vlan11 Outgoing Packets:0 Bytes:0
      Vlan12 Outgoing Packets:0 Bytes:0
      Vlan13 Outgoing Packets:0 Bytes:0
      Vlan14 Outgoing Packets:0 Bytes:0
      Vlan15 Outgoing Packets:0 Bytes:0
      Vlan16 Outgoing Packets:0 Bytes:0
      Vlan17 Outgoing Packets:0 Bytes:0
      Vlan18 Outgoing Packets:0 Bytes:0
      Vlan19 Outgoing Packets:0 Bytes:0
      Vlan20 Outgoing Packets:0 Bytes:0
```

Displays the Multicast Routing Table with Packet Counts and Bit Rates for All Sources:

```
dc31-102# show ip mroute 230.31.0.1 summary
IP Multicast Routing Table for VRF "default"

Total number of routes: 3111
Total number of (*,G) routes: 10
Total number of (S,G) routes: 3100
Total number of (*,G-prefix) routes: 1
Group count: 10, rough average sources per group: 310.0

Group: 230.31.0.1/32, Source count: 310
Source          packets      bytes           aps   pps       bit-rate     oifs
(*,G)           312          365136          1170  0         0.000   bps  10
131.30.11.11    217          13250           61    0         27.200  bps  10
```

Display IGMP Snooping Groups Information:

```
dc31-102# show ip igmp snooping groups 230.31.0.1 vlan 11
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
11    230.31.0.1         v2   D     Po101 Eth1/47 Po11
```

Displays Detected Multicast Routers for VLAN:

```
dc31-102# show ip igmp snooping mrouter vlan 11
Type: S - Static, D - Dynamic, V - vPC Peer Link,       I - Internal, F - FabricPath core port
      C - Co-learned, U - User Configured
Vlan  Router-port   Type      Uptime      Expires
11    Po101         SVD       1d20h       00:04:41
11    Vlan11        I         1d20h       never
```

Displays IGMP Snooping Querier Information for VLAN:

```
dc31-102# show ip igmp snooping querier vlan 11
Vlan  IP Address       Version   Expires     Port
11    131.10.11.2      v2        00:03:31    port-channel101
```

Display L3 FIB Entries:

```
dc31-102# show system internal forwarding ip multicast route  group 230.31.0.1 source 131.30.11.11 detail
Hardware Multicast FIB Entries:
 Flags Legend:
  * - no_dc_sup_redir
  S - sg_entry
  D - Non-RPF Drop
  B - Bi-dir route
  3 - RPF is L3lif
 W - Wildcard route
  A - Alt. Route exists
  R - RPF PTR or RP BD
  U - PD Route
 MET ENTRY FLAGS:
  O - Bridge Only
  P - Bridge Primary
  F - Ftag Hash Sel
  2 - L2_Update
  3 - L3_update
  I - inh_sg_from_starg
 IG - inh_from_gm

  ROUTERG ASIC_LIST: 4
(131.30.11.11/32, 230.31.0.1/32), Flags: *3
  Bigsur: 1, VPN: 1, MCTAG: 410, ALT MCTAG: 411  RPF Interface: Ethernet2/1 rpf_bd: 423 S Index: 0x23f0c G
Index: 0x29424 SG Index: 0x105c
  mtu_idx : 1, mc_port_mode : MC_PORT_MODE_10G, fabric_mc_en : 0, sup_copy : 0, rpf_fail_send_to_sup : 1
  MC PTR: 1038,   MET PTR: 1149,   MC PORT MODE: MC_PORT_MODE_10G
  MET Entries:
    Flags: 3I  BD: 391,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 392,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 393,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 394,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 395,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 396,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 397,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 398,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
```

```
    Flags: 3I  BD: 399,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3I  BD: 400,  MET TYPE: Svi  MCIDX: 345,  FTAG_BASE: 256,  NUM_FTAGS: 1,  ASIC_LIST: 1,4
    Flags: 3  BD: 421,  MET TYPE: VpnVlan  MCIDX: 11,  FTAG_BASE: 1023,  NUM_FTAGS: 1,  ASIC_LIST: 4
```

3.4.4 Layer-2/ Layer-3 Leaf/Access Layer Network Design Overview

3.4.4.1 vPC

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco NX-OS switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

vPC Peer Configurations:

```
N6K 1:
feature vpc

! vpc domain config
vpc domain 101
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  delay restore 150
  auto-recovery
  ip arp synchronize


! vpc peer-link config
interface port-channel102
  switchport mode trunk
  switchport trunk allowed vlan 1,11-410
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/42
  description Eth1/42==Eth1/42 dc31-102
  switchport mode trunk
  switchport trunk allowed vlan 1,11-410
  channel-group 102 mode active

! vpc peer-keepalive config
interface Ethernet1/41
  description Eth1/41==Eth1/41 dc31-102
  no switchport
  vrf member vpc-keepalive
  ip address 1.1.1.1/24

! vpc member port-channel config
interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  spanning-tree port type edge trunk
  vpc 11

! vpc member port config
interface Ethernet1/1
  description Eth1/1==Eth7/1 dc31-1001
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  channel-group 11 mode active

! PIM prebuild SPT
ip pim pre-build-spt
```

```
N6K 2:
feature vpc

! vpc domain config
vpc domain 101
  role priority 201
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  delay restore 150
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel101
  switchport mode trunk
  switchport trunk allowed vlan 1,11-410
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/42
  description Eth1/42==Eth1/42 dc31-101
  switchport mode trunk
  switchport trunk allowed vlan 1,11-410
  channel-group 101 mode active

! vpc peer-keepalive config
interface Ethernet1/41
  description Eth1/41==Eth1/41 dc31-101
  no switchport
  vrf member vpc-keepalive
  ip address 1.1.1.2/24

! vpc member port-channel config
interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  spanning-tree port type edge trunk
  vpc 11

! vpc member port config
interface Ethernet1/1
  description Eth1/1==Eth8/1 dc31-1001
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  channel-group 11 mode active

! PIM prebuild SPT
ip pim pre-build-spt
```

Display vPC Status:

```
dc31-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 40
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po101  up     1,11-410

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
12     Po12        up     success     success                    21-30
13     Po13        up     success     success                    31-40
14     Po14        up     success     success                    41-50
15     Po15        up     success     success                    51-60
16     Po16        up     success     success                    61-70
17     Po17        up     success     success                    71-80
18     Po18        up     success     success                    81-90
19     Po19        up     success     success                    91-100
```

3.4.4.1.1 LACP

DC31 makes use of LACP mode active for all link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
dc31-102# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
11    Po11(SU)    Eth      LACP      Eth1/1(P)
12    Po12(SU)    Eth      LACP      Eth1/2(P)
13    Po13(SU)    Eth      LACP      Eth1/3(P)
14    Po14(SU)    Eth      LACP      Eth1/4(P)
15    Po15(SU)    Eth      LACP      Eth1/5(P)
16    Po16(SU)    Eth      LACP      Eth1/6(P)
17    Po17(SU)    Eth      LACP      Eth1/7(P)
18    Po18(SU)    Eth      LACP      Eth1/8(P)
19    Po19(SU)    Eth      LACP      Eth1/9(P)
```

```
20    Po20(SU)    Eth      LACP      Eth1/10(P)
dc31-102# show lacp interface e1/1
Interface Ethernet1/1 is up
  Channel group is 11 port channel is Po11
  PDUs sent: 11195
  PDUs rcvd: 11192
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(7f9b, 0-23-4-ee-be-65, 800b, 8000, 101), (8000, 40-55-39-3-e3-42, a, 8000, 801)] ]
Operational as aggregated link since Sat Feb 15 19:45:17 2014

Local Port: Eth1/1   MAC Address= 0-2a-6a-35-a8-c1
  System Identifier=0x8000,  Port Identifier=0x8000,0x101
  Operational key=32779
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=(Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Neighbor: 0x801
  MAC Address= 40-55-39-3-e3-42
  System Identifier=0x8000,  Port Identifier=0x8000,0x801
  Operational key=10
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=(Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Aggregate or Individual(True=1)= 1
```

3.4.4.1.2 VLAN Trunking

DC31 makes use of VLAN trunking to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Display VLAN Information for Nexus 6000:

```
dc31-102# show vlan internal usage

VLANs                  DESCRIPTION
-------------------    -----------------
3968-4031              Multicast
4032-4035              Online Diagnostic
4036-4039              ERSPAN
4042                   Satellite
3968-4047,4094         Current
3000                   VPC bind-vrf
dc31-102# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po11, Po101, Eth1/1, Eth1/42
                                                Eth1/43, Eth1/44, Eth1/45
                                                Eth1/46, Eth1/47, Eth1/48

VLAN Type  Vlan-mode
---- ----- ----------
```

```
11   enet  CE

Primary  Secondary  Type            Ports
-------  ---------  --------------- -------------------------------------------
```

3.4.4.1.3 Spanning Tree

vPC technology helps build a loop free topology by leveraging port-channels from access devices to the vPC domain. A port-channel is seen as a logical link from the spanning tree's standpoint, so a vPC domain with vPC-attached access devices forms a star topology at Layer 2 (there are no STP blocked ports in this type of topology). In this case, STP is used as a fail-safe mechanism to protect against any network loops.

DC31 makes use of Rapid-PVST which is the default spanning tree protocol on NX-OS. For networks with larger logical port counts, MST is recommended.

Display Spanning Tree Information:

```
dc31-102# show spanning-tree vlan 11

VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    8203
             Address     002a.6a35.a8c1
             This bridge is the root
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8203   (priority 8192 sys-id-ext 11)
             Address     002a.6a35.a8c1
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po11             Desg FWD 1         128.4106 (vPC) P2p
Po101            Desg FWD 1         128.4196 (vPC peer-link) Network P2p
Eth1/46          Desg FWD 2         128.174  Edge P2p
Eth1/47          Desg FWD 2         128.175  Edge P2p
Eth1/48          Desg FWD 2         128.176  Edge P2p


dc31-102# show spanning-tree summary totals
Switch is in rapid-pvst mode
Root bridge for: VLAN0001, VLAN0011-VLAN0410
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short
STP-Lite                                 is enabled

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
401 vlans                     0         0        0       2001       2001
```

3.4.4.1.4 Configuration Parameters Consistency

After the vPC feature is enabled and the vPC peer-link on both peer devices is configured, Cisco Fabric Services messages provide a copy of the local vPC peer device configuration to the remote vPC peer device. The systems then determine whether any of the crucial configuration parameters differ on the two devices.

When a Type 1 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 1 consistency check failure, all vPC member ports are set to down state.
- For a vPC interface configuration Type 1 consistency check failure, the misconfigured vPC is set to down state.

When a Type 2 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 2 consistency check failure, all vPC member ports remain in up state and vPC systems trigger protective actions.
- For a vPC interface configuration Type 2 consistency check failure, the misconfigured vPC remains in up state. However, depending on the discrepancy type, vPC systems will trigger protective actions. The most typical misconfiguration deals with the allowed VLANs in the vPC interface trunking configuration. In this case, vPC systems will disable the vPC interface VLANs that do not match on both sides.

Display vPC Consistency Parameters:

```
dc31-102# sh vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
QoS                         2     ([], [], [], [], [],   ([], [], [], [], [],
                                  [])                    [])
Network QoS (MTU)           2     (9038, 0, 0, 0, 0, 0)  (9038, 0, 0, 0, 0, 0)
Network Qos (Pause)         2     (F, F, F, F, F, F)     (F, F, F, F, F, F)
Input Queuing (Bandwidth)   2     (100, 0, 0, 0, 0, 0)   (100, 0, 0, 0, 0, 0)
Input Queuing (Absolute     2     (F, F, F, F, F, F)     (F, F, F, F, F, F)
Priority)
Output Queuing (Bandwidth)  2     (100, 0, 0, 0, 0, 0)   (100, 0, 0, 0, 0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F)     (F, F, F, F, F, F)
Priority)
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
IGMP Snooping Group-Limit   2     4000                   4000
Interface-vlan admin up     2     1,11-410               1,11-410
Interface-vlan routing      2     1,11-410               1,11-410
capability
Allowed VLANs               -     1,11-410               1,11-410
Local suspended VLANs       -     -                      -

dc31-102# sh vpc consistency-parameters interface p11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch
```

```
Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Shut Lan                    1     No                     No
STP Port Type               1     Edge Trunk Port        Edge Trunk Port
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-65, 800b, 0-23-4-ee-be-65, 800b,
                                   0, 0), (8000,          0, 0), (8000,
                                  40-55-39-3-e3-42, a,   40-55-39-3-e3-42, a,
                                  0, 0)]                 0, 0)]
mode                        1     active                 active
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
Admin port mode             1
vPC card type               1     Empty                  Empty
Allowed VLANs               -     11-20                  11-20
Local suspended VLANs       -     -                      -
```

3.4.4.1.5 vPC Role Priority

There are two defined vPC roles: primary and secondary. The vPC role defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDUs) and responds to Address Resolution Protocol (ARP).
In case of a tie (same role priority value defined on both peer devices), the lowest system MAC will dictate the primary peer device.

Display vPC Role, System-MAC, System-Priority:

```
dc31-102# show vpc role

vPC Role status
----------------------------------------------------
vPC role                         : primary
Dual Active Detection Status     : 0
vPC system-mac                   : 00:23:04:ee:be:65
vPC system-priority              : 32667
vPC local system-mac             : 00:2a:6a:35:a8:c1
vPC local role-priority          : 201
```

3.4.4.1.6 vPC Peer-Link

The vPC peer-link is a standard 802.1Q trunk that performs the following actions:

- Carry vPC and non-vPC VLANs.
- Carry Cisco Fabric Services (CFS) messages that are tagged with CoS=4 for reliable communication CoS=4 for reliable communication.
- Carry flooded traffic between the vPC peer devices.
- Carry STP BPDUs, HSRP hello messages, and IGMP updates.

When the vPC peer-link fails and the vPC peer-keepalive link is still up, the vPC secondary peer device performs the following operations:

- Suspends its vPC member ports
- Shuts down the SVI associated to the vPC VLAN

Display vPC Peer-link Information:

```
dc31-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 40
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po101  up     1,11-410

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
12     Po12        up     success     success                    21-30
13     Po13        up     success     success                    31-40
14     Po14        up     success     success                    41-50
15     Po15        up     success     success                    51-60
16     Po16        up     success     success                    61-70
17     Po17        up     success     success                    71-80
18     Po18        up     success     success                    81-90
19     Po19        up     success     success                    91-100
20     Po20        up     success     success                    101-110
```

3.4.4.1.7 vPC Peer-Keepalive Link

The vPC peer-keepalive link is a Layer 3 link that joins one vPC peer device to the other vPC peer device and carries a periodic heartbeat between those devices. It is used at the boot up of the vPC systems to guarantee that both peer devices are up before forming the vPC domain. It is also used when the vPC peer-link fails, in which case, the vPC peer-keepalive link is leveraged to detect split brain scenario (both vPC peer devices are active-active).

Default Values for VPC Peer-Keepalive Links:

Timer	Default value
Keepalive interval	1 seconds
Keepalive hold timeout (on vPC peer-link loss)	3 seconds
Keepalive timeout	5 seconds

Display vPC Peer-Keepalive Information:

```
dc31-102# show vpc peer-keepalive

vPC keep-alive status             : peer is alive
```

```
--Peer is alive for             : (336696) seconds, (32) msec
--Send status                   : Success
--Last send at                  : 2014.02.19 17:07:46 62 ms
--Sent on interface             : Eth1/41
--Receive status                : Success
--Last receive at               : 2014.02.19 17:07:46 61 ms
--Received on interface         : Eth1/41
--Last update from peer         : (0) seconds, (340) msec

vPC Keep-alive parameters
--Destination                   : 1.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc-keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

3.4.4.1.8 vPC Member Link

As suggested by the name, a vPC member port is a port-channel member of a vPC. A port-channel defined as a vPC member port always contains the keywords *vpc <vpc id>.*

A vPC only supports Layer 2 port-channels. The port-channel can be configured in access or trunk switchport mode. Any VLAN allowed on the vPC member port is by definition called a vPC VLAN. Whenever a vPC VLAN is defined on a vPC member port, it must also be defined on the vPC peer-link. Not defining a vPC VLAN on the vPC peer-link will cause the VLAN to be suspended.

The configuration of the vPC member port must match on both the vPC peer devices. If there is an inconsistency, a VLAN or the entire port channel may be suspended (depending on Type-1 or Type-2 consistency check for the vPC member port). For instance, a MTU mismatch will suspend the vPC member port.

Display vPC Member Port-channel Information:

```
dc31-102# sh vpc br
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 40
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po101  up     1,11-410

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
```

```
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
12     Po12        up     success     success                    21-30
13     Po13        up     success     success                    31-40
14     Po14        up     success     success                    41-50
15     Po15        up     success     success                    51-60
16     Po16        up     success     success                    61-70
17     Po17        up     success     success                    71-80
18     Po18        up     success     success                    81-90
19     Po19        up     success     success                    91-100
dc31-102# show vpc consistency-parameters interface port-channel 11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Shut Lan                    1     No                     No
STP Port Type               1     Edge Trunk Port        Edge Trunk Port
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-65, 800b, 0-23-4-ee-be-65, 800b,
                                   0, 0), (8000,          0, 0), (8000,
                                  40-55-39-3-e3-42, a,   40-55-39-3-e3-42, a,
                                  0, 0)]                 0, 0)]
mode                        1     active                 active
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
Admin port mode             1
vPC card type               1     Empty                  Empty
Allowed VLANs               -     11-20                  11-20
Local suspended VLANs       -     -                      -
```

3.4.4.1.9 vPC ARP Synchronization

The vPC ARP Synchronization feature improves the convergence time for Layer 3 flows (North to South traffic). When the vPC peer-link fails and subsequently recovers, vPC ARP Synchronization performs an ARP bulk synchronization over Cisco Fabric Services (CFS) from the vPC primary peer device to the vPC secondary peer device.

Displays vPC ARP Synchronization Information:

```
dc31-101# sh ip arp sync-entries

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Address         Age       MAC Address     Interface
131.11.155.252  00:01:45  0000.8c43.eb64  Vlan410
131.11.155.253  00:01:45  0000.8c43.5e23  Vlan410
131.11.155.254  00:01:45  0000.8c44.59ef  Vlan410
131.11.154.252  00:01:45  0000.8c43.eb62  Vlan409
…
```

3.4.4.1.10 vPC Delay Restore

After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from the access to the core until the Layer 3 connectivity is reestablished.

The vPC Delay Restore feature delays the vPCs leg bringup on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on the vPC leg. The result provides a graceful restoration along with zero packet loss during the recovery phase (traffic still gets diverted to the alive vPC peer device).

This feature is enabled by default with a vPC restoration default timer of 30 seconds, which DC31 maintains in the testbed.

3.4.4.1.11 vPC Auto-Recovery

vPC auto-recovery feature was designed to address 2 enhancements to vPC.

- To provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure (vPC auto-recovery feature).
- To handle a specific case where both vPC peer devices reload but only one comes back to life (vPC auto-recovery reload-delay feature).

The switch which unsuspends its vPC role with vPC auto-recovery continues to remain primary even after peer-link is on. The other peer takes the role of secondary and suspends its own vPC until a consistency check is complete. Therefore, to avoid this situation from occurring erroneously, auto-recovery reload-delay-timer should be configured to be long enough for the system to fully complete its bootup sequence.

Helpful Commands for vPC Object Tracking:

Show vpc brief	Displays Auto-recovery status

Configuration Check:

```
dc31-102# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 40
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po101  up     1,11-410
```

```
vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
12     Po12        up     success     success                    21-30
13     Po13        up     success     success                    31-40
14     Po14        up     success     success                    41-50
15     Po15        up     success     success                    51-60
16     Po16        up     success     success                    61-70
17     Po17        up     success     success                    71-80
18     Po18        up     success     success                    81-90
19     Po19        up     success     success                    91-100
20     Po20        up     success     success                    101-110
```

3.4.4.1.12 PIM Pre-Build-SPT with vPC

PIM Pre-build SPT on non-forwarder attracts multicast traffic by triggering upstream PIM J/Ps (Join/Prune) without setting any interface in the OIF (Outgoing Interface) list. Multicast traffic is then always pulled to the non-active forwarder and finally dropped due to no OIFs.

The immediate effect of enabling PIM Pre-build SPT is to improve the convergence time upon active forwarder failure (1 to 3 seconds of convergence time). The other vPC peer device (which is the non-active forwarder) does not need to create any new upstream multicast state and can quickly transition to the active forwarder role by properly programming the OIF (Outgoing Interface) list.
The impact of enabling PIM prebuild SPT is the consumption of bandwidth and replication capacity on the primary and secondary data path (i.e. on vPC primary and secondary peer devices) in steady state.

As shown below, on the non-forwarder/secondary the (S,G) is created with no OIFs.

On Non-Forwarder:

```
dc31-101# show ip mroute 230.31.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.31.0.1/32), uptime: 2d02h, ip pim igmp
  Incoming interface: Ethernet2/1, RPF nbr: 31.101.11.1
  Outgoing interface list: (count: 10)
    Vlan19, uptime: 00:06:25, igmp
    Vlan18, uptime: 00:06:25, igmp
    Vlan20, uptime: 00:06:25, igmp
    Vlan13, uptime: 00:06:26, igmp
    Vlan17, uptime: 00:06:26, igmp
    Vlan15, uptime: 00:06:26, igmp
    Vlan14, uptime: 00:06:26, igmp
    Vlan16, uptime: 00:06:27, igmp
    Vlan11, uptime: 00:06:27, igmp
    Vlan12, uptime: 00:06:27, igmp

(131.50.11.11/32, 230.31.0.1/32), uptime: 2d02h, ip pim
  Incoming interface: Ethernet2/1, RPF nbr: 31.101.11.1
  Outgoing interface list: (count: 0)

DC5-DC101-5# sh ip pim intern vpc rp
PIM vPC RPF-Source Cache for Context "default" - Chassis Role Secondary


Source: 131.50.11.11
  Pref/Metric: 20/0
  Source role: secondary
  Forwarding state: Tie (not forwarding)
```

3.4.4.1.13 HSRP/HSRPv6 Active/Active with vPC

HSRP in the context of vPC has been improved from a functional and implementation standpoint to take full benefits of the L2 dual-active peer devices nature offered by vPC technology. HSRP operates in active-active mode from a data plane standpoint, as opposed to classical active/standby implementation with a STP based network. No additional configuration is required. As soon as a vPC domain is configured and interface VLAN with an associated HSRP group is activated, HSRP will behave by default in active/active mode (on the data plane side).

From a control plane standpoint, active-standby mode still applies for HSRP in context of vPC; the active HSRP instance responds to ARP request. ARP response will contain the HSRP vMAC which is the same on both vPC peer devices. The standby HSRP vPC peer device just relays the ARP request to active HSRP peer device through the vPC peer-link.

HSRPv4&v6 Configurations:

```
N6000 1:
interface Vlan11
  no shutdown
  mtu 9000
  ip address 131.10.11.2/24
  ipv6 address 2001:131:10:11::2/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 50
    ip 131.10.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 50
    ip 2001:131:10:11::1
```

```
N6000 2:
interface Vlan11
  no shutdown
  mtu 9000
  ip address 131.10.11.3/24
  ipv6 address 2001:131:10:11::3/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 150
    ip 131.10.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 150
    ip 2001:131:10:11::1
```

Helpful Commands for HSRP Active/Active with vPC:

Show hsrp brief	Displays hsrp status
Show mac address-table vlan <vlan id>	Displays mac addresses including HSRP vMAC; check for G-flag on vMAC for active/active HSRP

Configuration Check:

```
dc31-102# sh hsrp brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr      Standby addr     Group addr
Vlan11      1   150  P Active   local            131.10.11.2      131.10.11.1     (conf)
Vlan11      101 150  P Active   local            fe80::22a:6aff:fe37:d1bc  fe80::5:73ff:fea0:65 (impl auto
EUI64)
Vlan12      1   150  P Active   local            131.10.12.2      131.10.12.1     (conf)
Vlan12      101 150  P Active   local            fe80::22a:6aff:fe37:d1bc  fe80::5:73ff:fea0:65 (impl auto
EUI64)
…
```

3.5 DC32

3.5.1 Configuration of Platform Specific Features On DC32

3.5.1.1 Licensing

License Usage on Nexus 3548 in DC32:

```
N3548# sh license usage
Feature                        Ins  Lic   Status Expiry Date Comments
                                    Count
--------------------------------------------------------------------------------
24P_LIC_PKG                     No    -   Unused             -
24P_UPG_PKG                     No    -   Unused             -
LAN_BASE_SERVICES_PKG           Yes   -   In use Never       -
ALGO_BOOST_SERVICES_PKG         Yes   -   Unused Never       -
LAN1K9_ENT_SERVICES_PKG         No    -   Unused             -
LAN_ENTERPRISE_SERVICES_PKG     Yes   -   In use Never       -
--------------------------------------------------------------------------------
```

Although features can be enabled and configured in the CLI without licenses, they will not function until the license is installed.

3.5.1.2 Out-of-Band Management Network

DC32 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic.

Configuration:

```
vrf context management
  ip route 0.0.0.0/0 10.2.0.1

interface mgmt0
  vrf member management
  ip address 10.2.32.2/16
```

3.5.1.3 Common Configurations

3.5.1.3.1 SSH and TACACS+

SSH is enabled in DC32 to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration and Verification:

```
feature tacacs+


ip tacacs source-interface mgmt0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
    use-vrf management

N3548# show ssh server
ssh version 2 is enabled


N3548# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
```

```
admin    ttyS0        Feb 19 11:26 07:09        3804
interop  pts/0        Feb 18 10:37   .         14531 (taro.interop.cisco.com) session=ssh *
```

3.5.1.3.2 CDP and LLDP

CDP and LLDP are pervasively used on the DC32 testbed for inter-device discovery.

```
DC32-1# sh run cdp all
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
cdp format device-id system-name


interface mgmt0
  cdp enable

interface Ethernet1/1
  cdp enable

interface Ethernet1/2
  cdp enable
…

DC32-1# sh cdp nei
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute


Device-ID             Local Intrfce Hldtme Capability  Platform      Port ID
mgmt-sw3.interop.cisco.com
                    mgmt0          166    R S I     WS-C6504-E    Gig3/35
DC32-101.interop.cisco.com(FOC1704R08C)
                    Eth1/1         160    R S I s   N3K-C3548P-10 Eth1/1
DC32-101.interop.cisco.com(FOC1704R08C)
                    Eth1/2         156    R S I s   N3K-C3548P-10 Eth1/2
DC32-101.interop.cisco.com(FOC1704R08C)
…
```

```
DC32-1# sh run lldp all

feature lldp

lldp holdtime 120
lldp reinit 2
lldp timer 30
lldp tlv-select port-description
lldp tlv-select system-name
lldp tlv-select system-description
lldp tlv-select system-capabilities
lldp tlv-select management-address
lldp tlv-select dcbxp
lldp tlv-select port-vlan

interface mgmt0
  lldp transmit
  lldp receive

interface Ethernet1/1
  lldp transmit
  lldp receive
…
```

```
DC32-1# sh lldp nei
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID             Local Intf      Hold-time  Capability  Port ID
DC32-101.interop.cisco.com Eth1/1          120        BR          Eth1/1

DC32-101.interop.cisco.com Eth1/2          120        BR          Eth1/2
…
```

3.5.1.3.3 Syslog

Syslog is used to record all network events on the DC32 test bed. Whenever possible,DC32 makes use of a separate management VRF for syslog.

Configuration and Verification:

```
logging server syslog.interop.cisco.com 5 use-vrf management facility local6

N3548# sh logging server
Logging server:                 enabled
{syslog.interop.cisco.com}
        server severity:        notifications
        server facility:        local6
        server VRF:             management
```

3.5.1.3.4 SNMP

SNMP is used for system monitoring in DC32. Scripts are used to poll the systems asynchronously during the course of all DC32 test execution. Intense mibwalk processing may trigger SNMPD to crash on the Nexus 3548 (CSCum13379).

Configuration:

```
snmp-server source-interface trap mgmt0
snmp-server user admin network-admin auth md5 0x14383e3d6d3c3051fb7276c3c4874a91
priv 0x14383e3d6d3c3051fb7276c3c4874a91 localizedkey
snmp-server host 172.28.92.62 traps version 2c public
snmp-server enable traps callhome event-notify
snmp-server enable traps callhome smtp-send-fail
snmp-server enable traps cfs state-change-notif
snmp-server enable traps lldp lldpRemTablesChange
snmp-server enable traps cfs merge-failure
snmp-server enable traps aaa server-state-change
snmp-server enable traps upgrade UpgradeOpNotifyOnCompletion
snmp-server enable traps upgrade UpgradeJobStatusNotify
snmp-server enable traps feature-control FeatureOpStatusChange
snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended
snmp-server enable traps config ccmCLIRunningConfigChanged
snmp-server enable traps snmp authentication
snmp-server enable traps link cisco-xcvr-mon-status-chg
snmp-server enable traps vtp notifs
snmp-server enable traps vtp vlancreate
snmp-server enable traps vtp vlandelete
snmp-server enable traps bridge newroot
snmp-server enable traps bridge topologychange
snmp-server enable traps stpx inconsistency
snmp-server enable traps stpx root-inconsistency
snmp-server enable traps stpx loop-inconsistency
snmp-server enable traps poe portonoff
```

```
snmp-server enable traps poe pwrusageon
snmp-server enable traps poe pwrusageoff
snmp-server enable traps poe police
snmp-server community cisco group network-operator
snmp-server community private group network-admin
snmp-server community public group network-operator


N3548# sh snmp trap
--------------------------------------------------------------------------------
Trap type                      Description                   Enabled
--------------------------------------------------------------------------------
ospf-32              : OSPF base traps                              No
ospf-32              : OSPF LSA                                     Yes
BGP-32               :                                              No
entity               : entity_mib_change               Yes
entity               : entity_module_status_change     Yes
entity               : entity_power_status_change      Yes
entity               : entity_module_inserted          Yes
entity               : entity_module_removed           Yes
entity               : entity_unrecognised_module      Yes
entity               : entity_fan_status_change        Yes
entity               : entity_power_out_change         Yes
link                 : linkDown                        Yes
link                 : linkUp                          Yes
link                 : extended-linkDown               Yes
link                 : extended-linkUp                 Yes
link                 : cieLinkDown                     Yes
link                 : cieLinkUp                       Yes
link                 : delayed-link-state-change       Yes
callhome             : event-notify                    Yes
callhome             : smtp-send-fail                  Yes
cfs                  : state-change-notif              Yes
cfs                  : merge-failure                   Yes
rf                   : redundancy_framework            Yes
aaa                  : server-state-change             Yes
license              : notify-license-expiry           Yes
license              : notify-no-license-for-feature   Yes
license              : notify-licensefile-missing      Yes
license              : notify-license-expiry-warning   Yes
upgrade              : UpgradeOpNotifyOnCompletion     Yes
upgrade              : UpgradeJobStatusNotify          Yes
feature-control      : FeatureOpStatusChange           Yes
sysmgr               : cseFailSwCoreNotifyExtended     Yes
rmon                 : risingAlarm                     Yes
rmon                 : fallingAlarm                    Yes
rmon                 : hcRisingAlarm                   Yes
rmon                 : hcFallingAlarm                  Yes
config               : ccmCLIRunningConfigChanged      Yes
snmp                 : authentication                  Yes
link                 : cisco-xcvr-mon-status-chg       Yes
vtp                  : notifs                          Yes
vtp                  : vlancreate                      Yes
vtp                  : vlandelete                      Yes
bridge               : newroot                         Yes
bridge               : topologychange                  Yes
stpx                 : inconsistency                   Yes
stpx                 : root-inconsistency              Yes
stpx                 : loop-inconsistency              Yes
entity               : entity_sensor                   Yes
poe                  : portonoff                       Yes
poe                  : pwrusageon                      Yes
poe                  : pwrusageoff                     Yes
poe                  : police                          Yes
pim                  : pimNeighborLoss                 Yes
lldp                 : lldpRemTablesChange             Yes
```

3.5.1.3.5 NTP

NTP is used to synchronize the clocks on all DC32 devices to provide consistent timestamps on all network logs and events.

Configuration and Verification:

```
ntp distribute
ntp server 172.28.92.1
ntp commit

N3548# show ntp status
Distribution : Enabled
Last operational state: No session

N3548# show ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                 st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                 8    64      377   0.00200 management
```

3.5.1.3.6 SPAN

SPAN has been enabled on DC32 switches to provide packet captures to assist in network debugging.

Configuration and Verification:

```
monitor session 1
  source interface port-channel1031 both
  destination interface Ethernet1/37
  no shut
N3548# sh monitor session 1
   session 1
---------------
type              : local
state             : up
source intf       :
    rx            : Po1031
    tx            : Po1031
    both          : Po1031
source VLANs      :
    rx            :
destination ports : Eth1/37

Legend: f = forwarding enabled, l = learning enabled
```

3.5.1.3.7 DNS

DNS has been enabled to provide name lookup in DC32 network.

Configuration and Verification:

```
vrf context management
  ip domain-name interop.cisco.com
  ip domain-list interop.cisco.com
  ip domain-list cisco.com
  ip name-server 172.28.92.9 172.28.92.10

ip domain-lookup
ip domain-name interop.cisco.com
```

```
ip name-server 172.28.92.9 use-vrf management

DC32-1# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=0.961 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=0.731 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=1.54 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=1.542 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=1.515 ms
```

3.5.1.3.8 MTU

In order to configure the MTU to handle jumbo frames in the Nexus 3548 switches, the following policy-map has to be applied.

Configuration:

```
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216

system qos
  service-policy type network-qos jumbo

N3548# sh policy-map type network-qos jumbo


  Type network-qos policy-maps
  ===============================

  policy-map type network-qos jumbo
    class type network-qos class-default

      mtu 9216

DC32-1# sh queuing interface ethernet 1/1
Ethernet1/1 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0       WRR            100

  RX Queuing
    Multicast statistics:
        Mcast pkts dropped                      : 0
    Unicast statistics:
    qos-group 0
    HW MTU: 9216 (9216 configured)
    drop-type: drop, xon: 0, xoff: 0
    Statistics:
        Ucast pkts dropped                      : 0
```

3.5.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

The default PIMREG CoPP is 200pps. The PIMREG CoPP configuration at the multicast RP determines the rate of PIM source registration and periodic null-registers that can be processed. The PIMREG CoPP at the RP should be adjusted accordingly to accommodate the registration rates to prevent potential mroute states from timing out.

For example, there are 2000 active sources on the DC32 testbed, with bursts of 500 requiring registration. Testing found that a PIMREG CoPP of 1000pps was adequate to accommodate this number of multicast sources and burst pattern.

In addition, DC32 also modifies the police rate for ARP CoPP from 200 pps to 500 pps to allow ARP replies to be processed in a timely manner (CSCun37500).

The remaining values are kept to their default values.

```
N3548# sh copp status
Last Config Operation: None
Last Config Operation Timestamp: None
Last Config Operation Status: None
Policy-map attached to the control-plane: copp-system-policy
```

```
policy-map type control-plane copp-system-policy
  class copp-s-pimreg
    police pps 1000
```

Nexus 3548 CoPP:

```
policy-map type control-plane copp-system-policy
  class copp-s-default
    police pps 400
  class copp-s-ping
    police pps 100
  class copp-s-l3destmiss
    police pps 100
  class copp-s-glean
    police pps 500
  class copp-s-l3mtufail
    police pps 100
  class copp-s-ttl1
    police pps 100
  class copp-s-ip-options
    police pps 100
  class copp-s-ip-nat
    police pps 100
  class copp-s-ipmcmiss
    police pps 400
  class copp-s-ipmc-g-hit
    police pps 400
  class copp-s-ipmc-rpf-fail-g
    police pps 400
  class copp-s-ipmc-rpf-fail-sg
    police pps 400
  class copp-s-dhcpreq
    police pps 300
  class copp-s-dhcpresp
    police pps 300
  class copp-s-igmp
    police pps 400
  class copp-s-routingProto2
    police pps 1300
  class copp-s-eigrp
    police pps 200
  class copp-s-pimreg
    police pps 1000
  class copp-s-pimautorp
    police pps 200
  class copp-s-routingProto1
    police pps 5000
  class copp-s-arp
```

```
    police pps 500
  class copp-s-ptp
    police pps 1000
  class copp-s-bpdu
    police pps 12000
  class copp-s-cdp
    police pps 400
  class copp-s-lacp
    police pps 400
  class copp-s-lldp
    police pps 200
  class copp-icmp
    police pps 200
  class copp-telnet
    police pps 500
  class copp-ssh
    police pps 500
  class copp-snmp
    police pps 500
  class copp-ntp
    police pps 100
  class copp-tacacsradius
    police pps 400
  class copp-stftp
    police pps 400
  class copp-ftp
    police pps 100
  class copp-http
    police pps 100
```

3.5.1.5 ECMP for IPv4 host routes

ECMP support for host routes is enabled by default on the Nexus 3548 switches.

On the Nexus 3548 running 6.0(2)A1(1c), the *hardware profile unicast enable-host-ecmp* configuration is inconsistent (CSCuj95690). *Hardware profile unicast enable-host-ecmp* configuration shows that the unicast enable-host-ecmp is enabled, while the more specific IPv4 and IPv6 related commands indicate otherwise. The initial *enable-host-ecmp* takes precedence. It should be noted that disabling enable-host-ecmp is not supported by itself, but disabling ECMP as a whole is supported by setting the *maximum-paths* config to 1 in BGP.

```
N3548# sh run all | i profile
hardware profile multicast max-limit 8192
hardware profile multicast prefer-source-tree eternity
hardware profile unicast enable-host-ecmp
hardware profile multicast syslog-threshold 0
hardware profile unicast syslog-threshold 0
no hardware profile unicast enable-host-ecmp ipv4
no hardware profile unicast enable-host-ecmp ipv6
no hardware profile unicast enable-host-ecmp arp-nd
no hardware profile unicast enable-host-ecmp ipv4 arp
no hardware profile unicast enable-host-ecmp ipv6 nd
```

3.5.2 Routing Design Overview

3.5.2.1 Unicast Routing Design

3.5.2.1.1 BGP IPv4 Routing Design

The 3548 switches on DC32 do not support IPv6 routing on with version 6.0(2)A1(1c) of NX-OS. The network is split into three layers: core, spine and leaf. The layers are logically connected to each other through eBGP, as shown in Figure 32. The N7K core layer in BGP AS 3 is shared with other DC3 networks (DC31, DC33, and DC36). The spine layer runs OSPF to provide inter-switch connectivity to support iBGP sessions. The leaf layer is divided into multiple BGP ASes. This BGP logical design is easier to configure, maintain and debug than full mesh ibgp, route reflector, or confederations; the core can consolidate these as private ASes if there is a need to advertise to other BGP exchanges.

The leaf layer represents different top of rack topologies that can be deployed. AS 32101 employs N3548 in a traditional spanning tree topology, using HSRP for gateway redundancy for nodes. AS 32103 employs a routed top of rack with N3548. AS 32104 employs a routed Nexus 3048 ToR.

AS 32105 is used as a test tool rather than network under test. The Nexus 7000 is divided into multiple VRFs, with each VRF representing an extra ToR in the network. The goal is to test increasing number of ToR supported by the spine layer.

Figure 32 DC32 BGP Logical Design

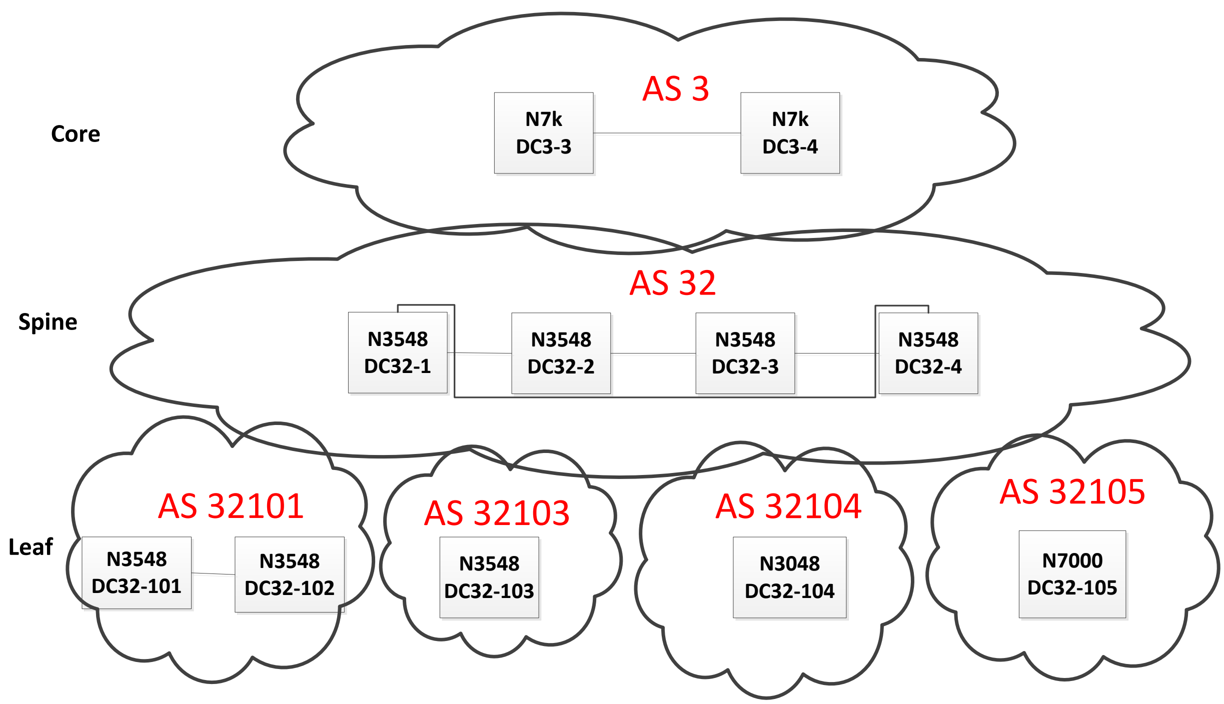

BGP peer templates are used to simplify the configuration.

DC32 Spine BGP configuration:

```
feature bgp

router bgp 32
  router-id 40.32.0.1
```

```
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
   network 32.101.11.0/24
   …
    network 40.32.254.1/32
    maximum-paths 32
  template peer BGPLEAF
    address-family ipv4 unicast
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  neighbor 32.101.11.101 remote-as 32101
    inherit peer BGPLEAF
  …
  neighbor 40.32.25.17 remote-as 3
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
```

DC32 Leaf BGP configuration:

```
router bgp 32103
  router-id 32.0.0.103
  address-family ipv4 unicast
    network 32.0.0.103/32
…
    network 132.103.100.0/24
    maximum-paths 32
  template peer BGPLEAF
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 32.103.11.1 remote-as 32
    inherit peer BGPLEAF
…
  neighbor 32.103.44.4 remote-as 32
    inherit peer BGPLEAF
```

3.5.2.1.1.1 BGP Router-ID

To establish BGP sessions between peers, BGP must have a router ID, which is sent to BGP peers in the OPEN message when a BGP session is established. On DC32, NVT has configured a loopback interface IP address as the BGP router-ID. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

To Verify the BGP Router-ID:

```
N3548# sh ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 134215, local router ID is 40.32.0.1
```

3.5.2.1.1.2 BGP Address Family

BGP address family for IPv4 has been configured to achieve BGP peering, load-balancing, default route injection.

To Verify the BGP Address Family:

```
N3548# sh ip bgp all summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 40.32.0.1, local AS number 32
BGP table version is 134215, IPv4 Unicast config peers 36, capable peers 36
459 network entries and 4417 paths using 255388 bytes of memory
BGP attribute entries [12/1632], BGP AS path entries [5/30]
BGP community entries [0/0], BGP clusterlist entries [0/0]
4381 received paths for inbound soft reconfiguration
4381 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
32.101.11.101   4 32101   43311   44664   134215    0    0     1w0d 168
32.101.12.101   4 32101   43313   44662   134215    0    0     1w0d 168
32.101.13.101   4 32101   43313   44663   134215    0    0     1w0d 168
32.101.14.101   4 32101   43310   44662   134215    0    0     1w0d 168
32.101.15.101   4 32101   43312   44668   134215    0    0     1w0d 168
32.101.16.101   4 32101   43309   44662   134215    0    0     1w0d 168
32.101.17.101   4 32101   43310   44664   134215    0    0     1w0d 168
```

3.5.2.1.1.3 BGP Load Sharing and ECMP

DC32 has configured the maximum-paths that BGP adds to the route table for equal-cost multipath load balancing as 32 for spine and leaf peers for IPv4 unicast address family.

3.5.2.1.1.4 BGP Authentication

DC32 has configured MD5 Authentication for BGP sessions.

To Verify the BGP Authentication:

```
N3548# sh ip bgp neighbors 32.101.11.101
BGP neighbor is 32.101.11.101,  remote AS 32101, ebgp link,  Peer index 6
  Inherits peer configuration from peer-template BGPLEAF
  BGP version 4, remote router ID 32.0.0.101
  BGP state = Established, up for 00:03:34
  Peer is directly attached, interface Ethernet1/1
  TCP MD5 authentication is enabled
```

3.5.2.1.1.5 BGP Update-Source

DC32 has configured BGP update-source to establish a BGP multi-hop sessions. DC32 has multi-hop sessions only on the iBGP peering between the spine switches.

To Verify the BGP Update-Source:

```
DC32-1# sh ip bgp neighbors 40.32.0.3
BGP neighbor is 40.32.0.3,  remote AS 32, ibgp link,  Peer index 2
  BGP version 4, remote router ID 40.32.0.3
  BGP state = Established, up for 4w0d
  Using loopback0 as update source for this peer
```

3.5.2.1.1.6 BGP Default Route

BGP default route is advertised from the spine peers to the leaf peers for Ipv4 address family.

To Verify the BGP Default Route:

```
DC32-1# sh ip bgp neighbors 32.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 134215, neighbor version 134215
  168 accepted paths consume 8736 bytes of memory
  354 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 32.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 00:00:06 after session start

  Local host: 32.101.11.1, Local port: 179
  Foreign host: 32.101.11.101, Foreign port: 22636
  fd = 60
```

3.5.2.1.1.7 BGP Next-Hop-Self

BGP next-hop-self is configured for BGP sessions between the spine switches for IPv4 address family.

To Verify the BGP Next-Hop-Self:

```
DC32-1# sh ip bgp neighbors 32.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 134215, neighbor version 134215
  168 accepted paths consume 8736 bytes of memory
  354 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 32.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 00:00:06 after session start

  Local host: 32.101.11.1, Local port: 179
  Foreign host: 32.101.11.101, Foreign port: 22636
  fd = 60
```

3.5.2.1.1.8 BGP Soft-Reconfiguration

BGP Soft reset is recommended because it allows routing tables to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis.

```
DC32-1# sh ip bgp neighbors 32.101.11.101  | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 134215, neighbor version 134215
  168 accepted paths consume 8736 bytes of memory
  354 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 32.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 00:00:06 after session start

  Local host: 32.101.11.1, Local port: 179
  Foreign host: 32.101.11.101, Foreign port: 22636
  fd = 60
```

3.5.2.1.2 OSPF Routing Design

OSPF is used as the IGP to provide reachability for establishing iBGP peering at the spine layer only. The OSPF process is enabled only on directly connected interfaces and the Loopback interface. All the OSPF enabled interfaces are in Area 0.0.0.0. Each OSPF network type is set to point-to-point to decrease OSPF neighbor setup latency. In order to improve OSPF convergence, SPF and LSA timers are throttled to (100 200 5000 and 50 100 300) respectively.

OSPF Router Configuration:

```
N3548# sh run ospf
feature ospf

router ospf 32
  router-id 40.32.0.1
  log-adjacency-changes
  timers throttle spf 100 200 5000
  timers throttle lsa 50 100 300

interface loopback0
  ip router ospf 32 area 0.0.0.0

interface port-channel1
  ip ospf network point-to-point
  ip router ospf 32 area 0.0.0.0

interface port-channel2
  ip ospf network point-to-point
  ip router ospf 32 area 0.0.0.0
…
```

3.5.2.1.3 Unicast Forwarding Verification

This Switch is the Authoritative Router for a Directly Connected Subnet on VLAN 11 132.101.52.0/24:

```
DC32-102# sh run int vlan 52

interface Vlan52
  no shutdown
  no ip redirects
  ip address 132.101.52.3/24
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
    ip 132.101.52.1
```

The host 132.101.52.51 has been Learned via ARP on this Subnet.

```
DC32-102# sh ip arp 132.101.52.51

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
132.101.52.51   00:07:48  0084.6534.3300  Vlan52
```

On NX-OS, “show ip route” will also Show Directly Connected Hosts as /32 Routes:

```
DC32-102# sh ip route 132.101.52.51
```

```
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

132.101.52.51/32, ubest/mbest: 1/0, attached
    *via 132.101.52.51, Vlan52, [250/0], 1w0d, am
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
DC32-102# sh ip adjacency  132.101.52.51

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
132.101.52.51   0084.6534.3300  50   arp        Vlan52
```

Find the PO Interface on which this MAC Address is Learnt:

```
DC32-102# sh mac address-table address 0084.6534.3300
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 52       0084.6534.3300    dynamic   133380     F    F  Po205
```

Display PO7 Member Interface with Module Information:

```
DC32-102# sh port-channel summary | i Po205
205   Po205(SU)   Eth      LACP      Eth1/40(P)
```

Display Adjacency Index for this Route 132.101.52.51 in the Hardware Table:

```
DC32-101# sh hardware internal libsdk mtc l3 host-tbl valid-only | i 132.101.052.051
================================================================================
                         HOST TABLE RESULT RAM INFORMATION
================================================================================

                                                   E          C C N I
                                                E  C          P P O N
                                                C  M            U T   R
                                                M  P          2   4 O O   R
                                                P           D   C   V U   O
                                                   M        R C O F F T  DU O
 HOST           VRF                  ADJACENCY  E  O STATS  O P D W L  EI ET O
 INDEX  CLASS    ID     IP ADDR       VID/PID   N  D   IDX  P U E D W RP FE O
 -----  ----- ----- --------------- ----------- - -- -----  - - - - - -- -- -
 40720  v4 uc     1 132.101.052.051   737/29488 0  0     0  0 0 0 0 0  0  0 0
```

Display DMAC Entry Programmed in the Adjacency Table:

```
DC32-101# sh hardware internal libsdk mtc l3 adj-tbl start-addr 29488 end-addr 29488

================================================================================
                              ADJACENCY TABLE INFORMATION
================================================================================

                                                           C  E L        E L
                                                           O  N O        X 3  S
                                                          RU AT C L  LE NI   LYE
   ADJ    KEY        BD/                                  EN DR A RU RC OS I RNC
 INDEX  CLASS  VLAN  FID  LIF DVIF  LID     MAC ADDR      FT JY L NC NH NT F NCH
```

```
-----   ------ ---- ---- ---- ---- ---- ----------------- -- -- - -- -- -- - ---
29488   eth uc    0   62   49    0   53 00:84:65:34:33:00  1  1 1  0  0  0 0   0
```

Route verification: "show ip route"

```
N3548# sh ip route 132.101.52.0/24
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

132.101.52.0/24, ubest/mbest: 16/0
    *via 32.101.11.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.12.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.13.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.14.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.15.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.16.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.17.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.101.18.101, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.102.11.102, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.102.12.102, [20/0], 00:11:35, bgp-32, external, tag 32101
    *via 32.102.13.102, [20/0], 00:11:35, bgp-32, external, tag 32101
    *via 32.102.14.102, [20/0], 00:11:35, bgp-32, external, tag 32101
    *via 32.102.15.102, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.102.16.102, [20/0], 00:11:35, bgp-32, external, tag 32101
    *via 32.102.17.102, [20/0], 00:11:36, bgp-32, external, tag 32101
    *via 32.102.18.102, [20/0], 00:11:35, bgp-32, external, tag 32101
```

Forwarding table verification: "show forwarding route"

```
N3548# show forwarding route 132.101.52.0/24

IPv4 routes for table default/base

------------------+------------------+----------------------+-----------------
Prefix            | Next-hop         | Interface            | Labels
------------------+------------------+----------------------+-----------------
*132.101.52.0/24    32.101.11.101      Ethernet1/1
                    32.101.12.101      Ethernet1/2
                    32.101.13.101      Ethernet1/3
                    32.101.14.101      Ethernet1/4
                    32.101.15.101      Ethernet1/5
                    32.101.16.101      Ethernet1/6
                    32.101.17.101      Ethernet1/7
                    32.101.18.101      Ethernet1/8
                    32.102.11.102      Ethernet1/9
                    32.102.12.102      Ethernet1/10
                    32.102.13.102      Ethernet1/11
                    32.102.14.102      Ethernet1/12
                    32.102.15.102      Ethernet1/13
                    32.102.16.102      Ethernet1/14
                    32.102.17.102      Ethernet1/15
                    32.102.18.102      Ethernet1/16
```

3.5.2.2 Multicast Routing Design

Multicast routing has been enabled across the entire DC32 network.

On the Nexus 3548 running the software release 6.0(2)A1(1c), the spine switch does not immediately remove OIF for interfaces that fail if they are routed ports. The consequence is that the switch could have mroutes with OIFs that are already down. The OIF will eventually get removed due to periodic PIM protocol state maintenance. However, the OIF is immediately removed if it is a routed port-channel – even a single member port channel. As a workaround, it is recommended to change all individual routed ports to single member port-channels when possible for Nexus 3548 (CSCul27880).

When multicast traffic is first sourced, the first hop router must send PIM registration messages to the RP. Any transit Nexus 3548 router will copy IPv4 unicast packets with protocol 103 to the CPU. Thus protocol 103 packets will be fordwarded through hardware and software resulting in packets being forwarded twice (CSCul883311). This issue has been resolved in the subsequent software releases.

On the Nexus 3548 when multicast source traffic is first sourced, the first hop router sends a PIM register towards the RP. Once the RP receives the register message, it will send a register stop to the first hop router which may discard these packets temporarily due to no (S,G) state created (CSCul56932).

On a network topology with anycast RP where multicast sources and receivers are on the same switch, the PIM RP may forward packets back toward the source DR due to the presence of receivers that joined to the (*,G). Because the source is also present, the DR has both (*,G) and (S,G) created for the local sources. The DR is expected to forward packets that match these sources using only the (S,G). The Nexus 3548 will forward packets using the (*,G) also – therefore, causing duplicate packets to be sent to the receiver (CSCum63413).

DC32 Multicast Configuration:

```
feature pim

ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
ip pim send-rp-announce loopback1 group-list 230.32.0.0/16
ip pim send-rp-discovery loopback1
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen

interface loopback1
  description dc32-RP
  ip address 40.32.254.1/32
  ip pim sparse-mode

feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.32.0.2 connect-source loopback0
ip msdp mesh-group 40.32.0.2 mesh32
ip msdp peer 40.32.0.3 connect-source loopback0
ip msdp mesh-group 40.32.0.3 mesh32
ip msdp peer 40.32.0.4 connect-source loopback0
ip msdp mesh-group 40.32.0.4 mesh32

interface loopback0
  ip address 40.32.0.1/32
  ip router ospf 32 area 0.0.0.0
  ip pim sparse-mode
```

3.5.2.2.1 PIM-ASM Rendezvous Point

PIM Sparse Mode has been configured as the protocol of choice for multicast routing. The RP is located at the spine layer.

3.5.2.2.1.1 Auto-RP

The DC32 testbed is designed to have the RP on the spine to support the groups sourced from that particular POD. DC32 makes use of Auto-RP to automate distribution of RP information in the network.

To Verify PIM RP:

```
N3548# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.32.254.1*, next Discovery message in: 00:00:50
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1, (0), uptime: 4w1d, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.32.254.1*, (0), uptime: 4w1d, expires: 00:02:10,
  priority: 0, RP-source: 40.32.254.1 (A), group ranges:
      230.32.0.0/16
DC32-1#


DC32-1# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.32.0.0/16      ASM       40.32.254.1      -
DC32-1#
```

3.5.2.2.1.1.1 Auto-RP Forward Listen

DC32 has enabled the Auto-RP listening and forwarding feature so that the Auto-RP mechanism can dynamically inform routers in the PIM domain of the group-to-RP mapping since PIM dense mode is not supported on NX-OS. By default, listening or forwarding of Auto-RP messages is not enabled on NX-OS.

3.5.2.2.1.2 Static RP

For the groups with a Rendezvous Point on the core, the RP is statically configured on all routers in the DC32 network.

To Verify PIM RP:

```
DC32-1# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.32.254.1*, next Discovery message in: 00:00:50
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
```

```
RP: 40.3.254.1, (0), uptime: 4w1d, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.32.254.1*, (0), uptime: 4w1d, expires: 00:02:10,
  priority: 0, RP-source: 40.32.254.1 (A), group ranges:
      230.32.0.0/16
DC32-1#


DC32-1# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.32.0.0/16      ASM       40.32.254.1      -
DC32-1#
```

3.5.2.2.1.3 Anycast RP with MSDP

DC32 has configured Anycast RP with MSDP at the spine layer. DC32 has also configured Anycast RP with MSDP among the spine routers.

DC32 Anycast RP and MSDP Configuration:

```
!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list 230.32.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  description dc32-RP
  ip address 40.32.254.1/32
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.32.0.2 connect-source loopback0
ip msdp mesh-group 40.32.0.2 mesh32
ip msdp peer 40.32.0.3 connect-source loopback0
ip msdp mesh-group 40.32.0.3 mesh32
ip msdp peer 40.32.0.4 connect-source loopback0
ip msdp mesh-group 40.32.0.4 mesh32
```

To Verify MSDP Peer and SA_Cache:

```
N3548# sh ip msdp sa-cache
MSDP SA Route Cache for VRF "default" - 767 entries
Source          Group           RP              ASN        Uptime
132.101.11.41   230.32.0.1      40.32.0.4       0          00:00:06
132.101.12.41   230.32.0.1      40.32.0.4       0          00:00:05
132.101.21.41   230.32.0.1      40.32.0.4       0          00:00:06
132.101.22.41   230.32.0.1      40.32.0.4       0          00:00:05
132.101.31.41   230.32.0.1      40.32.0.4       0          00:00:06
132.101.32.41   230.32.0.1      40.32.0.4       0          00:00:05
132.101.41.41   230.32.0.1      40.32.0.4       0          00:00:05
132.101.42.41   230.32.0.1      40.32.0.4       0          00:00:05
132.101.51.41   230.32.0.1      40.32.0.4       0          00:00:05
…

N3548# sh ip msdp sum
MSDP Peer Status Summary for VRF "default"
Local ASN: 32, originator-id: 40.32.0.1
```

```
Number of configured peers:  3
Number of established peers: 3
Number of shutdown peers:    0

Peer             Peer        Connection      Uptime/   Last msg  (S,G)s
Address          ASN         State           Downtime  Received  Received
40.32.0.2        0           Established     2d09h     00:00:04  0
40.32.0.3        0           Established     2d09h     00:00:06  1000
40.32.0.4        0           Established     2d09h     00:00:05  500
```

3.5.2.2.1.3.1 MSDP Mesh Group

MSDP Mesh Group is configured on the spines to prevent each MSDP peer from advertising SA's learned from other peers i.e., only locally registered sources.

```
feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.32.0.2 connect-source loopback0
ip msdp mesh-group 40.32.0.2 mesh32
ip msdp peer 40.32.0.3 connect-source loopback0
ip msdp mesh-group 40.32.0.3 mesh32
ip msdp peer 40.32.0.4 connect-source loopback0
ip msdp mesh-group 40.32.0.4 mesh32


DC32-1# sh run int lo 0
interface loopback0
  ip address 40.32.0.1/32
  ip router ospf 32 area 0.0.0.0
  ip pim sparse-mode
```

3.5.2.2.2 PIM SPT-Threshold

DC32 has enabled *ip pim spt-threshold infinity* on all last hop PIM routers to decrease the multicast entries.

3.5.2.2.3 Multicast Multipath

Cisco NX-OS Multicast Multipath is enabled by default and the load sharing selection algorithm is based on the source and group addresses.

3.5.2.2.4 Multicast Forwarding Verification

The following sequence of commands illustrates the verification of the Cisco NX-OS multicast L2 and L3 forwarding.

Displays a Specific Multicast Route 230.101.0.1 with Incoming Interface Information:

```
N3548# sh ip mroute 230.32.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.32.0.1/32), uptime: 3w2d, ip pim igmp
  Incoming interface: Ethernet1/28, RPF nbr: 32.102.44.4
```

```
  Outgoing interface list: (count: 10)
    Vlan51, uptime: 1w0d, igmp
    Vlan52, uptime: 1w0d, igmp
    Vlan41, uptime: 1w0d, igmp
    Vlan42, uptime: 1w0d, igmp
    Vlan32, uptime: 1w0d, igmp
    Vlan31, uptime: 1w0d, igmp
    Vlan22, uptime: 1w0d, igmp
    Vlan21, uptime: 1w0d, igmp
    Vlan12, uptime: 1w0d, igmp
    Vlan11, uptime: 1w0d, igmp


(132.101.11.41/32, 230.32.0.1/32), uptime: 00:01:08, ip mrib pim
  Incoming interface: Vlan11, RPF nbr: 132.101.11.41
  Outgoing interface list: (count: 10)
    Vlan11, uptime: 00:01:08, mrib, (RPF)
    Vlan12, uptime: 00:01:08, mrib
    Vlan21, uptime: 00:01:08, mrib
    Vlan22, uptime: 00:01:08, mrib
    Vlan31, uptime: 00:01:08, mrib
    Vlan32, uptime: 00:01:08, mrib
    Vlan41, uptime: 00:01:08, mrib
    Vlan42, uptime: 00:01:08, mrib
    Vlan51, uptime: 00:01:08, mrib
    Vlan52, uptime: 00:01:08, mrib
…
```

Display DR Information:

```
N3548# sh ip pim int brief
PIM Interface Status for VRF "default"
Interface            IP Address      PIM DR Address  Neighbor  Border
                                                     Count     Interface
port-channel1        40.32.1.1       40.32.1.2       1         no
port-channel2        40.32.4.1       40.32.4.4       1         no
port-channel3        40.32.21.1      40.32.21.15     1         no
port-channel4        40.32.25.1      40.32.25.17     1         no
port-channel1031     32.103.11.1     32.103.11.103   1         no
port-channel1032     32.103.12.1     32.103.12.103   1         no
port-channel1033     32.103.13.1     32.103.13.103   1         no
port-channel1034     32.103.14.1     32.103.14.103   1         no
port-channel1041     32.104.11.1     32.104.11.104   1         no
port-channel1051.1   32.105.11.1     32.105.11.105   1         no
port-channel1051.2   32.105.12.1     32.105.12.105   1         no
port-channel1051.3   32.105.13.1     32.105.13.105   1         no
port-channel1051.4   32.105.14.1     32.105.14.105   1         no
port-channel1051.5   32.105.15.1     32.105.15.105   1         no
port-channel1051.6   32.105.16.1     32.105.16.105   1         no
port-channel1051.7   32.105.17.1     32.105.17.105   1         no
port-channel1051.8   32.105.18.1     32.105.18.105   1         no
port-channel1051.9   32.105.19.1     32.105.19.105   1         no
port-channel1051.10  32.105.51.1     32.105.51.105   1         no
loopback0            40.32.0.1       40.32.0.1       0         no
loopback1            40.32.254.1     40.32.254.1     0         no
Ethernet1/1          32.101.11.1     32.101.11.101   1         no
Ethernet1/2          32.101.12.1     32.101.12.101   1         no
Ethernet1/3          32.101.13.1     32.101.13.101   1         no
Ethernet1/4          32.101.14.1     32.101.14.101   1         no
Ethernet1/5          32.101.15.1     32.101.15.101   1         no
Ethernet1/6          32.101.16.1     32.101.16.101   1         no
Ethernet1/7          32.101.17.1     32.101.17.101   1         no
Ethernet1/8          32.101.18.1     32.101.18.101   1         no
Ethernet1/9          32.102.11.1     32.102.11.102   1         no
Ethernet1/10         32.102.12.1     32.102.12.102   1         no
Ethernet1/11         32.102.13.1     32.102.13.102   1         no
```

```
Ethernet1/12         32.102.14.1     32.102.14.102   1         no
Ethernet1/13         32.102.15.1     32.102.15.102   1         no
Ethernet1/14         32.102.16.1     32.102.16.102   1         no
Ethernet1/15         32.102.17.1     32.102.17.102   1         no
Ethernet1/16         32.102.18.1     32.102.18.102   1         no
```

Displays Mroute RPF Interface and Forwarding Counters in L3 Hardware Table:

```
N3548# show forwarding multicast route group 230.32.0.1 source 132.101.11.41

  (132.101.11.41/32, 230.32.0.1/32), RPF Interface: Vlan11, flags:
    Received Packets: 5863 Bytes: 3001856
    Number of Outgoing Interfaces: 9
    Outgoing Interface List Index: 17
      Vlan12 Outgoing Packets:0 Bytes:0
      Vlan21 Outgoing Packets:0 Bytes:0
      Vlan22 Outgoing Packets:0 Bytes:0
      Vlan31 Outgoing Packets:0 Bytes:0
      Vlan32 Outgoing Packets:0 Bytes:0
      Vlan41 Outgoing Packets:0 Bytes:0
      Vlan42 Outgoing Packets:0 Bytes:0
      Vlan51 Outgoing Packets:0 Bytes:0
      Vlan52 Outgoing Packets:0 Bytes:0
```

Displays the Multicast Routing Table with Packet Counts and Bit Rates for All Sources:

```
N3548# sh ip mroute 230.32.0.1 summary
IP Multicast Routing Table for VRF "default"

Total number of routes: 511
Total number of (*,G) routes: 10
Total number of (S,G) routes: 500
Total number of (*,G-prefix) routes: 1
Group count: 10, rough average sources per group: 50.0

Group: 230.32.0.1/32, Source count: 50
Source          packets      bytes           aps   pps       bit-rate     oifs
(*,G)           8871476161   5136769483616   579   3002      14.101  mbps 10
132.101.11.41   7833         3639192         464   24        88.041  kbps 10
132.101.11.42   7637         3541584         463   20        79.746  kbps 12
132.101.11.43   7694         3438898         446   24        88.041  kbps 11
132.101.11.44   7494         3339298         445   24        79.269  kbps 10
132.101.11.45   7038         3177491         451   24        79.269  kbps 11
132.101.12.41   7809         3627240         464   24        88.041  kbps 12
132.101.12.42   7643         3544572         463   20        79.746  kbps 10
132.101.12.43   7699         3441388         446   24        87.974  kbps 11
132.101.12.44   7500         3342286         445   24        79.202  kbps 12
132.101.12.45   7031         3174005         451   24        79.202  kbps 11
132.101.21.41   7832         3638694         464   24        88.041  kbps 11
132.101.21.42   7637         3541584         463   20        79.746  kbps 12
132.101.21.43   7693         3438400         446   24        88.041  kbps 11
132.101.21.44   7491         3337804         445   24        79.269  kbps 10
132.101.21.45   7038         3177491         451   24        79.269  kbps 11
132.101.22.41   7807         3560453         456   24        79.202  kbps 11
132.101.22.42   7641         3543576         463   20        79.680  kbps 10
132.101.22.43   7698         3440890         446   24        87.974  kbps 11
132.101.22.44   7499         3341788         445   24        79.202  kbps 12
132.101.22.45   7044         3246270         460   24        87.974  kbps 11
…
```

Display IGMP Snooping Groups Information:

```
N3548# sh ip igmp snooping groups 230.32.0.1 vlan 11
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
11    230.32.0.1         v2   D     Po201
```

Displays Detected Multicast Routers for VLAN:

```
N3548# sh ip igmp snooping mrouter vlan 11
Type: S - Static, D - Dynamic, I - Internal
Type: S - Static, D - Dynamic, I - Internal
Vlan  Router-port  Type      Uptime      Expires
11    Vlan11       I         1w1d        never
11    Po200        D         1d12h       00:04:38
```

Displays IGMP Snooping Querier Information for VLAN:

```
N3548# sh ip igmp snooping querier vlan 11
Vlan  IP Address       Version    Expires     Port
11    132.101.11.2     v2         00:02:49    port-channel200
```

3.5.3 Layer-2/ Layer-3 Leaf/Access Layer Network Design Overview

vPC is not supported on Nexus 3548 running software version 6.0(2)A1(1c)

3.5.3.1.1 Spanning Tree

Multiple spanning tree protocol (MSTP) has been configured on DC32-101 and DC32-102 Leaf switches to avoid loops. VLANs 11-110 are configured on the leaf switches. Eight MST instances have been configured with 10 VLANs each. The root is configured on DC32-102 with DC32-101 as the secondary root. Eight port-channels are configured carrying 10 VLANs each.

Spanning Tree Configuration:

```
spanning-tree mode mst
spanning-tree mst 0-8 priority 28672
spanning-tree mst configuration
  instance 1 vlan 11-20
  instance 2 vlan 21-30
  instance 3 vlan 31-40
  instance 4 vlan 41-50
  instance 5 vlan 51-60
  instance 6 vlan 61-70
  instance 7 vlan 71-80
  instance 8 vlan 81-90
interface Ethernet1/45
  spanning-tree port type edge trunk
interface Ethernet1/46
  spanning-tree port type edge trunk
interface Ethernet1/47
  spanning-tree port type edge trunk
```

3.5.3.1.2 LACP

DC32 makes use of LACP mode active for all link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
N3548# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
200   Po200(SU)   Eth      LACP      Eth1/33(P)   Eth1/34(P)
201   Po201(SU)   Eth      LACP      Eth1/36(P)
202   Po202(SU)   Eth      LACP      Eth1/37(P)
203   Po203(SU)   Eth      LACP      Eth1/38(P)
204   Po204(SU)   Eth      LACP      Eth1/39(P)
205   Po205(SU)   Eth      LACP      Eth1/40(P)
206   Po206(SU)   Eth      LACP      Eth1/41(P)
207   Po207(SU)   Eth      LACP      Eth1/42(P)
208   Po208(SU)   Eth      LACP      Eth1/43(P)
1003  Po1003(RU)  Eth      LACP      Eth1/44(P)   Eth1/48(P)

N3548# show lacp interface e 1/36
Interface Ethernet1/36 is up
  Channel group is 201 port channel is Po201
  PDUs sent: 203
  PDUs rcvd: 205
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(8000, 40-55-39-26-35-c3, c8, 8000, 209), (8000, 44-3-a7-7a-b9-bc, c8
, 8000, 124)] ]
Operational as aggregated link since Wed Feb 19 17:15:52 2014

Local Port: Eth1/36   MAC Address= 44-3-a7-7a-b9-bc
  System Identifier=0x8000,44-3-a7-7a-b9-bc
  Port Identifier=0x8000,0x124
  Operational key=200
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=(Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Neighbor: 0x209
  MAC Address= 40-55-39-26-35-c3
  System Identifier=0x8000,  Port Identifier=0x8000,0x209
  Operational key=200
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=(Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
```

3.5.3.1.3 VLAN Trunking

DC32 makes use of VLAN trunking to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Vlan Configuration and Display Information:

```
vlan configuration 1,11-110
vlan 1,11-110

N3548# show vlan internal usage

VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035               Online Diagnostic
4036-4039               ERSPAN
4042                    Satellite
3968-4047,4094          Current

N3548# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po200, Po201, Eth1/33, Eth1/34
                                                Eth1/36, Eth1/45, Eth1/46
                                                Eth1/47

VLAN Type  Vlan-mode
---- ----- ----------
11   enet  CE

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------

N3548# sh int po 205 trunk

--------------------------------------------------------------------------------
Port          Native  Status        Port
              Vlan                  Channel
--------------------------------------------------------------------------------
Po205         1       trunking      --

--------------------------------------------------------------------------------
Port          Vlans Allowed on Trunk
--------------------------------------------------------------------------------
Po205         51-60

--------------------------------------------------------------------------------
Port          Vlans Err-disabled on Trunk
--------------------------------------------------------------------------------
Po205         none

--------------------------------------------------------------------------------
Port          STP Forwarding
--------------------------------------------------------------------------------
Po205         51-60

--------------------------------------------------------------------------------
Port          Vlans in spanning tree forwarding state and not pruned
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
Port          Vlans Forwarding on FabricPath
--------------------------------------------------------------------------------
Po205         none
```

3.5.3.1.4 HSRP Active/Standby

HSRP provides default gateway redundancy for hosts, ensuring that user traffic immediately and transparently recovers from first hop failures in spine layer. Preempt delay is configured to allow the router to populate its routing table before becoming the active router

HSRP configuration and verification:

N3548-1#	N3548-2#

interface Vlan52
  no shutdown
  no ip redirects
  ip address 132.101.52.3/24
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
    ip 132.101.52.1
	interface Vlan52
	  no shutdown
	  no ip redirects
	  ip address 132.101.52.2/24
	  ip pim sparse-mode
	  hsrp version 2
	  hsrp 1
	    authentication md5 key-string cisco
	    preempt delay minimum 120
	    priority 99
	 ip 132.101.52.1

```
N3548# sh hsrp summary

HSRP Summary:

Extended-hold (NSF) disabled
Global HSRP-BFD disabled

Total Groups: 100
     Version::    V1-IPV4: 0       V2-IPV4: 100    V2-IPV6: 0
       State::     Active: 100     Standby: 0       Listen: 0
       State::  V6-Active: 0    V6-Standby: 0    V6-Listen: 0

Total HSRP Enabled interfaces: 100

Total Packets:
             Tx - Pass: 78124295 Fail: 0
             Rx - Good: 62475165

Packet for unknown groups: 0

Total MTS: Rx: 1882
```

3.5.3.1.5 L2/L3 TCAM Tables

Nexus 3000/3548 platforms display MAC age as "seconds since first seen." This behavior differs from the Nexus 5000, 6000 and 7000 platforms which are displayed as "seconds since last seen" and should be taken into account when reading the table. (CSCun37474)

When topology change notifications or MAC address clears are initiated on the Nexus 3000/3548 the ARP address table also gets flushed (CSCun31859/CSCun32115). As a result, the ARP table will be re-learned.

```
DC32-102# sh mac address-table vlan 52
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
```

```
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 52       0000.061a.ee62    dynamic   18810      F    F  Po205
* 52       0084.6534.3300    dynamic   18830      F    F  Po205
* 52       0084.6534.3301    dynamic   18830      F    F  Po205
* 52       0084.6534.3302    dynamic   18830      F    F  Po205
* 52       0084.6534.3303    dynamic   18830      F    F  Po205
* 52       0084.6534.3304    dynamic   18830      F    F  Po205
* 52       0084.6534.3305    dynamic   18830      F    F  Po205
* 52       0084.6534.3306    dynamic   18830      F    F  Po205
* 52       0084.6534.3307    dynamic   18830      F    F  Po205
* 52       0084.6534.3308    dynamic   18830      F    F  Po205
* 52       0084.6534.3309    dynamic   18830      F    F  Po205
* 52       0084.6534.330a    dynamic   18830      F    F  Po205
* 52       0084.6534.330b    dynamic   18830      F    F  Po205
* 52       0084.6534.330c    dynamic   18830      F    F  Po205
* 52       0084.6534.330d    dynamic   18830      F    F  Po205
* 52       0084.6534.330e    dynamic   18830      F    F  Po205
* 52       0084.6534.330f    dynamic   18830      F    F  Po205
* 52       0084.6534.3310    dynamic   18830      F    F  Po205
* 52       0084.6534.3311    dynamic   18830      F    F  Po205…
```

Display Hardware MAC Table Entries:

```
DC32-102# show hardware internal libsdk mtc l2 mac-table-ce valid-only
    C L A S S I C A L   E T H E R N E T   M A C   T A B L E   E N T R I E S

                                                       N
                                                       O
                                                       N
                                                       E

                                                 .STA. E
                            ..LEARNED..                X
                                                   N   I
                                   A     ..CFG..   O   S
                                   G               T A T
                                   E  A              G E A
                            S     A   G      C   M N E N D
                            Y     GG  E      P   O O   T J
                            N     ER     SS    S D T Q
                            C      P  T  ET  2 W I I U E E L
                                E U   S  CATN    F F E N N 3
      V                     EE  D NS  T  UTRTC B I I I T T
  TBL L                     CC UG IE  M  RIAFP I E E E R R I  LIF            PBP
 ADDR D   FID    MAC-ADDRESS HH CE TL  P  ECPYU T D D D Y Y F  IDX DVIF  LID  LID
----- - ---- -------------- -- -- -- --- ----- - - - - - - - ---- ---- ---- ----
    0 1    0 0000:0000:0000 00 00 00   0 01000 1 0 0 0 0 0 0    0    0   51    0
   48 1   95 fc99:4752:07fc 00 00 00  80 00000 0 1 0 0 0 0 0   58    0   62    0
  192 1   92 0084:6516:3397 00 00 00  74 00000 1 0 0 0 0 0 0   55    0   59    0
  224 1   83 0084:651f:3365 00 00 00  74 00000 1 0 0 0 0 0 0   54    0   58    0
  240 1   93 0084:6515:33a3 00 00 00  74 00000 1 0 0 0 0 0 0   55    0   59    0
  308 1   82 0084:6520:3356 00 00 00  74 00000 1 0 0 0 0 0 0   54    0   58    0
  320 1   83 0084:651f:3389 00 00 00  74 00000 1 0 0 0 0 0 0   54    0   58    0
  324 1   82 0084:6520:333f 00 00 00  74 00000 1 0 0 0 0 0 0   54    0   58    0
  336 1   93 0084:6515:334f 00 00 00  74 00000 1 0 0 0 0 0 0   55    0   59    0
  352 1   92 0084:6516:337b 00 00 00  74 00000 1 0 0 0 0 0 0   55    0   59    0
  384 1   93 0084:6515:3339 00 00 00  74 00000 1 0 0 0 0 0 0   55    0   59    0
…
```

Display VLAN to BD Mapping for VLAN52:

```
DC32-102# sh hardware internal libsdk mtc vlan sw-bd-2-vlan valid-only
ADDR      BD      VLAN
----    ----      ------------
```

```
…
  57      57        57
  58      58        56
  59      59        55
  60      60        54
  61      61        53
  62      62        52
  63      63        51
  64      64        50
  65      65        49
  66      66        48
…
```

Display VLAN BDDB Hardware Programming and Flood Index for VLAN 52:

```
DC32-102# sh hardware internal libsdk mtc vlan bddb start-addr 62 end-addr 62
             C
        F    A        I
        C    C        G U
       IO    H     F  M S
       PE    E     L SP E
                K  O   O       DDNN                   A
       AA    U  I  W MSPL      RRTT      A            L                  M
       CCA   SD L  H ANT2      OOFF      G            T                  C
       LLC   EE LM   TO        PPYY  D   E
         L    S  DRU COFF            E  A             F                  F
       KK   LE  I BS HPLO      SDSD  F  GG            T  U  B            T
       EEF F2TR PIAE   OR      AAAA   D ERL           A  C  C            A
       YYC CMHO  IC  LLO             GE  PD           G                  G
         Z FP U MCLO AADF      MMMMD WF U                F  F  MC   MC
(BD) V SSS   KT C  X BB C SD S IIIIR    NSB           S  T  T FTAG FTAG  N
 TBL L EEE EEEE SEEI EEEO TPDD SSSSO EG IEA           E  A  A BASE BASE  U FLOOD
ADDR D LLN NNYR TNND LLNE LLLL SSSSP NW TLL VSAN  FID L  G  G   0    1   M INDEX
---- - --- ---- ---- ---- ---- ----- -- --- ---- ---- - -- -- ---- ---- -- -----
  62 1 100 0000 0000 0010 0000 00100 00 000   52   62 0 63 63    0    0  0 12350
```

Display BD Flood Ports for VLAN 52:

```
DC32-102# show hardware internal libsdk mtc vlan flood2bd-front-port start-addr 62 end-addr 62
V                          [[[[[[[[[[[[[[[  FRONT PORT BITMAP  ]]]]]]]]]]]]]]]
L ****      FLOOD CP2  CPU        4        3        2        1
D VLAN   BD INDEX CPU CODE        1        3        5        7        9        1
- ---- ---- ----- --- ---- -------- -------- -------- -------- -------- --------
1   52   62 12350   0    0 00000000 10000011 00000000 00000000 00000000 00000000
```

Display CBL for VLAN 52:

```
DC32-102# show hardware internal libsdk mtc vlan cbl-front-port start-addr 62 end-addr 62

D = Disable, B = blocking, L = learning, F = forwarding

V           C  [[[[[[[[[[[[[[[[[  FRONT PORT CBL  ]]]]]]]]]]]]]]]]]
L ****      P         4        3        2        1
D VLAN   BD U         1        3        5        7        9        1
- ---- ---- - -------- -------- -------- -------- -------- --------
1   52   62 B BBBBBBBB FBBBBBFF BBBBBBBB BBBBBBBB BBBBBBBB BBBBBBBB
```

Display Egress SVI-BD, Front Ports for VLAN 52:

```
DC32-102# show hardware internal libsdk mtc vlan svi-bd-front-port start-addr 62 end-addr 62
MAC SA index: router mac index 0-3
Egr Cbl Drop: 0: dont drop, 1:drop
QinQ type   : 0:no qtag, 1:one qtag; 2: two qtag
Customer ID : Qtag insert into packet

. .... .... . . . .... .. .... [[[[[   F R O N T   P O R T   01-04   ]]]]]
```

```
                               [[[[[[   F R O N T   P O R T   05-08   ]]]]]]
                               [[[[[[   F R O N T   P O R T   09-12   ]]]]]]
                               [[[[[[   . . . . . . . . . .   . . .   ]]]]]]
                               [[[[[[   F R O N T   P O R T   45-48   ]]]]]]
                               |N E       |N E       |N E       |N E       |
                               |A G     C |A G     C |A G     C |A G     C |
              M                |T R     U |T R     U |T R     U |T R     U |
              A                |    Q   S |    Q   S |    Q   S |    Q   S |
              C         F      |O C I   T |O C I   T |O C I   T |O C I   T |
            C   N       T      |U B N   O |U B N   O |U B N   O |U B N   O |
            E S O       A   T  |T L Q   M |T L Q   M |T L Q   M |T L Q   M |
              A         G  ME  |        E |        E |        E |        E |
            M   S          UN  |S D T   R |S D T   R |S D T   R |S D T   R |
V           O I T       M  LA  |I R Y     |I R Y     |I R Y     |I R Y     |
L ****      D D A FTAG  O  TNI |D O P   I |D O P   I |D O P   I |D O P   I |
D VLAN  BD  E X T BASE  D  ITD |E P E   D |E P E   D |E P E   D |E P E   D |
- ---- ---- - - - ---- -- ---- +- - - ----+- - - ----+- - - ----+- - - ----+
1   52   62 1 0 0 1023  0    0
                     PORT 01-04|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 05-08|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 09-12|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 13-16|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 17-20|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 21-24|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 25-28|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 29-32|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 33-36|0 0 1   52|0 0 1   52|0 1 0 4095|0 1 0 4095|
                     PORT 37-40|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 0 1   52|
                     PORT 41-44|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
                     PORT 45-48|0 1 0 4095|0 1 0 4095|0 1 0 4095|0 1 0 4095|
```

3.6 DC33

3.6.1 Configuration of Platform Specific Features On DC36

3.6.1.1 Licensing

License Usage on Nexus 3000 in DC33:

```
DC33-1# sh license usage
Feature                        Ins  Lic   Status Expiry Date Comments
                                   Count
--------------------------------------------------------------------------------
LAN_BASE_SERVICES_PKG           Yes   -   In use Never        -
ALGO_BOOST_SERVICES_PKG         No    -   Unused              -
LAN_ENTERPRISE_SERVICES_PKG     Yes   -   In use Never        -
--------------------------------------------------------------------------------
```

Although features can be enabled and configured in the CLI without licenses, they will not function until the license is installed.

3.6.1.2 Out-of-Band Management Network

DC33 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic.

Configuration:

```
interface mgmt0
  vrf member management
  ip address 10.2.33.1/16
```

3.6.1.3 Common Configurations

3.6.1.3.1 SSH and TACACS+

SSH is enabled in DC33 to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration and Verification:

```
feature tacacs+


ip tacacs source-interface mgmt0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
    use-vrf management


DC33-1# sh ssh server
ssh version 2 is enabled
DC33-1# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
interop  pts/0        Feb 10 11:37   .          3995 (taro.interop.cisco.com) session=ssh *
```

3.6.1.3.2 CDP and LLDP

CDP and LLDP are pervasively used on the DC33 testbed for inter-device discovery.

Configuration and Verification:

```
DC33-1# sh run cdp all

version 6.0(2)U1(3)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
cdp format device-id system-name


interface mgmt0
  cdp enable

interface Ethernet1/1
  cdp enable

<TRUNCATED>

interface Ethernet1/52
  cdp enable

DC33-1# sh cdp nei
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute


Device-ID             Local Intrfce Hldtme Capability  Platform      Port ID
mgmt-sw3.interop.cisco.com
                    mgmt0          157    R S I     WS-C6504-E    Gig3/2
DC33-101.interop.cisco.com(FOC1711R1GX)
                    Eth1/1         161    R S I s   N3K-C3048TP-1 Eth1/1
DC33-101.interop.cisco.com(FOC1711R1GX)
                    Eth1/2         161    R S I s   N3K-C3048TP-1 Eth1/2
DC33-101.interop.cisco.com(FOC1711R1GX)
                    Eth1/3         162    R S I s   N3K-C3048TP-1 Eth1/3
DC33-101.interop.cisco.com(FOC1711R1GX)
                    Eth1/4         162    R S I s   N3K-C3048TP-1 Eth1/4
<TRUNCATED>

DC33-1# sh run lldp all

feature lldp

lldp holdtime 120
lldp reinit 2
lldp timer 30
lldp tlv-select port-description
lldp tlv-select system-name
lldp tlv-select system-description
lldp tlv-select system-capabilities
lldp tlv-select management-address
lldp tlv-select dcbxp
lldp tlv-select port-vlan

interface mgmt0
  lldp transmit
  lldp receive

interface Ethernet1/1
  lldp transmit
  lldp receive

<TRUNCATED>

interface Ethernet1/52
  lldp transmit
```

```
  lldp receive

DC33-1# sh lldp nei
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID              Local Intf      Hold-time  Capability  Port ID
DC33-101.interop.cisco.com Eth1/1         120        BR          Eth1/1

DC33-101.interop.cisco.com Eth1/2         120        BR          Eth1/2

DC33-101.interop.cisco.com Eth1/3         120        BR          Eth1/3

DC33-101.interop.cisco.com Eth1/4         120        BR          Eth1/4

DC33-101.interop.cisco.com Eth1/5         120        BR          Eth1/5
```

3.6.1.3.3 Syslog

Syslog is used to record all network events on the DC33 test bed. Whenever possible, DC33 makes use of a separate management VRF for syslog.

Configuration and Verification:

```
logging server syslog.interop.cisco.com 5 use-vrf management facility local6

DC33-1# sh logging server
Logging server:                 enabled
{syslog.interop.cisco.com}
        server severity:        notifications
        server facility:        local6
        server VRF:             management
```

3.6.1.3.4 SNMP

SNMP is used for system monitoring in DC33. Scripts are used to poll the systems asynchronously during the course of all DC33 test execution.

Configuration:

```
version 6.0(2)U1(3)
snmp-server user admin network-admin auth md5 0xeeea7f7d446b7958c520b61b33df1cbd
 priv 0xeeea7f7d446b7958c520b61b33df1cbd localizedkey
snmp-server community cisco group network-operator
snmp-server community private group network-admin
snmp-server community public group network-operator
```

3.6.1.3.5 NTP

NTP is used to synchronize the clocks on all DC33 devices to provide consistent timestamps on all network logs and events.

Configuration and Verification:

```
ntp distribute
ntp server 172.28.92.1
ntp commit

DC33-1# sh ntp status
Distribution : Enabled
Last operational state: No session
```

```
DC33-1# sh ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                st   poll   reach delay   vrf
-------------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                8    64      377   0.00092 management
```

3.6.1.3.6 SPAN

SPAN has been enabled on DC33 switches to provide packet captures to assist in network debugging.

Configuration and Verification:

```
monitor session 1
  source interface port-channel11 both
  destination interface Ethernet1/50
  no shut

DC33-1# sh monitor session 1
   session 1
---------------
type              : local
state             : up
acl-name          : acl-name not specified
source intf       :
    rx            : Po11
    tx            : Po11
    both          : Po11
source VLANs      :
    rx            :
destination ports : Eth1/50

Legend: f = forwarding enabled, l = learning enabled
```

3.6.1.3.7 DNS

DNS has been enabled to provide name lookup in DC33 network.

Configuration and Verification:

```
vrf context management
  ip domain-lookup
  ip domain-name interop.cisco.com
  ip domain-list cisco.com
  ip domain-list interop.cisco.com
  ip name-server 172.28.92.9 172.28.92.10


DC33-1# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=1.631 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=1.754 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=1.578 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=1.409 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=1.374 ms

--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.374/1.549/1.754 ms
```

3.6.1.3.8 MTU

In the Nexus 3000 routers, in order to configure the MTU to handle jumbo frames the following policy-map has to be applied.

Configuration and Verification:

```
  policy-map type network-qos jumbo
    class type network-qos class-default
      match qos-group 0
      mtu 9216

DC33-1# sh policy-map type network-qos jumbo


  Type network-qos policy-maps
  ===============================

  policy-map type network-qos jumbo
    class type network-qos class-default
      mtu 9216

DC33-1# sh queuing interface ethernet 1/1
Ethernet1/1 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0       WRR            100

  RX Queuing
    qos-group 0
    HW MTU: 9216 (9216 configured)
    drop-type: drop, xon: 0, xoff: 0
    Statistics:
        Ucast pkts sent over the port           : 581711504
        Ucast bytes sent over the port          : 812069209800
        Mcast pkts sent over the port           : 129918846
        Mcast bytes sent over the port          : 181366709016
        Ucast pkts dropped                      : 0
        Ucast bytes dropped                     : 0
        Mcast pkts dropped                      : 0
        Mcast bytes dropped                     : 0

DC33-1# sh int e1/1
Ethernet1/1 is up
 Dedicated Interface
  Hardware: 10/100/1000 Ethernet, address: b0fa.eb5f.dc7c (bia b0fa.eb5f.dc28)
  Internet Address is 33.101.11.1/24
  MTU 9216 bytes, BW 1000000 Kbit, DLY 10 usec
```

3.6.1.4 Debugging on the Broadcom Shell

Nexus 3000 offers a very powerful tool that allows an easy access to the Broadcom shell. This allows to access to a big variety of commands hence enhancing the debug capabilities of the chipset. These commands should be used with caution as they are backdoors to program the hardware and bypass NX-OS.

Accessing the Broadcom Shell:

```
DC33-102# test hardware internal bcm-usd bcm-diag-shell
Available Unit Numbers: 0
bcm-shell.0> help
Help: Type help "command" for detailed command usage
Help: Upper case letters signify minimal match

Commands common to all modes:
```

```
        ?                   Display list of commands
        ASSert              Assert
        BackGround          Execute a command in the background.
        BCM                 Set shell mode to BCM.
        BCMX                Set shell mode to BCMX.
        break               place to hang a breakpoint
        CASE                Execute command based on string match
        CD                  Change current working directory
        cint                Enter the C interpreter
        CONFig              Configure Management interface
        CONSole             Control console options
<TRUNCATED>
```

3.6.1.5 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU.

The default PIMREG CoPP is 200pps. The PIMREG CoPP configuration at the multicast RP determines the rate of PIM source registration and periodic null-registers that can be processed. The PIMREG CoPP at the RP should be adjusted accordingly to accommodate the registration rates to prevent potential mroute states from timing out.

For example, there are 2000 active sources on the DC33 testbed, with bursts of 500 requiring registration. Testing found that a CoPP of 1000pps was adequate to accommodate this number of multicast sources and burst pattern.

The remaining values are kept to their default values.

Configuration of CoPP on Nexus 3000 Software Release 6.0(2)U1(3) as Used in DC33:

```
policy-map type control-plane copp-system-policy
  class copp-s-selfIp
    police pps 500
  class copp-s-default
    police pps 400
  class copp-s-l2switched
    police pps 200
  class copp-s-ping
    police pps 100
  class copp-s-l3destmiss
    police pps 100
  class copp-s-glean
    police pps 500
  class copp-s-l3mtufail
    police pps 100
  class copp-s-ttl1
    police pps 100
  class copp-s-ipmcmiss
    police pps 400
  class copp-s-l3slowpath
    police pps 100
  class copp-s-dhcpreq
    police pps 300
  class copp-s-dhcpresp
    police pps 300
  class copp-s-dai
    police pps 300
  class copp-s-igmp
    police pps 400
  class copp-s-routingProto2
    police pps 1300
```

```
  class copp-s-v6routingProto2
    police pps 1300
  class copp-s-eigrp
    police pps 200
  class copp-s-pimreg
    police pps 1000
  class copp-s-pimautorp
    police pps 200
  class copp-s-routingProto1
    police pps 1000
  class copp-s-arp
    police pps 200
  class copp-s-ptp
    police pps 1000
  class copp-s-bfd
    police pps 350
  class copp-s-bpdu
    police pps 12000
  class copp-s-dpss
    police pps 1000
  class copp-icmp
    police pps 200
  class copp-telnet
    police pps 500
  class copp-ssh
    police pps 500
  class copp-snmp
    police pps 500
  class copp-ntp
    police pps 100
  class copp-tacacsradius
    police pps 400
  class copp-stftp
    police pps 400
  class copp-s-vxlan
    police pps 1000
```

3.6.1.6 ECMP for IPv4 and IPv6 host routes

ECMP support for host routes is disabled by default on the Nexus 3000 switches. On DC33 ECMP for host routes is enabled to program all unicast host routes into the longest-prefix match algorithm (LPM) table. ECMP for host routes is provided in the switch hardware.

```
DC33-1# sh run all | inc profile

hardware profile multicast max-limit 6000
no hardware profile multicast prefer-source-tree
hardware profile unicast enable-host-ecmp
hardware profile multicast syslog-threshold 90
hardware profile unicast syslog-threshold 90
hardware profile unicast enable-host-ecmp ipv4
hardware profile unicast enable-host-ecmp ipv6
no hardware profile unicast enable-host-ecmp arp-nd
no hardware profile unicast enable-host-ecmp ipv4 arp
no hardware profile unicast enable-host-ecmp ipv6 nd

DC33-1# show hardware profile status
Total LPM Entries = 8191.
Total Host Entries = 16384.
Reserved LPM Entries = 1024.
Max Host4/Host6 Limit Entries (shared)=  4384/2192*
Max Mcast Limit Entries = 6000.
Used LPM Entries (Total) = 724.
Used IPv4 LPM Entries =  377.
Used IPv6 LPM Entries =  347.
```

```
Used IPv6 LPM_128 Entries =  13.
Used Host Entries in LPM (Total) = 245.
Used Host4 Entries in LPM = 188.
Used Host6 Entries in LPM = 57.
Used Mcast Entries = 889.
Used Mcast OIFL Entries = 61.
Used Host Entries in Host (Total) = 0.
Used Host4 Entries in Host = 0.
Used Host6 Entries in Host = 0.
MFIB prefer-source-tree = Disabled/0/0.

*Unicast Host Table is in shared mode b/n v4 & v6...
```

3.6.1.6.1 Repartition of the TCAM for Multicast Entries

Figure 33 Nexus 3000 – L3 TCAM Allocation

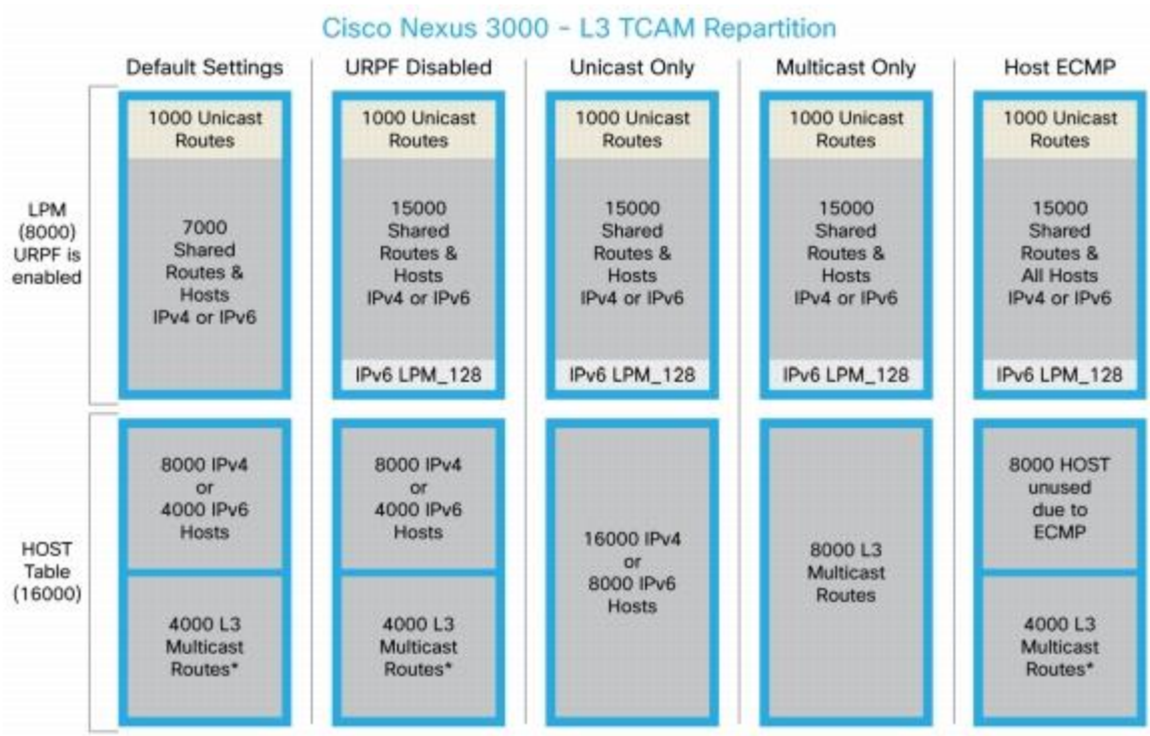

As shown above, the default TCAM allocation for multicast routes is 4000 multicast entries.

In order to accommodate all of the multicast entries deployed on DC33, the TCAM table had to be repartitioned in order to increase the allocated region for multicast entries to 6000.

```
DC33-1# sh run all | inc profile

hardware profile multicast max-limit 6000
no hardware profile multicast prefer-source-tree
DC33-1# show hardware profile status
Total LPM Entries = 8191.
Total Host Entries = 16384.
Reserved LPM Entries = 1024.
Max Host4/Host6 Limit Entries (shared)=  4384/2192*
Max Mcast Limit Entries = 6000.
```

```
Used LPM Entries (Total) = 724.
Used IPv4 LPM Entries =  377.
Used IPv6 LPM Entries =  347.
Used IPv6 LPM_128 Entries =  13.
Used Host Entries in LPM (Total) = 245.
Used Host4 Entries in LPM = 188.
Used Host6 Entries in LPM = 57.
Used Mcast Entries = 889.
Used Mcast OIFL Entries = 61.
Used Host Entries in Host (Total) = 0.
Used Host4 Entries in Host = 0.
Used Host6 Entries in Host = 0.
MFIB prefer-source-tree = Disabled/0/0.

*Unicast Host Table is in shared mode b/n v4 & v6...
```

3.6.1.6.2 Preventing Multicast Packet Duplication

On a network topology with anycast RP where multicast sources and receivers are on the same switch, the PIM RP may forward packets back toward the source DR due to the presense of receivers that joined to the (*,G). Because the source is also present, the DR has both (*,G) and (S,G) created for the local sources. The DR is expected to forward packets that match these sources using only the (S,G). The N3K will forward packets using the (*,G) also – therefore, causing duplicate packets to be send to the receiver (CSCub70536).

As a workaround, on DC33, the following command has been tested to prevent duplicate packets when both (S,G) and (*,G) with different RPF interfaces are on the switch.

```
hardware profile multicast prefer-source-tree eternity
```

When this command is used, the switch supports source (S, G) route injections at a slower rate which will cause slower switchover from shared to source tree. The multicast routing table must have at least 500 entries free for source (S, G) routes.

In the vPC topology with Nexus 3048, some multicast packet duplication might still be seen under certain conditions described in CSCul14373.

3.6.2 Image Upgrade and Downgrade

On the DC33 testbed both “install all” and “reload” commands have been used to upgrade/downgrade software images.

```
DC33-1# install all kickstart bootflash:n3000-uk9-kickstart.6.0.2.U2.0.8.bin system bootflash:n3000-uk9.6.0.2.U2.0.8.bin

Verifying image bootflash:/n3000-uk9-kickstart.6.0.2.U2.0.8.bin for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
```

```
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/n3000-uk9-kickstart.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive          none



Images will be upgraded according to following table:
Module             Image         Running-Version             New-Version  Upg-Required
------  ----------------  ----------------------  ----------------------  ------------
     1            system             6.0(2)U2(1)             6.0(2)U2(1)            no
     1         kickstart             6.0(2)U2(1)             6.0(2)U2(1)            no
     1              bios      v2.5.0(06/27/2013)      v2.5.0(06/27/2013)            no
     1         power-seq                    v4.1                    v4.1            no

Additional info for this installation:
--------------------------------------

Service "fwm" : vPC is L3 enabled.  Upgrade needs to be disruptive.



Do you want to continue with the installation (y/n)?  [n]y

Switch will be reloaded for disruptive upgrade.

Install is in progress, please wait.
Performing runtime checks.
SUCCESS
Setting boot variables.
SUCCESS
Performing configuration copy.
SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

NVT recommends saving the configuration prior to any image upgrade/downgrade and comparing the configurations before and after to ensure a successful migration. In some situations some differences might be observed (CSCul45536, CSCuj74966).

In order to restore the proper configuration from such situations as well as any PSS corruption, execute the following procedure:

1. Change/Check boot variable;

2. Write memory;
3. Write erase and reload;
4. copy <device-storage>:<saved-config> running-config;
5. Change boot variable;
6. Write memory;
7. Reload [optional]

The third step will cause the Nexus 3000 to be set to the factory default values.

3.6.3 Routing Design Overview

3.6.3.1 Unicast Routing Design

The network is split into three layers: core, spine and leaf. The layers are logically connected to each other through eBGP, as shown in Figure 34. The N7K core layer in BGP AS 3 is shared with other DC3 networks (DC31, DC32, and DC36). The spine layer runs OSPF to provide inter-switch connectivity to support iBGP sessions. The leaf layer is divided into multiple BGP ASes. This BGP logical design is easier to configure, maintain and debug than full mesh ibgp, route reflector, or confederations; the core can consolidate these as private ASes if there is a need to advertise to other BGP exchanges.

The spine layer is eBGP connected to the ASes configured at the Leaf layer over both IPv4 and IPv6 address families (eBGP dual stack). The spine routers also inject the default route down to the leaf ASes for both IPv4 and IPv6 address families (default-originate). ECMP is enabled on both IPv4 and IPv6 address families (maximum-path 32) across the DC33 network.

The leaf layer represents different top of rack topologies that can be deployed. AS 33101 employs two Nexus 3048 in a vPC topology, using HSRP for gateway redundancy for nodes. AS 33103 employs a routed top of rack with N3048. AS 33104 is used as a test tool rather than network under test. The Catalyst 6500 is divided into multiple VRFs, with each VRF representing an extra ToR in the network. The goal is to test increasing number of ToR supported by the spine layer.

Figure 34 DC33 BGP Logical Design

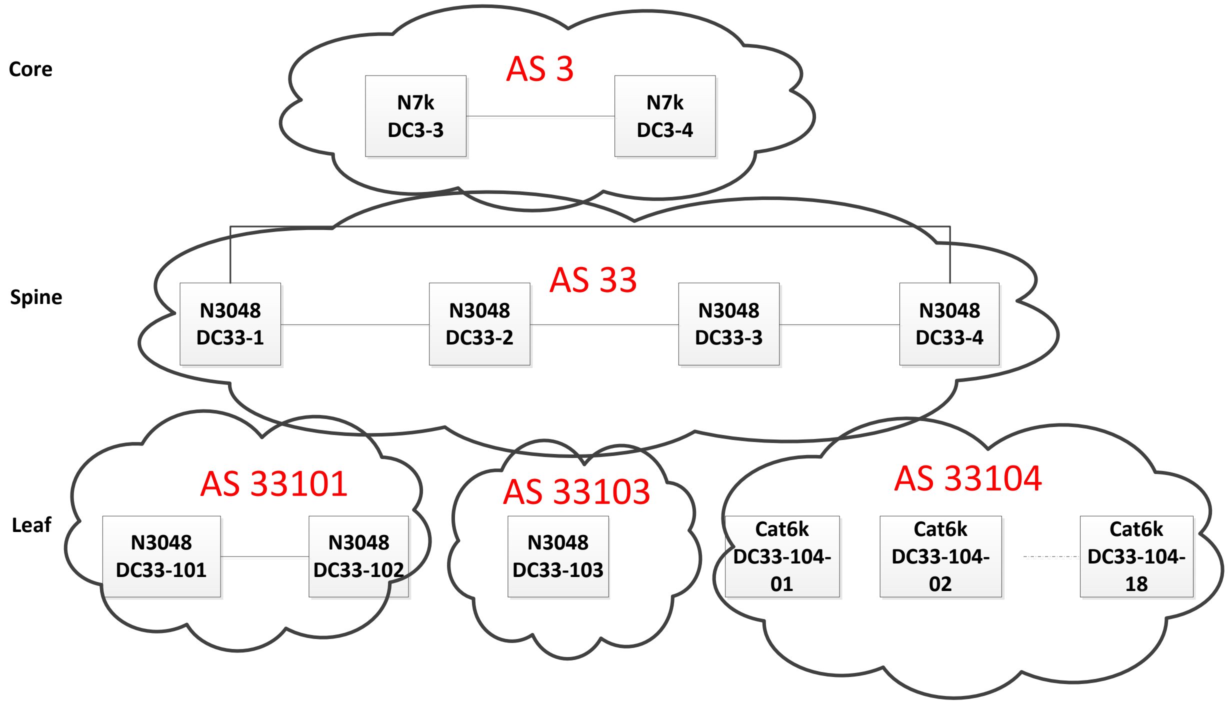

BGP peer templates are used to simplify configuration.

BGP Spine Router Configuration:

```
feature bgp
router bgp 33
  router-id 40.33.0.1
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
    network 33.1.11.1/32
<TRUNCATED>
    network 33.1.28.1/32
    network 33.101.11.0/24
<TRUNCATED>
    network 33.103.18.0/24
    network 33.114.1.0/24
<TRUNCATED>
    network 33.114.18.0/24
    network 40.33.0.1/32
    network 40.33.1.0/24
    network 40.33.4.0/24
    network 40.33.31.0/24
    network 40.33.41.0/24
    network 40.33.254.1/32
    maximum-paths 64
  address-family ipv6 unicast
    network 2001:1:33:1:11:1::1/128
<TRUNCATED>
    network 2001:1:33:1:28:1::1/128
    network 2001:1:40:33::1:0:1/128
    network 2001:1:40:33:31::/80
    network 2001:1:40:33:41::/80
    network 2001:33:101:11::/64
<TRUNCATED>
```

```
    network 2001:33:114:18::/64
    network 2001:33:114:1::/64
<TRUNCATED>
    network 2001:33:114:9::/64
    maximum-paths 64
  template peer BGPLEAF
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      default-originate
      soft-reconfiguration inbound
    address-family ipv6 unicast
      default-originate
      soft-reconfiguration inbound
  neighbor 33.101.11.101 remote-as 33101
    inherit peer BGPLEAF
<TRUNCATED>
  neighbor 33.102.18.102 remote-as 33101
    inherit peer BGPLEAF
  neighbor 33.103.11.103 remote-as 33103
    inherit peer BGPLEAF
  neighbor 33.103.12.103 remote-as 33103
    inherit peer BGPLEAF
  neighbor 33.114.1.104 remote-as 104000
    inherit peer BGPLEAF
    address-family ipv4 unicast
    address-family ipv6 unicast
  neighbor 33.114.2.104 remote-as 104000
    inherit peer BGPLEAF
    address-family ipv4 unicast
    address-family ipv6 unicast
  <TRUNCATED>
  neighbor 33.114.18.104 remote-as 104000
    inherit peer BGPLEAF
    address-family ipv4 unicast
    address-family ipv6 unicast
  neighbor 40.33.0.2 remote-as 33
    password 3 a667d47acc18ea6b
    update-source loopback0
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 40.33.0.3 remote-as 33
    password 3 a667d47acc18ea6b
    update-source loopback0
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 40.33.0.4 remote-as 33
    password 3 a667d47acc18ea6b
    update-source loopback0
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 40.33.31.15 remote-as 3
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
```

```
  neighbor 40.33.41.17 remote-as 3
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
```

BGP Leaf router configurations:

```
feature bgp
router bgp 33101
  router-id 33.0.0.102
  log-neighbor-changes
  address-family ipv4 unicast
    network 33.0.0.102/32
    network 33.102.11.0/24
<TRUNCATED>
    network 33.102.48.0/24
    network 133.101.1.0/24
    network 133.101.11.0/24
<TRUNCATED>
    network 133.101.110.0/24
    maximum-paths 64
  address-family ipv6 unicast
    network 2001:133:101:100::/64
<TRUNCATED>
    network 2001:133:101:110::/64
    network 2001:133:101:11::/64
<TRUNCATED>
    network 2001:133:101:19::/64
    network 2001:133:101:1::/64
    network 2001:133:101:20::/64
<TRUNCATED>
    network 2001:133:101:99::/64
    network 2001:1:33::102:0:102/128
    maximum-paths 64
  template peer BGPLEAF
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 33.102.11.1 remote-as 33
    inherit peer BGPLEAF
<TRUNCATED>
  neighbor 33.102.48.4 remote-as 33
    inherit peer BGPLEAF
  neighbor 133.101.1.2 remote-as 33101
    inherit peer BGPLEAF
```

```
router bgp 104000
 bgp router-id 33.0.0.104
 bgp log-neighbor-changes
 maximum-paths 32
 !
 address-family ipv4 vrf 104001
  network 33.1.1.104 mask 255.255.255.255
  network 33.104.1.0 mask 255.255.255.0
  network 33.114.1.0 mask 255.255.255.0
  network 33.124.1.0 mask 255.255.255.0
  network 33.134.1.0 mask 255.255.255.0
  network 33.144.1.0 mask 255.255.255.0
  network 133.104.1.0 mask 255.255.255.0
```

```
 neighbor 33.114.1.1 remote-as 33
 neighbor 33.114.1.1 password 3 cisco123
 neighbor 33.114.1.1 activate
 neighbor 33.114.1.1 soft-reconfiguration inbound
 neighbor 33.124.1.2 remote-as 33
 neighbor 33.124.1.2 password 3 cisco123
 neighbor 33.124.1.2 activate
 neighbor 33.124.1.2 soft-reconfiguration inbound
 neighbor 33.134.1.3 remote-as 33
 neighbor 33.134.1.3 password 3 cisco123
 neighbor 33.134.1.3 activate
 neighbor 33.134.1.3 soft-reconfiguration inbound
 neighbor 33.144.1.4 remote-as 33
 neighbor 33.144.1.4 password 3 cisco123
 neighbor 33.144.1.4 activate
 neighbor 33.144.1.4 soft-reconfiguration inbound
 maximum-paths 32
exit-address-family
!
address-family ipv6 vrf 104001
 maximum-paths 32
 neighbor 33.114.1.1 remote-as 33
 neighbor 33.114.1.1 password 3 cisco123
 neighbor 33.114.1.1 activate
 neighbor 33.114.1.1 soft-reconfiguration nbound
 neighbor 33.124.1.2 remote-as 33
 neighbor 33.124.1.2 password 3 cisco123
 neighbor 33.124.1.2 activate
 neighbor 33.124.1.2 soft-reconfiguration inbound
 neighbor 33.134.1.3 remote-as 33
 neighbor 33.134.1.3 password 3 cisco123
 neighbor 33.134.1.3 activate
 neighbor 33.134.1.3 soft-reconfiguration inbound
 neighbor 33.144.1.4 remote-as 33
 neighbor 33.144.1.4 password 3 cisco123
 neighbor 33.144.1.4 activate
 neighbor 33.144.1.4 soft-reconfiguration inbound
exit-address-family
```

3.6.3.1.1.1 BGP Router-ID

To establish BGP sessions between peers, BGP must have a router ID, which is sent to BGP peers in the OPEN message when a BGP session is established. On DC33, NVT has configured a loopback interface IP address as the BGP router-ID. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

To Verify the BGP Router-ID:

```
DC33-1# sh ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 59144, local router ID is 40.33.0.1
```

3.6.3.1.1.2 BGP Address Family

BGP address family for IPv4 and Ipv6 have been configured to achieve BGP peering, load-balancing, default route injection.

To Verify the BGP Address Family:

```
DC33-1# sh ip bgp all sum
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 40.33.0.1, local AS number 33
BGP table version is 28188, IPv4 Unicast config peers 41, capable peers 41
426 network entries and 3915 paths using 227436 bytes of memory
BGP attribute entries [10/1360], BGP AS path entries [4/24]
BGP community entries [0/0], BGP clusterlist entries [0/0]
3849 received paths for inbound soft reconfiguration
3849 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
33.101.11.101   4 33101    3287   19988    28188    0    0    2d06h 167
33.101.12.101   4 33101    3287   19985    28188    0    0    2d06h 167
<TRUNCATED>
33.103.12.103   4 33103    3286   19891    28188    0    0    2d06h 109
33.114.1.104    4 104000
                          12895    3489    28188    0    0    2d06h 6
<TRUNCATED>
33.114.18.104   4 104000
                           3712    3487    28188    0    0    2d06h 6
40.33.0.2       4    33   12817   12817    28188    0    0    2d06h 283
40.33.0.3       4    33    3489   12817    28188    0    0    2d06h 283
40.33.0.4       4    33   12809   12817    28188    0    0    2d06h 283
40.33.31.15     4     3    3278   19566    28188    0    0    2d06h 1
40.33.41.17     4     3    3281   19545    28188    0    0    2d06h 1

BGP summary information for VRF default, address family IPv6 Unicast
BGP router identifier 40.33.0.1, local AS number 33
BGP table version is 1138, IPv6 Unicast config peers 41, capable peers 41
385 network entries and 2663 paths using 164656 bytes of memory
BGP attribute entries [8/1088], BGP AS path entries [3/18]
BGP community entries [0/0], BGP clusterlist entries [0/0]
2600 received paths for inbound soft reconfiguration
2600 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
33.101.11.101   4 33101    3287   19988     1138    0    0    2d06h 103
33.101.12.101   4 33101    3287   19985     1138    0    0    2d06h 103
<TRUNCATED>
33.103.12.103   4 33103    3286   19891     1138    0    0    2d06h 109
33.114.1.104    4 104000
                          12895    3489     1138    0    0    2d06h 0
33.114.2.104    4 104000
                           3839    3489     1138    0    0    2d06h 0
<TRUNCATED>
33.114.18.104   4 104000
                           3712    3487     1138    0    0    2d06h 0
40.33.0.2       4    33   12817   12817     1138    0    0    2d06h 244
40.33.0.3       4    33    3489   12817     1138    0    0    2d06h 244
40.33.0.4       4    33   12809   12817     1138    0    0    2d06h 244
40.33.31.15     4     3    3278   19566     1138    0    0    2d06h 1
40.33.41.17     4     3    3281   19545     1138    0    0    2d06h 1
```

3.6.3.1.1.3 BGP Load Sharing and ECMP

DC33 has configured the maximum-paths that BGP adds to the route table for equal-cost multipath load balancing as 32 for both spine and leaf peers for IPv4/IPv6 address families.

3.6.3.1.1.4 BGP Authentication

DC33 has configured MD5 Authentication for BGP sessions.

To Verify the BGP Authentication:

```
DC33-1# sh ip bgp neighbors 40.33.0.2
BGP neighbor is 40.33.0.2,  remote AS 33, ibgp link,  Peer index 1
  BGP version 4, remote router ID 40.33.0.2
  BGP state = Established, up for 2d06h
  Using loopback0 as update source for this peer
  TCP MD5 authentication is enabled
```

3.6.3.1.1.5 BGP Update-Source

DC33 has configured BGP update-source to establish a BGP multi-hop sessions. DC33 has multi-hop sessions only on the iBGP peering between the spine switches.

To Verify the BGP Update-Source:

```
DC33-1# sh ip bgp neighbors 40.33.0.2
BGP neighbor is 40.33.0.2,  remote AS 33, ibgp link,  Peer index 1
  BGP version 4, remote router ID 40.33.0.2
  BGP state = Established, up for 2d06h
  Using loopback0 as update source for this peer
```

3.6.3.1.1.6 BGP Default Route

BGP default route is advertised from the spine peers to the leaf peers for both Ipv4 and Ipv6 address families.

To Verify the BGP Default Route:

```
DC33-1# sh ip bgp neighbors 33.114.18.104 | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 59144, neighbor version 59144
  8 accepted paths consume 416 bytes of memory
  480 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1
  Default information originate, default sent

  For address family: IPv6 Unicast
  BGP table version 32935, neighbor version 0
  0 accepted paths consume 0 bytes of memory
  0 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1
  Default information originate, default sent
```

3.6.3.1.1.7 BGP Next-Hop-Self

BGP next-hop-self is configured for iBGP sessions between the spine switches for both IPv4 and IPv6 address families.

To Verify the BGP Next-Hop-Self:

```
DC33-1# sh ip bgp neighbors 33.114.18.104 | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 59144, neighbor version 59144
  8 accepted paths consume 416 bytes of memory
  480 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1
  Default information originate, default sent

  For address family: IPv6 Unicast
  BGP table version 32935, neighbor version 0
  0 accepted paths consume 0 bytes of memory
  0 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1
  Default information originate, default sent
```

3.6.3.1.1.8 BGP Soft-Reconfiguration

BGP Soft reset is recommended because it allows routing tables to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis.

```
DC33-1# sh ip bgp neighbors 33.114.18.104 | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 59144, neighbor version 59144
  8 accepted paths consume 416 bytes of memory
  480 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1

  For address family: IPv6 Unicast
  BGP table version 32935, neighbor version 0
  0 accepted paths consume 0 bytes of memory
  0 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 33.114.18.1
```

3.6.3.1.2 OSPF/OSPFv3 Routing Design

At the spine layer (AS 33), OSPF/OSPFv3 is used as an IGP to grant reachability within the AS itself. OSPF router-ID and MD5 area authentication are enabled. The OSPF process is enable d only on directly connected interfaces and the Loopback interface. All the OSPF enabled interfaces are in Area 0.0.0.0. Each OSPF network type is set to point-to-point to decrease OSPF neighbor setup latency. In order to improve OSPF convergence, SPF and LSA timers are throttled to (100 200 5000 and 50 100 300) respectively.

OSPF Router Configuration:

```
feature ospf

router ospf 3
  router-id 40.33.0.1
  area 0.0.0.0 authentication message-digest
  log-adjacency-changes
  timers throttle spf 100 200 5000
  timers throttle lsa 50 100 300

interface loopback0
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip router ospf 3 area 0.0.0.0
```

```
interface port-channel1
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip ospf network point-to-point
  ip router ospf 3 area 0.0.0.0

interface port-channel2
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip ospf network point-to-point
  ip router ospf 3 area 0.0.0.0
```

OSPFv3 router configuration:

```
feature ospfv3

router ospfv3 33
  router-id 40.33.0.1
  log-adjacency-changes detail

interface loopback0
  ipv6 router ospfv3 33 area 0.0.0.0

interface port-channel1
  ospfv3 network point-to-point
  ipv6 router ospfv3 33 area 0.0.0.0

interface port-channel2
  ospfv3 network point-to-point
  ipv6 router ospfv3 33 area 0.0.0.0
```

3.6.3.1.3 Unicast Forwarding Verification

On NX-OS platforms, routing is performed using hardware forwarding engines. The following sequence of commands illustrates verification of the programming of a host on a directly connected subnet on the Nexus 3000.

This Switch is the Authoritative Router for a Directly Connected Subnet on VLAN 11 133.101.11.0/24:

```
DC33-102# show running-config interface vlan 11

version 6.0(2)U1(3)

interface Vlan11
  no shutdown
  mtu 9216
  no ip redirects
  ip address 133.101.11.3/24
  ipv6 address 2001:133:101:11::3/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 133.101.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 2001:133:101:11::1
```

The Host 133.101.52.51 has been Learned via ARP on this Subnet.

```
DC33-102# sh ip arp 133.101.52.51

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
133.101.52.51   00:04:38  0085.6534.3300  Vlan52
```

On NX-OS, “show ip route” will also Show Directly Connected Hosts as /32 Routes:

```
DC33-102# sh ip route 133.101.52.51
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

133.101.52.51/32, ubest/mbest: 1/0, attached
    *via 133.101.52.51, Vlan52, [250/0], 00:07:17, am
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
DC33-102# sh ip adjacency 133.101.52.51

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
133.101.52.51   0085.6534.3300  50   arp        Vlan52           G
```

Find the PO Interface on which this MAC Address is Learnt:

```
DC33-102# sh mac address-table address 0085.6534.3300
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 52       0085.6534.3300    dynamic   502070     F    F   Po51
```

Display Po51 member interface with module information

```
DC33-102# sh port-channel summary | inc Po51
51    Po51(SU)    Eth      LACP      Eth1/43(P)
```

Display Adjacency Index for this Route in Hardware Table:

```
DC33-102# sh system internal forwarding ip route 133.101.52.51

Routes for table default/base

----+---------------------+----------+----------+-----------
Dev | Prefix              | PfxIndex | AdjIndex | LIF
----+---------------------+----------+----------+-----------
 1    133.101.52.51/32      0xaab8f430    0x18ab7    0x84
```

Display DMAC Entry Programmed in Adjacency Table:

```
DC33-102# sh system internal forwarding adjacency entry 0x18ab7 det
Device: 1   Index: 0x18ab7   dmac: 0085.6534.3300 smac: b0fa.eb5f.dafc
        e-lif: 0x84
```

3.6.3.1.3.1 Unicast ECMP verification

RIB&FIB Verification “show ip route”:

```
DC33-1(config)# sh ip route 133.101.11.0/24
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

133.101.11.0/24, ubest/mbest: 16/0
    *via 33.101.11.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.12.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.13.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.14.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.15.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.16.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.17.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.101.18.101, [20/0], 03:51:58, bgp-33, external, tag 33101
    *via 33.102.11.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.12.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.13.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.14.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.15.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.16.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.17.102, [20/0], 03:52:01, bgp-33, external, tag 33101
    *via 33.102.18.102, [20/0], 03:52:01, bgp-33, external, tag 33101
```

RIB&FIB Verification “show forwarding route”:

```
DC33-1(config)# show forwarding route 133.101.11.0/24

IPv4 routes for table default/base

------------------+------------------+----------------------+-----------------
Prefix            | Next-hop         | Interface            | Labels
------------------+------------------+----------------------+-----------------
*133.101.11.0/24     33.101.11.101      Ethernet1/1
                     33.101.12.101      Ethernet1/2
                     33.101.13.101      Ethernet1/3
                     33.101.14.101      Ethernet1/4
                     33.101.15.101      Ethernet1/5
                     33.101.16.101      Ethernet1/6
                     33.101.17.101      Ethernet1/7
                     33.101.18.101      Ethernet1/8
                     33.102.11.102      Ethernet1/9
                     33.102.12.102      Ethernet1/10
                     33.102.13.102      Ethernet1/11
                     33.102.14.102      Ethernet1/12
                     33.102.15.102      Ethernet1/13
                     33.102.16.102      Ethernet1/14
                     33.102.17.102      Ethernet1/15
                     33.102.18.102      Ethernet1/16
```

Programming Verification “show system internal forwarding ipv4 route”:

```
DC33-1(config)# show system internal forwarding ipv4 route 133.101.11.0/24

Routes for table default/base

----+--------------------+----------+----------+-----------
```

```
Dev | Prefix              | PfxIndex | AdjIndex | LIF
----+---------------------+----------+----------+-----------
 1    133.101.11.0/24       0xaabe3814    0x30d62    0x3
 1         "                  "       0x30d62    0xc
 1         "                  "       0x30d62    0x6
 1         "                  "       0x30d62    0x9
 1         "                  "       0x30d62    0x8
 1         "                  "       0x30d62    0xa
 1         "                  "       0x30d62    0x4
 1         "                  "       0x30d62    0x7
 1         "                  "       0x30d62    0x5
 1         "                  "       0x30d62    0xe
 1         "                  "       0x30d62    0xd
 1         "                  "       0x30d62    0x11
 1         "                  "       0x30d62    0xb
 1         "                  "       0x30d62    0xf
 1         "                  "       0x30d62    0x10
 1         "                  "       0x30d62    0x12
```

Programming Verification “show system internal forwarding ipv4 route”:

```
DC33-1(config)# show system internal forwarding adjacency entry 0x30d62
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x3
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0xc
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x6
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x9
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x8
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0xa
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x4
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x7
Device: 1   Index: 0x30d62   dmac: 4403.a7a3.bdfc  smac: b0fa.eb5f.dc7c
        e-lif: 0x5
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0xe
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0xd
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0x11
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0xb
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0xf
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0x10
Device: 1   Index: 0x30d62   dmac: b0fa.eb5f.dafc  smac: b0fa.eb5f.dc7c
        e-lif: 0x12
```

Programming Verification “show system internal forwarding ipv4 route”:

```
DC33-1(config)# show platform fwm info l3lif all | grep 0x12
Eth1/16:sdb: lif_index-2-ifindex key = 0x12 data = 0x1a00f000
```

Programming Verification on the Broadcom Shell:

```
bcm-shell.0> l3 defip show
Unit 0, Total Number of DEFIP entries: 16385
#     VRF     Net addr             Next Hop Mac        INTF MODID PORT PRIO CLASS HIT VLAN
<TRUNCATED>
2698  1        133.101.11.0/24      00:00:00:00:00:00 200034    0     0     0    0 y      (ECMP)
<TRUNCATED>
```

3.6.3.2 Multicast Routing Design

Multicast routing has been enabled across the entire DC33 network. Multicast with multipath is enabled since from each leaf router there are multiple PIM enabled interfaces to each of the spine routers (all configured as anycast RP).

On the Nexus 3000 running the software release 6.0(2)U1(3), the spine switch does not immediately remove OIFs for interfaces that fail if they are routed ports; the multicast route table may have entries pointing to OIFs that are in operationally down state. OIF will eventually get removed due to periodic PIM protocol state maintenance. However, the OIF is immediately removed if it is a routed port-channel – even single member port channel. As a workaround, it is recommended to change all individual routed ports to single member port-channels when possible for Nexus 3000 (CSCul28087, CSCum21940).

DC33 Multicast Configuration:

```
feature pim

ip pim rp-address 40.3.254.1 group-list 230.3.0.0/16
ip pim send-rp-announce loopback1 group-list 230.33.0.0/16
ip pim send-rp-discovery loopback1
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen

interface loopback0
  ip pim sparse-mode

interface loopback1
  description dc33-RP
  ip address 40.33.254.1/32
  ip pim sparse-mode

interface port-channel1
  ip pim sparse-mode

interface port-channel2
  ip pim sparse-mode

interface port-channel3
  ip pim sparse-mode

interface port-channel4
  ip pim sparse-mode

interface port-channel1031
  ip pim sparse-mode

interface port-channel1032
  ip pim sparse-mode

interface Ethernet1/1
  ip pim sparse-mode

<TRUNCATED>

interface Ethernet1/48
  ip pim sparse-mode
```

```
feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.33.0.2 connect-source loopback0
ip msdp mesh-group 40.33.0.2 MESH33
ip msdp peer 40.33.0.3 connect-source loopback0
ip msdp mesh-group 40.33.0.3 MESH33
ip msdp peer 40.33.0.4 connect-source loopback0
ip msdp mesh-group 40.33.0.4 MESH33
interface loopback0
  ip address 40.33.0.1/32
  ipv6 address 2001:1:40:33:0:1:0:1/128
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip router ospf 3 area 0.0.0.0
  ipv6 router ospfv3 33 area 0.0.0.0
  ip pim sparse-mode
```

3.6.3.2.1 PIM-ASM Rendezvous Point

PIM Sparse Mode has been configured as the protocol of choice for multicast routing. NX-OS does not support PIM SSM and PIM Bidir operating over vPC. The RP is located at the spine layer.

3.6.3.2.1.1 Auto-RP

The DC33 testbed is designed to have the RP located at the spine layer to support the groups sourced from each different type of leaf router. Each RP is configured at the spine routers. DC33 makes use of Auto-RP to automate distribution of RP information in the network.

To Verify PIM RP:

```
DC33-1# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.33.254.1*, next Discovery message in: 00:00:21
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1, (0), uptime: 3d07h, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.33.254.1*, (0), uptime: 3d07h, expires: 00:02:47,
  priority: 0, RP-source: 40.33.254.1 (A), group ranges:
      230.33.0.0/16

DC33-1# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.33.0.0/16      ASM       40.33.254.1      -
```

3.6.3.2.1.1.1 Auto-RP Forward Listen

DC33 has enabled the Auto-RP listening and forwarding feature so that the Auto-RP mechanism can dynamically inform routers in the PIM domain of the group-to-RP mapping since PIM dense mode is not supported on NX-OS. By default, listening or forwarding of Auto-RP messages is not enabled on NX-OS.

3.6.3.2.1.2 Static RP

For the groups with a Rendezvous Point on the core, the RP is statically configured on all routers in the DC33 network.

To Verify PIM RP:

```
DC33-1# sh ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP RPA: 40.33.254.1*, next Discovery message in: 00:00:21
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None

RP: 40.3.254.1, (0), uptime: 3d07h, expires: never,
  priority: 0, RP-source: (local), group ranges:
      230.3.0.0/16
RP: 40.33.254.1*, (0), uptime: 3d07h, expires: 00:02:47,
  priority: 0, RP-source: 40.33.254.1 (A), group ranges:
      230.33.0.0/16

DC33-1# sh ip pim group-range
PIM Group-Range Configuration for VRF "default"
Group-range        Mode      RP-address       Shared-tree-only range
232.0.0.0/8        SSM       -                -
230.3.0.0/16       ASM       40.3.254.1       -
230.33.0.0/16      ASM       40.33.254.1      -
```

3.6.3.2.1.3 Anycast RP with MSDP

DC33 has configured Anycast RP with MSDP at the spine layer.

DC33 Anycast RP and MSDP Configuration:

```
!Anycast RP configuration
ip pim send-rp-announce loopback1 group-list 230.33.0.0/16
ip pim send-rp-discovery loopback1
interface loopback1
  description dc33-RP
  ip address 40.33.254.1/32
  ip pim sparse-mode

! MSDP configuration
ip msdp originator-id loopback0
ip msdp peer 40.33.0.2 connect-source loopback0
ip msdp mesh-group 40.33.0.2 MESH33
ip msdp peer 40.33.0.3 connect-source loopback0
ip msdp mesh-group 40.33.0.3 MESH33
ip msdp peer 40.33.0.4 connect-source loopback0
ip msdp mesh-group 40.33.0.4 MESH33
interface loopback0
  ip address 40.33.0.1/32
  ipv6 address 2001:1:40:33:0:1:0:1/128
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip router ospf 3 area 0.0.0.0
  ipv6 router ospfv3 33 area 0.0.0.0
```

```
  ip pim sparse-mode
```

To Verify MSDP Peer and SA_Cache:

```
DC33-1# sh ip msdp sa-cache
MSDP SA Route Cache for VRF "default" - 640 entries
Source          Group           RP              ASN         Uptime
133.101.11.41   230.33.0.1      40.33.0.3       0           3d07h
133.101.11.42   230.33.0.1      40.33.0.3       0           3d07h
133.101.11.43   230.33.0.1      40.33.0.3       0           3d07h
133.101.11.44   230.33.0.1      40.33.0.3       0           3d07h
133.101.11.45   230.33.0.1      40.33.0.3       0           3d07h
133.101.12.41   230.33.0.1      40.33.0.3       0           3d07h

DC33-1# sh ip msdp sum
MSDP Peer Status Summary for VRF "default"
Local ASN: 33, originator-id: 40.33.0.1

Number of configured peers:  3
Number of established peers: 3
Number of shutdown peers:    0

Peer            Peer        Connection      Uptime/   Last msg  (S,G)s
Address         ASN         State           Downtime  Received  Received
40.33.0.2       0           Established     3d07h     00:00:56  50
40.33.0.3       0           Established     3d07h     00:00:20  550
40.33.0.4       0           Established     3d07h     00:00:04  40
```

3.6.3.2.1.3.1 MSDP Mesh Group

MSDP Mesh Group is configured on the spines to prevent each MSDP peer from advertising SA learned from other peers i.e., only locally registered sources.

```
feature msdp

ip msdp originator-id loopback0
ip msdp peer 40.33.0.2 connect-source loopback0
ip msdp mesh-group 40.33.0.2 MESH33
ip msdp peer 40.33.0.3 connect-source loopback0
ip msdp mesh-group 40.33.0.3 MESH33
ip msdp peer 40.33.0.4 connect-source loopback0
ip msdp mesh-group 40.33.0.4 MESH33

interface loopback0
  ip address 40.33.0.1/32
  ipv6 address 2001:1:40:33:0:1:0:1/128
  ip ospf message-digest-key 33 md5 3 a667d47acc18ea6b
  ip router ospf 3 area 0.0.0.0
  ipv6 router ospfv3 33 area 0.0.0.0
  ip pim sparse-mode
```

3.6.3.2.2 PIM SPT-Threshold

DC33 has enabled *ip pim spt-threshold infinity* on the last hop non-vPC PIM routers to decrease the multicast entries hardware usage across the network. Nexus 3000 vPC does not support PIM spt-threshold configuration.

3.6.3.2.3 Multicast Multipath

Cisco NX-OS Multicast Multipath is enabled by default; the load sharing selection algorithm is based on the source and group addresses.

3.6.3.2.4 Static OIF

On DC33 network, NVT has configured and tested the static-oif feature on the leaf layer to statically designate a receiver on a given subnet.

Static-OIF configuration and verification:

```
route-map EW_STATIC_JOIN permit 10
  match ip multicast group-range 230.33.0.1 to 230.33.0.10

interface Vlan110
  no shutdown
  mtu 9216
  no ip redirects
  ip address 133.101.110.3/24
  ipv6 address 2001:133:101:110::3/64
  ip pim sparse-mode
  ip igmp static-oif route-map EW_STATIC_JOIN
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 133.101.110.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 2001:133:101:110::1

DC33-102(config-if)# show ip mroute 230.33.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.33.0.1/32), uptime: 00:13:12, igmp pim ip static
  Incoming interface: Ethernet1/30, RPF nbr: 33.102.46.4
  Outgoing interface list: (count: 11)
    Vlan110, uptime: 00:00:33, static
    Vlan51, uptime: 00:11:52, igmp
    Vlan22, uptime: 00:11:53, igmp
    Vlan31, uptime: 00:11:53, igmp
    Vlan21, uptime: 00:12:00, igmp
    Vlan52, uptime: 00:12:01, igmp
    Vlan42, uptime: 00:12:01, igmp
    Vlan41, uptime: 00:12:04, igmp
    Vlan32, uptime: 00:12:07, igmp
    Vlan11, uptime: 00:13:10, igmp
    Vlan12, uptime: 00:13:12, igmp

(133.101.11.41/32, 230.33.0.1/32), uptime: 00:13:14, pim ip mrib
  Incoming interface: Vlan11, RPF nbr: 133.101.11.41
  Outgoing interface list: (count: 12)
    Vlan110, uptime: 00:00:33, mrib
    Vlan51, uptime: 00:11:52, mrib
    Vlan22, uptime: 00:11:53, mrib
    Vlan31, uptime: 00:11:53, mrib
    Vlan21, uptime: 00:12:00, mrib
    Vlan52, uptime: 00:12:01, mrib
    Vlan42, uptime: 00:12:01, mrib
```

```
     Vlan41, uptime: 00:12:04, mrib
     Vlan32, uptime: 00:12:07, mrib
     Ethernet1/25, uptime: 00:12:26, pim
     Vlan11, uptime: 00:13:10, mrib, (RPF)
     Vlan12, uptime: 00:13:12, mrib
```

Reloading the switch may lead to the configuration loss of the *static-oif* from the interface configuration (CSCul45536). It is good practice to save the configurations in order to retrieve the potentially lost configurations.

3.6.3.2.4.1 Static IGMP Snooping entry

In order to properly receive the multicast traffic on the statically configured subnet, L2 multicast programming has to be properly configured by either disabling IGMP snooping or configuring additional IGMP snooping static entries for each multicast group:

Static IGMP Snooping configuration and verification:

```
vlan configuration 110
  ip igmp snooping static-group 230.33.0.1 interface port-channel101
  ip igmp snooping static-group 230.33.0.2 interface port-channel101
  ip igmp snooping static-group 230.33.0.3 interface port-channel101
  ip igmp snooping static-group 230.33.0.4 interface port-channel101
  ip igmp snooping static-group 230.33.0.5 interface port-channel101
  ip igmp snooping static-group 230.33.0.6 interface port-channel101
  ip igmp snooping static-group 230.33.0.7 interface port-channel101
  ip igmp snooping static-group 230.33.0.8 interface port-channel101
  ip igmp snooping static-group 230.33.0.9 interface port-channel101
  ip igmp snooping static-group 230.33.0.10 interface port-channel101
vlan 110

DC33-102(config-if)# sh ip igmp snooping vlan 110
IGMP Snooping information for vlan 110
  IGMP snooping enabled
  Optimised Multicast Flood (OMF) disabled
  IGMP querier present, address: 133.101.110.2, version: 2, i/f Po200
  Switch-querier disabled
  IGMPv3 Explicit tracking enabled
  IGMPv2 Fast leave disabled
  IGMPv1/v2 Report suppression enabled
  IGMPv3 Report suppression disabled
  Link Local Groups suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 2
  Number of groups: 10
  VLAN vPC function enabled
  Active ports:
    Eth1/49   Po101   Po200

DC33-102(config-if)# sh ip igmp snooping groups 230.33.0.1 vlan 110
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
110   230.33.0.1         v2   S     Po101
DC33-102# show mac address-table multicast igmp-snooping
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY    Ports
---------+-----------------+--------+---------+------+----+------------------
  110      0100.5e21.0001    igmp      0          F    F  Po101 Po200
```

```
  110      0100.5e21.0002    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0003    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0004    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0005    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0006    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0007    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0008    igmp      0          F    F   Po101 Po200
  110      0100.5e21.0009    igmp      0          F    F   Po101 Po200
  110      0100.5e21.000a    igmp      0          F    F   Po101 Po200
<TRUNCATED>
DC33-102# show ip igmp snooping mrouter vlan 110
Type: S - Static, D - Dynamic, I - Internal
Type: S - Static, D - Dynamic, I - Internal
Vlan  Router-port   Type      Uptime      Expires
110   Po200         SVD       04:19:19    00:03:48
110   Vlan110       I         04:19:17    never
```

On a vPC setup running 6.0(2)U1(3), the multicast RP should be equidistant from both vPC peers. If this condition is not met, static-oif on vPC peers is not supported (CSCum01506).

3.6.3.2.5 Multicast Forwarding Verification

The following sequence of commands illustrates the verification of the Cisco NX-OS multicast L2 and L3 forwarding.

Displays a Specific Multicast Route 230.33.0.1:

```
DC33-102# show ip mroute 230.33.0.1
IP Multicast Routing Table for VRF "default"

(*, 230.33.0.1/32), uptime: 1w0d, igmp pim ip
  Incoming interface: Ethernet1/9, RPF nbr: 33.102.21.2
  Outgoing interface list: (count: 10)
    Vlan12, uptime: 1w0d, igmp
    Vlan31, uptime: 1w0d, igmp
    Vlan21, uptime: 1w0d, igmp
    Vlan52, uptime: 1w0d, igmp
    Vlan22, uptime: 1w0d, igmp
    Vlan42, uptime: 1w0d, igmp
    Vlan32, uptime: 1w0d, igmp
    Vlan51, uptime: 1w0d, igmp
    Vlan41, uptime: 1w0d, igmp
    Vlan11, uptime: 1w0d, igmp

(133.101.11.41/32, 230.33.0.1/32), uptime: 5d23h, pim mrib ip
  Incoming interface: Vlan11, RPF nbr: 133.101.11.41
  Outgoing interface list: (count: 11)
    Ethernet1/25, uptime: 5d23h, pim
    Vlan11, uptime: 5d23h, mrib, (RPF)
    Vlan12, uptime: 5d23h, mrib
    Vlan21, uptime: 5d23h, mrib
    Vlan22, uptime: 5d23h, mrib
    Vlan31, uptime: 5d23h, mrib
    Vlan32, uptime: 5d23h, mrib
    Vlan41, uptime: 5d23h, mrib
    Vlan42, uptime: 5d23h, mrib
    Vlan51, uptime: 5d23h, mrib
    Vlan52, uptime: 5d23h, mrib
```

Displays the Internal Forwarding Adjacency 230.33.0.1:

```
DC33-102# show system internal forwarding adjacency multicast group 230.33.0.1
```

```
230.33.0.1              0.0.0.0
oif-list:                                  0x2000002
                 Ref-Count:          690
                     Port:             0       Encap:          50
                     Port:             0       Encap:          51
                     Port:             1       Encap:          60
                     Port:             1       Encap:          61
                     Port:             2       Encap:          70
                     Port:             2       Encap:          71
                     Port:             3       Encap:          80
                     Port:             3       Encap:          81
                     Port:             4       Encap:          90
                     Port:             4       Encap:          91
                     Port:             9       Encap:          49
                     Port:            10       Encap:          49
                     Port:            10       Encap:          50
                     Port:            10       Encap:          51
                     Port:            10       Encap:          60
                     Port:            10       Encap:          61
                     Port:            10       Encap:          70
                     Port:            10       Encap:          71
                     Port:            10       Encap:          80
                     Port:            10       Encap:          81
                     Port:            10       Encap:          90
                     Port:            10       Encap:          91
230.33.0.1              133.101.11.41
oif-list:                                  0x2000003
                 Ref-Count:            4
                     Port:            25       Encap:          28
                     Port:             0       Encap:          -1
                     Port:             0       Encap:          51
                     Port:             1       Encap:          60
                     Port:             1       Encap:          61
                     Port:             2       Encap:          70
                     Port:             2       Encap:          71
                     Port:             3       Encap:          80
                     Port:             3       Encap:          81
                     Port:             4       Encap:          90
                     Port:             4       Encap:          91
                     Port:             9       Encap:          49
                     Port:            10       Encap:          -1
<TRUNCATED>
230.33.0.1              133.103.1.41
oif-list:                                  0x2000002
                 Ref-Count:          690
                     Port:             0       Encap:          50
                     Port:             0       Encap:          51
                     Port:             1       Encap:          60
                     Port:             1       Encap:          61
                     Port:             2       Encap:          70
                     Port:             2       Encap:          71
                     Port:             3       Encap:          80
                     Port:             3       Encap:          81
                     Port:             4       Encap:          90
                     Port:             4       Encap:          91
                     Port:             9       Encap:          49
                     Port:            10       Encap:          49
                     Port:            10       Encap:          50
                     Port:            10       Encap:          51
                     Port:            10       Encap:          60
                     Port:            10       Encap:          61
                     Port:            10       Encap:          70
                     Port:            10       Encap:          71
                     Port:            10       Encap:          80
                     Port:            10       Encap:          81
                     Port:            10       Encap:          90
                     Port:            10       Encap:          91
<TRUNCATED>
```

Display DR Information for Interface Vlan11:

```
DC33-102# sh ip pim interface brief
PIM Interface Status for VRF "default"
Interface            IP Address      PIM DR Address  Neighbor  Border
                                                     Count     Interface
Vlan110              133.101.110.3   133.101.110.3   1         no
<TRUNCATED>
Vlan1                133.101.1.3     133.101.1.3     1         no
Ethernet1/1          33.102.11.102   33.102.11.102   1         no
<TRUNCATED>
Ethernet1/32         33.102.48.102   33.102.48.102   1         no
```

Displays Mroute RPF Interface and Forwarding Counters in L3 Hardware Table:

```
DC33-102# sh forwarding multicast route group 230.33.0.1 source 133.101.11.41

  (133.101.11.41/32, 230.33.0.1/32), RPF Interface: Vlan11, flags:
    Received Packets: 34450 Bytes: 2239250
    Number of Outgoing Interfaces: 10
    Outgoing Interface List Index: 25
      Vlan12 Outgoing Packets:0 Bytes:0
      Vlan21 Outgoing Packets:0 Bytes:0
      Vlan22 Outgoing Packets:0 Bytes:0
      Vlan31 Outgoing Packets:0 Bytes:0
      Vlan32 Outgoing Packets:0 Bytes:0
      Vlan41 Outgoing Packets:0 Bytes:0
      Vlan42 Outgoing Packets:0 Bytes:0
      Vlan51 Outgoing Packets:0 Bytes:0
      Vlan52 Outgoing Packets:0 Bytes:0
      Ethernet1/25 Outgoing Packets:0 Bytes:0
```

Displays the Multicast Routing Table with Packet Counts and Bit Rates for All Sources:

```
DC33-102# sh ip mr 230.33.0.1 sum
IP Multicast Routing Table for VRF "default"

Total number of routes: 1191
Total number of (*,G) routes: 10
Total number of (S,G) routes: 1180
Total number of (*,G-prefix) routes: 1
Group count: 10, rough average sources per group: 118.0

Group: 230.33.0.1/32, Source count: 118
Source          packets      bytes           aps   pps       bit-rate     oifs
(*,G)           1633         700411          428   0         0.000   bps  10
133.101.11.41   165173       136402146       825   3         38.060  kbps 11
133.101.11.42   189730       170241692       897   0         27.200  bps  12
133.101.11.43   163217       133708105       819   0         27.200  bps  12
133.101.11.44   151341       117341650       775   0         27.200  bps  11
133.101.11.45   225765       219899249       974   0         27.200  bps  11
<TRUNCATED>
133.120.17.41   69048        3943223         57    0         27.200  bps  0
133.121.18.41   68909        3751681         54    0         27.200  bps  0
```

Display IGMP Snooping Groups Information:

```
DC33-102# sh ip igmp snooping groups 230.33.0.1 vlan 11
Type: S - Static, D - Dynamic, R - Router port, F - FabricPath core port

Vlan  Group Address      Ver  Type  Port list
11    230.33.0.1         v2   D     Po11
```

Displays Detected Multicast Routers for VLAN:

```
DC33-102# sh ip igmp snooping mrouter vlan 11
Type: S - Static, D - Dynamic, I - Internal
```

```
Type: S - Static, D - Dynamic, I - Internal
Vlan  Router-port   Type      Uptime      Expires
11    Po200         SVD       2w1d        00:04:09
11    Vlan11        I         2w1d        never
```

Displays IGMP Snooping Querier Information for VLAN:

```
DC33-102# sh ip igmp snooping querier vlan 11
Vlan  IP Address       Version   Expires     Port
11    133.101.11.2     v2        00:02:58    port-channel200
```

3.6.4 Layer-2/ Layer-3 Leaf/Access Layer Network Design Overview

3.6.4.1 vPC

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco NX-OS switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

vPC Peer Configurations:

```
N3000-1:
feature vpc

! vpc domain config
vpc domain 101

peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  delay restore 150
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel102
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/42
  description Eth1/42==Eth1/42 dc33-102
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  channel-group 102 mode active

! vpc peer-keepalive config
interface Ethernet1/41
  description Eth1/41==Eth1/41 dc33-102
  no switchport
  vrf member vpc-keepalive
  ip address 1.1.1.1/24

! vpc member port-channel config
interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  spanning-tree port type edge trunk
  vpc 11

! vpc member port config
interface Ethernet1/1
  description Eth1/1==Eth7/1 dc33-1001
  switchport mode trunk
  switchport trunk allowed vlan 11-20
```

```
N3000-2:
feature vpc

! vpc domain config
vpc domain 101
  role priority 201
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  delay restore 150
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel101
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/42
  description Eth1/42==Eth1/42 dc33-101
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  channel-group 101 mode active

! vpc peer-keepalive config
interface Ethernet1/41
  description Eth1/41==Eth1/41 dc33-101
  no switchport
  vrf member vpc-keepalive
  ip address 1.1.1.2/24

! vpc member port-channel config
interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  spanning-tree port type edge trunk
  vpc 11

! vpc member port config
interface Ethernet1/1
  description Eth1/1==Eth8/1 dc33-1001
  switchport mode trunk
  switchport trunk allowed vlan 11-20
```

channel-group 11 mode active ! PIM prebuild SPT ip pim pre-build-spt	channel-group 11 mode active ! PIM prebuild SPT ip pim pre-build-spt

Display vPC Status:

```
DC33-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                      : 301
Peer status                        : peer adjacency formed ok
vPC keep-alive status              : peer is alive
Configuration consistency status  : success
Per-vlan consistency status        : success
Type-2 consistency status          : success
vPC role                           : primary, operational secondary
Number of vPCs configured          : 10
Peer Gateway                       : Enabled
Peer gateway excluded VLANs      : -
Dual-active excluded VLANs         : -
Graceful Consistency Check         : Enabled
Auto-recovery status               : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ -------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.6.4.1.1 LACP

DC33 makes use of LACP mode active for all link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
DC33-102# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
11    Po11(SU)    Eth      LACP      Eth1/39(P)
21    Po21(SU)    Eth      LACP      Eth1/40(P)
31    Po31(SU)    Eth      LACP      Eth1/41(P)
```

```
41    Po41(SU)    Eth      LACP      Eth1/42(P)
51    Po51(SU)    Eth      LACP      Eth1/43(P)
61    Po61(SU)    Eth      LACP      Eth1/44(P)
71    Po71(SU)    Eth      LACP      Eth1/45(P)
81    Po81(SU)    Eth      LACP      Eth1/46(P)
91    Po91(SU)    Eth      LACP      Eth1/47(P)
101   Po101(SU)   Eth      LACP      Eth1/48(P)
200   Po200(SU)   Eth      LACP      Eth1/50(P)   Eth1/51(P)
DC33-102# show lacp interface ethernet 1/39
Interface Ethernet1/39 is up
  Channel group is 11 port channel is Po11
  PDUs sent: 44463
  PDUs rcvd: 48063
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(7f9b, 0-23-4-ee-bf-2d, 800b, 8000, 127), (8000, 0-1e-f6-e7-6c-0, b,
8000, 228)] ]
Operational as aggregated link since Wed Jan 29 11:38:49 2014

Local Port: Eth1/39   MAC Address= 0-23-4-ee-bf-2d
  System Identifier=0x8000,0-23-4-ee-bf-2d
  Port Identifier=0x8000,0x127
  Operational key=32779
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=(Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Neighbor: 0x228
  MAC Address= 0-1e-f6-e7-6c-0
  System Identifier=0x8000,  Port Identifier=0x8000,0x228
  Operational key=11
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=(Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
```

3.6.4.1.2 VLAN Trunking

DC33 makes use of VLAN trunking to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Display VLAN Information for Nexus 3000:

```
DC33-102# sh vlan internal usage

VLANs                   DESCRIPTION
-------------------     -----------------
3968-4031               Multicast
4032-4035               Online Diagnostic
4036-4039               ERSPAN
4042                    Satellite
3968-4047,4094          Current
DC33-102# show vlan id 11

VLAN Name                             Status    Ports
```

```
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po11, Po200, Eth1/33, Eth1/34
                                                Eth1/39, Eth1/49, Eth1/50
                                                Eth1/51

VLAN Type  Vlan-mode
---- ----- ----------
11   enet  CE

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------

DC33-102# sh int po101 trunk

--------------------------------------------------------------------------------
Port          Native  Status        Port
              Vlan                  Channel
--------------------------------------------------------------------------------
Po101         1       trunking      --

--------------------------------------------------------------------------------
Port          Vlans Allowed on Trunk
--------------------------------------------------------------------------------
Po101         101-110

--------------------------------------------------------------------------------
Port          Vlans Err-disabled on Trunk
--------------------------------------------------------------------------------
Po101         none

--------------------------------------------------------------------------------
Port          STP Forwarding
--------------------------------------------------------------------------------
Po101         101-110

--------------------------------------------------------------------------------
Port          Vlans in spanning tree forwarding state and not pruned
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
Port          Vlans Forwarding on FabricPath
--------------------------------------------------------------------------------
Po101         none
```

3.6.4.1.3 Spanning Tree

vPC technology helps build a loop free topology by leveraging port-channels from access devices to the vPC domain. A port-channel is seen as a logical link from the spanning tree's standpoint, so a vPC domain with vPC-attached access devices forms a star topology at Layer 2 (there are no STP blocked ports in this type of topology). In this case, STP is used as a fail-safe mechanism to protect against any network loops.

DC33 makes use of Rapid-PVST which is the default spanning tree protocol on NX-OS. For networks with larger logical port counts, MST is recommended.

Display Spanning Tree Information:

```
DC33-102# sh spanning-tree vlan 11

VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    8203
```

```
            Address     4403.a7a3.bdfc
            Cost        2
            Port        4295 (port-channel200)
            Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8203   (priority 8192 sys-id-ext 11)
             Address     b0fa.eb5f.dafc
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po11             Desg FWD 1         128.4106 (vPC) P2p Peer(STP)
Po200            Root FWD 2         128.4295 (vPC peer-link) Network P2p
Eth1/49          Desg FWD 2         128.177  Edge P2p

DC33-102# sh spanning-tree summary totals
Switch is in rapid-pvst mode
Root bridge for: none
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short
STP-Lite                                 is enabled

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
101 vlans                    0         0        0        302        302
```

Display L2 Table-VLAN and L2 Table-STG Tables Information from Broadcom Shell:

```
bcm-shell.0> vlan show 11
vlan 11 ports cpu,ge38,xe0-xe2 (0x0000000000000000000000000000000000000000000000000e008000000001),
untagged none (0x0000000000000000000000000000000000000000000000000000000000000000) MCAST_FLOOD_UNKNOWN
bcm-shell.0> dump vlan 11
VLAN.ipipe0[11]:
<VP_GROUP_BITMAP=0,VLAN_PROFILE_PTR=3,VLAN_CLASS_ID=0x1c,VIRTUAL_PORT_EN=0,VALID=1,UUC_TRILL_NETWORK_RECEI
VERS_PRESENT=0,UUC_IDX=0,UMC_TRILL_NETWORK_RECEIVERS_PRESENT=0,UMC_IDX=0,TRILL_TRANSIT_IGMP_MLD_PAYLOAD_TO
_CPU=1,TRILL_RBRIDGE_NICKNAME_INDEX=0,TRILL_DOMAIN_NONUC_REPL_INDEX=0,TRILL_ACCESS_RECEIVERS_PRESENT=0,STG
=0x63,SRC_PVLAN_PORT_TYPE=0,SERVICE_CTR_IDX=0x64,PORT_BITMAP_W2=0,PORT_BITMAP_W1=0xe0080,PORT_BITMAP_W0=1,
PORT_BITMAP=0xe008000000001,L2_ENTRY_KEY_TYPE=0,ING_PORT_BITMAP_W2=0,ING_PORT_BITMAP_W1=0xe0080,ING_PORT_B
ITMAP_W0=1,ING_PORT_BITMAP=0xe008000000001,HIGIG_TRUNK_OVERRIDE_PROFILE_PTR=0,FID_ID=0xb,EVEN_PARITY_1=0,E
VEN_PARITY_0=0,ENABLE_IGMP_MLD_SNOOPING=0,BC_TRILL_NETWORK_RECEIVERS_PRESENT=0,BC_IDX=0>
bcm-shell.0> stg stp 63
STG 63:
    Block: ge32-ge33,ge38-ge40,ge42-ge47,xe3
  Forward: ge0-ge31,ge34-ge37,ge41,xe0-xe2
```

Display L2 Table-L2UserEntry:

```
DC33-102(config)# show hardware internal bcm-usd info tables l2 l2-user-entry all slot-num 0 | exclude
0000.0000.0000
Slot number 0

                    [ L2_USER_ENTRY (all info) - B549 TABLE ]

+----+----+--+-+-+----+---+----+----+--+----------------+-+----+--------------+
|    |  P |  | | |                |D   |  |                | |    |              |
|    | ER |  |D| |                |S   |  |                | |    |              |
|    | VO |  |O| |                |T   |  |                | |    |              |
|    | ET |  |N| |                |    |  |                | |    |              |
|    | NO |  |T| |                |D   |  |                |K|    |              |
|    |  C |C | | |                |I   |  |                |E|    |              |
|    | PO |L |L| |                |S   |  |                |Y|    |              |
|    |VAL |A |R|T|                |C   |  |                |T|    |              |
|    |AR B|S |N|R|    TRUNK       |    |  |                | |    |              |
```

```
|    |LIPP|S | |U|============|ACR |P |                |Y|    |              |
|    |ITKD|  |S|N|     MOD PORT|RPPL|R |                |P|    |              |
|ADDR|DYTU|ID|A|K|TGID ID  NUM |DUE3|I |MASK            |E|VLAN|  MAC ADDRESS |
+----+----+--+-+-+----+---+----+----+--+----------------+-+----+--------------+
   0 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
   1 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0100.0ccc.ccc
   2 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0100.0ccc.ccc
   3 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
   4 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
```

3.6.4.1.4 Configuration Parameters Consistency

After the vPC feature is enabled and the vPC peer-link on both peer devices is configured, Cisco Fabric Services messages provide a copy of the local vPC peer device configuration to the remote vPC peer device. The systems then determine whether any of the crucial configuration parameters differ on the two devices.

When a Type 1 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 1 consistency check failure, all vPC member ports are set to down state.
- For a vPC interface configuration Type 1 consistency check failure, the misconfigured vPC is set to down state.

When a Type 2 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 2 consistency check failure, all vPC member ports remain in up state and vPC systems trigger protective actions.
- For a vPC interface configuration Type 2 consistency check failures, the misconfigured vPC remains in up state. However, depending on the discrepancy type, vPC systems will trigger protective actions. The most typical misconfiguration deals with the allowed VLANs in the vPC interface trunking configuration. In this case, vPC systems will disable the vPC interface VLANs that do not match on both sides.

Display vPC Consistency Parameters:

```
DC33-102# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
QoS                         2     ([], [], [], [], [],   ([], [], [], [], [],
                                  [], [], [])            [], [], [])
Network QoS (MTU)           2     (9216, 0, 0, 0, 0, 0,  (9216, 0, 0, 0, 0, 0,
                                  0, 0)                  0, 0)
Network Qos (Pause)         2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Network Qos (WRED)          2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Network Qos (ECN)           2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Output Queuing (Bandwidth)  2     (100, 0, 0, 0, 0, 0,   (100, 0, 0, 0, 0, 0,
                                  0, 0)                  0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
Priority)                         F)                     F)
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
```

```
STP MST Region Name          1     ""                     ""
STP MST Region Revision      1     0                      0
STP MST Region Instance to   1
 VLAN Mapping
STP Loopguard                1     Disabled               Disabled
STP Bridge Assurance         1     Enabled                Enabled
STP Port Type, Edge          1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard         Disabled               Disabled
STP MST Simulate PVST        1     Enabled                Enabled
IGMP Snooping Group-Limit    2     8000                   8000
Interface-vlan admin up      2     1,11-110               1,11-110
Interface-vlan routing       2     1,11-110               1,11-110
capability
Allowed VLANs                -     1,11-110               1,11-110
Local suspended VLANs        -     -                      -

DC33-102# show vpc consistency-parameters interface port-channel 11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                         Type  Local Value            Peer Value
-------------                ----  ---------------------- -----------------------
Shut Lan                     1     No                     No
STP Port Type                1     Edge Trunk Port        Edge Trunk Port
STP Port Guard               1     None                   None
STP MST Simulate PVST        1     Default                Default
lag-id                       1     [(7f9b,                [(7f9b,
                                   0-23-4-ee-bf-2d, 800b, 0-23-4-ee-bf-2d, 800b,
                                    0, 0), (8000,          0, 0), (8000,
                                   0-1e-f6-e7-6c-0, b, 0, 0-1e-f6-e7-6c-0, b, 0,
                                    0)]                    0)]
mode                         1     active                 active
Speed                        1     1000 Mb/s              1000 Mb/s
Duplex                       1     full                   full
Port Mode                    1     trunk                  trunk
Native Vlan                  1     1                      1
MTU                          1     1500                   1500
Admin port mode              1
vPC card type                1     Empty                  Empty
Allowed VLANs                -     11-20                  11-20
Local suspended VLANs        -     -                      -
```

3.6.4.1.5 vPC Role Priority

There are two defined vPC roles: primary and secondary. The vPC role defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDUs) and responds to Address Resolution Protocol (ARP).
In case of a tie (same role priority value defined on both peer devices), the lowest system MAC will dictate the primary peer device.

Display vPC Role, System-MAC, System-Priority:

```
DC33-102# show vpc role

vPC Role status
----------------------------------------------------
vPC role                        : primary, operational secondary
Dual Active Detection Status    : 0
vPC system-mac                  : 00:23:04:ee:bf:2d
vPC system-priority             : 32667
vPC local system-mac            : b0:fa:eb:5f:da:fc
vPC local role-priority         : 201
```

3.6.4.1.6 vPC Peer-Link

The vPC peer-link is a standard 802.1Q trunk that performs the following actions:

- Carry vPC and non-vPC VLANs.
- Carry Cisco Fabric Services (CFS) messages that are tagged with CoS=4 for reliable communication CoS=4 for reliable communication.
- Carry flooded traffic between the vPC peer devices.
- Carry STP BPDUs, HSRP hello messages, and IGMP updates.

When the vPC peer-link fails and the vPC peer-keepalive link is still up, the vPC secondary peer device performs the following operations:

- Suspends its vPC member ports
- Shuts down the SVI associated to the vPC VLAN

Display vPC Peer-link Information:

```
DC33-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.6.4.1.7 vPC Peer-Keepalive Link

The vPC peer-keepalive link is a Layer 3 link that joins one vPC peer device to the other vPC peer device and carries a periodic heartbeat between those devices. It is used at the boot up of the vPC systems to guarantee that both peer devices are up before forming the vPC domain. It is also used when the vPC peer-link fails, in which case, the vPC peer-keepalive link is leveraged to detect split brain scenario (both vPC peer devices are active-active).

Default Values for VPC Peer-Keepalive Links:

Timer	Default value
Keepalive interval	1 seconds
Keepalive hold timeout (on vPC peer-link loss)	3 seconds
Keepalive timeout	5 seconds

Display vPC Peer-Keepalive Information:

```
DC33-102# sh vpc peer-keepalive

vPC keep-alive status             : peer is alive
--Peer is alive for             : (1334740) seconds, (86) msec
--Send status                   : Success
--Last send at                  : 2014.02.13 22:23:09 342 ms
--Sent on interface             : Eth1/35
--Receive status                : Success
--Last receive at               : 2014.02.13 22:23:08 889 ms
--Received on interface         : Eth1/35
--Last update from peer         : (0) seconds, (453) msec

vPC Keep-alive parameters
--Destination                   : 1.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc-keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

3.6.4.1.8 vPC Member Link

As suggested by the name, a vPC member port is a port-channel member of a vPC. A port-channel defined as a vPC member port always contains the keywords *vpc <vpc id>.*

A vPC only supports Layer 2 port-channels. The port-channel can be configured in access or trunk switchport mode. Any VLAN allowed on the vPC member port is by definition called a vPC VLAN. Whenever a vPC VLAN is defined on a vPC member port, it must also be defined on the vPC peer-link. Not defining a vPC VLAN on the vPC peer-link will cause the VLAN to be suspended.

The configuration of the vPC member port must match on both the vPC peer devices. If there is an inconsistency, a VLAN or the entire port channel may be suspended (depending on Type-1 or Type-2 consistency check for the vPC member port). For instance, a MTU mismatch will suspend the vPC member port.

Display vPC Member Port-channel Information:

```
DC33-102# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
```

```
vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
DC33-102# show vpc consistency-parameters interface port-channel 11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Shut Lan                    1     No                     No
STP Port Type               1     Edge Trunk Port        Edge Trunk Port
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-bf-2d, 800b, 0-23-4-ee-bf-2d, 800b,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-1e-f6-e7-6c-0, b, 0, 0-1e-f6-e7-6c-0, b, 0,
                                   0)]                    0)]
mode                        1     active                 active
Speed                       1     1000 Mb/s              1000 Mb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
Admin port mode             1
vPC card type               1     Empty                  Empty
Allowed VLANs               -     11-20                  11-20
Local suspended VLANs       -     -                      -
```

3.6.4.1.9 vPC ARP Synchronization

The vPC ARP Synchronization feature improves the convergence time for Layer 3 flows (North to South traffic). When the vPC peer-link fails and subsequently recovers, vPC ARP Synchronization performs an

ARP bulk synchronization over Cisco Fabric Services (CFS) from the vPC primary peer device to the vPC secondary peer device.

Displays vPC IP ARP sync information on the secondary vPC:

```
dc31-101# sh ip arp sync-entries

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Address         Age       MAC Address     Interface
131.11.155.252  00:01:45  0000.8c43.eb64  Vlan410
131.11.155.253  00:01:45  0000.8c43.5e23  Vlan410
131.11.155.254  00:01:45  0000.8c44.59ef  Vlan410
131.11.154.252  00:01:45  0000.8c43.eb62  Vlan409
…
```

Although the *ip arp synchronization* feature is configured on the Nexus 3000 platform running software release 6.0(2)U1(3), the command does not appear on the running-config of both vPC peers (CSCun29189).

3.6.4.1.10 vPC Delay Restore

After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from the access to the core until the Layer 3 connectivity is reestablished.

The vPC Delay Restore feature delays the vPCs leg bringup on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on the vPC leg. The result provides a graceful restoration along with zero packet loss during the recovery phase (traffic still gets diverted to the alive vPC peer device).

This feature is enabled by default with a vPC restoration default timer of 30 seconds, which DC33 maintains in the testbed.

3.6.4.1.11 vPC Auto-Recovery

vPC auto-recovery feature was designed to address 2 enhancements to vPC.

- To provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure (vPC auto-recovery feature).
- To handle a specific case where both vPC peer devices reload but only one comes back to life (vPC auto-recovery reload-delay feature).

The switch which unsuspends its vPC role with vPC auto-recovery continues to remain primary even after peer-link is on. The other peer takes the role of secondary and suspends its own vPC until a consistency check is complete. Therefore, to avoid this situation from occurring erroneously, auto-recovery reload-delay-timer should be configured to be long enough for the system to fully complete its bootup sequence.

Helpful Commands for vPC Object Tracking:

Show vpc brief	Displays Auto-recovery status

Configuration Check:

```
DC33-102# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.6.4.1.12 PIM Pre-Build-SPT with vPC

PIM Pre-build SPT on non-forwarder attracts multicast traffic by triggering upstream PIM J/Ps (Join/Prune) without setting any interface in the OIF (Outgoing Interface) list. Multicast traffic is then always pulled to the non-active forwarder and finally dropped due to no OIFs.

The immediate effect of enabling PIM Pre-build SPT is to improve the convergence time upon active forwarder failure (1 to 3 seconds of convergence time). The other vPC peer device (which is the non-active forwarder) does not need to create any new upstream multicast state and can quickly transition to the active forwarder role by properly programming the OIF (Outgoing Interface) list.
The impact of enabling PIM prebuild SPT is the consumption of bandwidth and replication capacity on the primary and secondary data path (i.e. on vPC primary and secondary peer devices) in steady state. As shown below, on the non-forwarder/secondary the (S,G) is created with no OIFs.

On the vPC peers:

N3000 1:	N3000 2:
DC33-101# **show ip mroute 230.33.0.1 shared** IP Multicast Routing Table for VRF "default"	DC33-102# **show ip mroute 230.33.0.1 shared** IP Multicast Routing Table for VRF "default"

```
(*, 230.33.0.1/32), uptime: 01:03:15, igmp ip pim
static
  Incoming interface: Ethernet1/29, RPF nbr:
33.101.45.4
  Outgoing interface list: (count: 11)
    Vlan110, uptime: 00:50:37, static
    Vlan51, uptime: 01:01:54, igmp
    Vlan22, uptime: 01:01:56, igmp
    Vlan31, uptime: 01:01:56, igmp
    Vlan21, uptime: 01:02:03, igmp
    Vlan42, uptime: 01:02:04, igmp
    Vlan52, uptime: 01:02:04, igmp
    Vlan41, uptime: 01:02:06, igmp
    Vlan32, uptime: 01:02:09, igmp
    Vlan11, uptime: 01:03:13, igmp
    Vlan12, uptime: 01:03:15, igmp

DC33-101# show ip mroute 230.33.0.1 133.103.1.41
IP Multicast Routing Table for VRF "default"

(133.103.1.41/32, 230.33.0.1/32), uptime: 01:07:38,
ip pim
  Incoming interface: Ethernet1/19, RPF nbr:
33.101.33.3
  Outgoing interface list: (count: 0)

DC33-101# sh ip pim intern vpc rpf-source
PIM vPC RPF-Source Cache for Context "default" -
Chassis Role Secondary

Source: 133.101.11.41
  Pref/Metric: 0/0
  Source role: secondary
  Forwarding state: Win-force (forwarding)

Source: 133.101.11.42
  Pref/Metric: 0/0
  Source role: secondary
  Forwarding state: Win-force (forwarding)

<TRUNCATED>
```

```
(*, 230.33.0.1/32), uptime: 01:03:20, igmp pim ip
static
  Incoming interface: Ethernet1/30, RPF nbr:
33.102.46.4
  Outgoing interface list: (count: 11)
    Vlan110, uptime: 00:50:41, static
    Vlan51, uptime: 01:02:00, igmp
    Vlan22, uptime: 01:02:01, igmp
    Vlan31, uptime: 01:02:01, igmp
    Vlan21, uptime: 01:02:08, igmp
    Vlan52, uptime: 01:02:09, igmp
    Vlan42, uptime: 01:02:09, igmp
    Vlan41, uptime: 01:02:12, igmp
    Vlan32, uptime: 01:02:15, igmp
    Vlan11, uptime: 01:03:18, igmp
    Vlan12, uptime: 01:03:20, igmp

DC33-102# show ip mroute 230.33.0.1 133.103.1.41
IP Multicast Routing Table for VRF "default"

(133.103.1.41/32, 230.33.0.1/32), uptime: 01:06:18,
ip pim mrib
  Incoming interface: Ethernet1/17, RPF nbr:
33.102.31.3
  Outgoing interface list: (count: 11)
    Vlan110, uptime: 00:54:14, mrib
    Vlan51, uptime: 01:05:33, mrib
    Vlan22, uptime: 01:05:34, mrib
    Vlan31, uptime: 01:05:35, mrib
    Vlan21, uptime: 01:05:41, mrib
    Vlan52, uptime: 01:05:42, mrib
    Vlan42, uptime: 01:05:43, mrib
    Vlan41, uptime: 01:05:45, mrib
    Vlan32, uptime: 01:05:48, mrib
    Vlan12, uptime: 01:06:18, mrib
    Vlan11, uptime: 01:06:18, mrib

DC33-102# sh ip pim intern vpc rpf-source
PIM vPC RPF-Source Cache for Context "default" -
Chassis Role Primary

Source: 133.101.11.41
  Pref/Metric: 0/0
  Source role: primary
  Forwarding state: Win-force (forwarding)

Source: 133.101.11.42
  Pref/Metric: 0/0
  Source role: primary
  Forwarding state: Win-force (forwarding)

<TRUNCATED>
```

3.6.4.1.13 HSRP/HSRPv6 Active/Active with vPC

HSRP in the context of vPC has been improved from a functional and implementation standpoint to take full benefits of the L2 dual-active peer devices nature offered by vPC technology. HSRP operates in active-active mode from a data plane standpoint, as opposed to classical active/standby implementation with a STP based network. No additional configuration is required. As soon as a vPC domain is configured and interface VLAN with an associated HSRP group is activated, HSRP will behave by default in active/active mode (on the data plane side).

From a control plane standpoint, active-standby mode still applies for HSRP in context of vPC; the active HSRP instance responds to ARP request. ARP response will contain the HSRP vMAC which is the same on both vPC peer devices. The standby HSRP vPC peer device just relays the ARP request to active HSRP peer device through the vPC peer-link.

HSRPv4&v6 Configurations:

N3000 1:	N3000 2:

`interface Vlan11`
`  no shutdown`
`  mtu 9216`
`  no ip redirects`
`  ip address 133.101.11.2/24`
`  ipv6 address 2001:133:101:11::2/64`
`  ip pim sparse-mode`
`  hsrp version 2`
`  hsrp 1`
`    authentication md5 key-string cisco`
`    preempt delay minimum 120`
`    priority 101`
`    ip 133.101.11.1`
`  hsrp 101 ipv6`
`    authentication md5 key-string cisco`
`    preempt delay minimum 120`
`    priority 101`
`    ip 2001:133:101:11::1`
	`interface Vlan11`
	`  no shutdown`
	`  mtu 9216`
	`  no ip redirects`
	`  ip address 133.101.11.3/24`
	`  ipv6 address 2001:133:101:11::3/64`
	`  ip pim sparse-mode`
	`  hsrp version 2`
	`  hsrp 1`
	`    authentication md5 key-string cisco`
	`    preempt delay minimum 120`
	`    priority 99`
	`    ip 133.101.11.1`
	`  hsrp 101 ipv6`
	`    authentication md5 key-string cisco`
	`    preempt delay minimum 120`
	`    priority 99`
	`    ip 2001:133:101:11::1`

Helpful Commands for HSRP Active/Active with vPC:

Show hsrp brief	Displays hsrp status
Show mac address-table vlan <vlan id>	Displays mac addresses including HSRP vMAC; check for G-flag on vMAC for active/active HSRP

Configuration Check:

```
DC33-102# sh hsrp brief
                    P indicates configured to preempt.
                    |
Interface   Grp Prio P State    Active addr      Standby addr     Group addr
Vlan11      1   99   P Standby  133.101.11.2     local            133.101.11.1
  (conf)
Vlan11      101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local           fe80
::5:73ff:fea0:65 (impl auto EUI64)
<TRUNCATED>
Vlan109     1   99   P Standby  133.101.109.2    local            133.101.109.1
  (conf)
Vlan109     101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local           fe80
::5:73ff:fea0:65 (impl auto EUI64)
Vlan110     1   99   P Standby  133.101.110.2    local            133.101.110.1
  (conf)
Vlan110     101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local           fe80
::5:73ff:fea0:65 (impl auto EUI64)
```

3.6.4.2 L2/L3 TCAM Tables

Nexus 3000/3548 platforms display MAC age as “seconds since first seen.” This behavior differs from the Nexus 5000, 6000 and 7000 platforms which are displayed as “seconds since last seen” and should be taken into account when reading the table. (CSCun37434)

When topology change notifications or MAC address clears are initiated on the Nexus 3000 the ARP address table also gets flushed (CSCun32115). As a result, the ARP table will be re-learned.

List of useful commands for TCAM table

Display L2 table-L2EntryTable:

```
DC33-102(config)# sh mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0000.0028.2a0c    dynamic   40         F    F  Po11
* 11       0085.650b.2900    dynamic   40         F    F  Po11
* 11       0085.650b.2901    dynamic   2320       F    F  Po11
* 11       0085.650b.2902    dynamic   40         F    F  Po11
<TRUNCATED>
* 11       0085.650b.33c5    dynamic   2320       F    F  Po11
* 11       0085.650b.33c6    dynamic   40         F    F  Po11
* 11       0085.650b.33c7    dynamic   2320       F    F  Po11
  11       0100.5e21.0001    igmp      0          F    F  Po11 Po200
<TRUNCATED>
  11       0100.5e21.0009    igmp      0          F    F  Po11 Po200
  11       0100.5e21.000a    igmp      0          F    F  Po11 Po200
```

Display L2 table-L2EntryTable:

```
DC33-102(config)# show platform fwm info hw-stm
VLAN   MAC Address      Port    PC
------+----------------+-------+---------
32    00:85:65:20:33:aa 2       Y[Po-1355349985
12    00:85:65:0c:33:58 0       Y[Po-1355354805
11    00:85:65:0b:33:22 0       Y[Po-1355354805
22    00:85:65:16:33:89 1       Y[Po-1355349995
42    00:85:65:2a:33:56 3       Y[Po-1355349975
<TRUNCATED>
```

Display L2 table-L2EntryTable from Broadcom shell:

```
bcm-shell.0> l2 show
<TRUNCATED>
mac=00:85:65:0b:33:50 vlan=11 GPORT=0x0 Trunk=0 Hit
mac=00:85:65:29:33:22 vlan=41 GPORT=0x0 Trunk=3 Hit
mac=00:85:65:34:33:89 vlan=52 GPORT=0x0 Trunk=4 Hit
mac=00:85:65:1f:33:01 vlan=31 GPORT=0x0 Trunk=2
mac=01:00:5e:21:00:03 vlan=41 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2070
mac=00:00:2f:64:fa:86 vlan=21 GPORT=0x0 Trunk=1
mac=00:85:65:0b:33:7c vlan=11 GPORT=0x0 Trunk=0
mac=00:85:65:2a:33:08 vlan=42 GPORT=0x0 Trunk=3
mac=01:00:5e:21:00:04 vlan=32 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2064
mac=00:85:65:0c:33:06 vlan=12 GPORT=0x0 Trunk=0 Hit
mac=01:00:5e:21:00:04 vlan=11 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2072
mac=00:85:65:0b:33:46 vlan=11 GPORT=0x0 Trunk=0 Hit
mac=00:85:65:2a:33:32 vlan=42 GPORT=0x0 Trunk=3
<TRUNCATED>
```

Display L3 TCAM from Broadcom shell:

```
bcm-shell.0> l3 defib show
<TRUNCATED>
3403  1        33.101.47.0/24       00:00:00:00:00:00 102180    0     0     0    0 n
3404  1        33.101.48.0/24       00:00:00:00:00:00 102180    0     0     0    0 n
3404  1        133.108.5.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3405  1        133.101.39.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
```

```
3405   1         133.101.25.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3406   1         133.106.3.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3406   1         133.109.6.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3407   1         133.101.108.0/24     00:00:00:00:00:00 100003    0     0     0    0 n
3407   1         133.101.57.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3408   1         133.112.9.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3408   1         133.110.7.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3409   1         133.101.43.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3409   1         133.101.75.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3410   1         133.111.8.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3410   1         133.107.4.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3411   1         133.101.61.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3411   1         133.101.49.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3412   1         133.116.13.0/24      00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3412   1         133.115.12.0/24      00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3413   1         133.101.12.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3413   1         133.101.93.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
<TRUNCATED>

bcm-shell.0> l3 egress show
<TRUNCATED>
102612  00:00:2f:65:af:cb   67   67     1t   0         -1   no   no
102613  00:00:2f:64:fa:6c  119  119     1t   0         -1   no   no
102614  00:00:2f:64:fa:66  119  119     1t   0         -1   no   no
102615  00:00:2f:65:af:bb   67   67     1t   0         -1   no   no
<TRUNCATED>

bcm-shell.0> ipmc table show
<TRUNCATED>
133.101.42.44   230.33.0.8        97   -1 1  -1 0    1     0 y
133.101.42.41   230.33.0.8        97   -1 1  -1 0    1     0 y
133.101.11.45   230.33.0.5       127   -1 1  -1 0    1     0 y
133.101.12.43   230.33.0.4        76   -1 1  -1 0    1     0 y
133.101.51.43   230.33.0.8        85   -1 1  -1 0    1     0 y
133.101.52.45   230.33.0.8       132   -1 1  -1 0    1     0 y
<TRUNCATED>
```

3.7 DC36

3.7.1 Configuration of Platform Specific Features On DC36

3.7.1.1 Licensing

License Usage on Nexus 3000 in DC36:

```
N3064# sh license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
--------------------------------------------------------------------------------
LAN_BASE_SERVICES_PKG         Yes   -   In use Never       -
ALGO_BOOST_SERVICES_PKG       Yes   -   Unused Never       -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   In use Never       -
--------------------------------------------------------------------------------
```

Although features can be enabled and configured in the CLI without licenses, they will not function until the license is installed.

3.7.1.2 Out-of-Band Management Network

DC36 makes use of out-of-band method to manage the chassis in the network to separate management traffic from production traffic.

Configuration:

```
interface mgmt0
  vrf member management
  ip address 10.2.36.1/16
```

3.7.1.3 Common Configurations

3.7.1.3.1 SSH and TACACS+

SSH is enabled in DC36 to provide connectivity for network device management. Authentication is provided through TACACS+.

Configuration and Verification:

```
feature tacacs+


ip tacacs source-interface mgmt0
tacacs-server host 172.28.92.17 key 7 "fewhg123"
aaa group server tacacs+ AAA-Servers
    server 172.28.92.17
    use-vrf management


N3064# sh ssh server
ssh version 2 is enabled
N3064# sh users
NAME     LINE         TIME         IDLE          PID COMMENT
interop  pts/0        Feb 10 11:37   .          3995 (taro.interop.cisco.com) session=ssh *
```

3.7.1.3.2 CDP and LLDP

CDP and LLDP are pervasively used on the DC36 test bed for inter-device discovery.

CDP Configuration and Verification:

```
DC36-5# sh run cdp all

!Command: show running-config cdp all
!Time: Tue Feb 18 10:29:16 2014

version 6.0(2)U2(1)
cdp advertise v2
cdp enable
cdp holdtime 180
cdp timer 60
cdp format device-id system-name

interface mgmt0
  cdp enable

interface Ethernet1/1
  cdp enable

interface Ethernet1/2
  cdp enable

interface Ethernet1/3
  cdp enable
…
DC36-5# sh cdp neighbors interface mgmt 0
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute


Device-ID             Local Intrfce Hldtme Capability  Platform      Port ID
mgmt-sw3.interop.cisco.com
                    mgmt0          130    R S I     WS-C6504-E    Gig4/3
```

LLDP Configuration and Verification:

```
feature lldp

lldp timer 30
lldp holdtime 120
lldp reinit 2
lldp tlv-select port-description
lldp tlv-select system-name
lldp tlv-select system-description
lldp tlv-select system-capabilities
lldp tlv-select management-address
lldp tlv-select dcbxp
lldp tlv-select port-vlan

interface mgmt0
  lldp transmit
  lldp receive

interface Ethernet1/1
  lldp transmit
  lldp receive

interface Ethernet1/2
  lldp transmit
  lldp receive

interface Ethernet1/3
  lldp transmit
  lldp receive
…
```

```
DC36-5# sh lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID              Local Intf      Hold-time  Capability  Port ID
DC36-105.interop.cisco.com Eth1/1          120        BR          Ethernet1/1
DC36-105.interop.cisco.com Eth1/2          120        BR          Ethernet1/2
DC36-105.interop.cisco.com Eth1/3          120        BR          Ethernet1/3
```

3.7.1.3.3 Syslog

Syslog is used to record all network events on the DC36 test bed. Whenever possible,DC36 makes use of a separate management VRF for syslog.

Configuration and Verification:

```
logging server syslog.interop.cisco.com 7 use-vrf management facility local6

N3064# sh log server
Logging server:                 enabled
{syslog.interop.cisco.com}
        server severity:        debugging
        server facility:        local6
        server VRF:             management
```

3.7.1.3.4 SNMP

SNMP is used for system monitoring in DC36. Scripts are used to poll the systems asynchronously during the course of all DC36 test execution.

Configuration:

```
version 6.0(2)U2(1)
snmp-server user admin network-admin auth md5 0x390c81441d991e0ba96d533f3ad69e68
 priv 0x390c81441d991e0ba96d533f3ad69e68 localizedkey
snmp-server host 172.28.92.62 traps version 2c public
snmp-server enable traps callhome event-notify
snmp-server enable traps callhome smtp-send-fail
snmp-server enable traps cfs state-change-notif
snmp-server enable traps lldp lldpRemTablesChange
snmp-server enable traps cfs merge-failure
snmp-server enable traps aaa server-state-change
snmp-server enable traps upgrade UpgradeOpNotifyOnCompletion
snmp-server enable traps upgrade UpgradeJobStatusNotify
snmp-server enable traps feature-control FeatureOpStatusChange
snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended
snmp-server enable traps config ccmCLIRunningConfigChanged
snmp-server enable traps snmp authentication
snmp-server enable traps link cisco-xcvr-mon-status-chg
snmp-server enable traps vtp notifs
snmp-server enable traps vtp vlancreate
snmp-server enable traps vtp vlandelete
snmp-server enable traps bridge newroot
snmp-server enable traps bridge topologychange
snmp-server enable traps stpx inconsistency
snmp-server enable traps stpx root-inconsistency
snmp-server enable traps stpx loop-inconsistency
snmp-server community public group network-operator
snmp-server community private group network-admin
snmp-server community cisco group network-operator
```

3.7.1.3.5 NTP

NTP is used to synchronize the clocks on all DC36 devices to provide consistent timestamps on all network logs and events.

Configuration and Verification:

```
ntp distribute
ntp server 172.28.92.1 use-vrf management
ntp commit

N3064# sh ntp status
Distribution : Enabled
Last operational state: No session

N3064# sh ntp peer-status
Total peers : 1
* - selected for sync, + -  peer mode(active),
- - peer mode(passive), = - polled in client mode
    remote               local                st   poll   reach delay   vrf
-----------------------------------------------------------------------------
*172.28.92.1           0.0.0.0                8    64      377   0.00092 management
```

3.7.1.3.6 SPAN

SPAN has been enabled on DC36 Nexus 3048 and Nexus 3064 switches to provide packet captures to assist in network debugging. Packets sourced by CPU cannot be monitored in SPAN session on Nexus3048 Nexus 3064 switches (CSCul38909). The embedded Ethanalyzer tool can be used instead.

Configuration and Verification:

```
monitor session 1
  source interface port-channel11 both
  destination interface Ethernet1/50
  no shut

N3064# sh monitor session 1
   session 1
---------------
type              : local
state             : up
acl-name          : acl-name not specified
source intf       :
    rx            : Po11
    tx            : Po11
    both          : Po11
source VLANs      :
    rx            :
destination ports : Eth1/50

Legend: f = forwarding enabled, l = learning enabled
```

3.7.1.3.7 DNS

DNS has been enabled to provide name lookup in the DC36 network.

Configuration and Verification:

```
vrf context management
  ip domain-lookup
  ip domain-name interop.cisco.com
  ip domain-list cisco.com
```

```
  ip domain-list interop.cisco.com
  ip name-server 172.28.92.9 172.28.92.10


N3064# ping karo vrf management
PING karo.interop.cisco.com (172.28.92.48): 56 data bytes
64 bytes from 172.28.92.48: icmp_seq=0 ttl=62 time=1.631 ms
64 bytes from 172.28.92.48: icmp_seq=1 ttl=62 time=1.754 ms
64 bytes from 172.28.92.48: icmp_seq=2 ttl=62 time=1.578 ms
64 bytes from 172.28.92.48: icmp_seq=3 ttl=62 time=1.409 ms
64 bytes from 172.28.92.48: icmp_seq=4 ttl=62 time=1.374 ms

--- karo.interop.cisco.com ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.374/1.549/1.754 ms
```

3.7.1.3.8 UDLD

UDLD is used to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a device detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops. UDLD aggressive mode is used across the DC36 network.

Configuration:

```
feature udld

udld aggressive

N3064# sh udld neighbors
Port            Device Name    Device ID   Port ID         Neighbor State
-------------------------------------------------------------------------
Ethernet1/1     FOC1723R2W2    1           Ethernet1/9     bidirectional
Ethernet1/2     FOC1723R2W2    1           Ethernet1/10    bidirectional
Ethernet1/3     FOC1723R2W2    1           Ethernet1/11    bidirectional
```

3.7.1.3.9 MTU

In order to configure the MTU to handle jumbo frames in DC36 the following policy-map has to be applied.

Configuration and Verification:

```
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
system qos
  service-policy type network-qos jumbo

interface Ethernet1/1
  no switchport
  mtu 9216

N3064# sh int e1/1
Ethernet1/1 is up
 Dedicated Interface
  Hardware: 10/100/1000 Ethernet, address: 4403.a7a3.c441 (bia 4403.a7a3.c408)
  Internet Address is 36.102.11.102/24
  MTU 9216 bytes, BW 1000000 Kbit, DLY 10 usec

N3064# sh queuing interface ethernet 1/1
Ethernet1/1 queuing information:
```

```
    qos-group  sched-type  oper-bandwidth
        0       WRR            100
    qos-group 0
    HW MTU: 9216 (9216 configured)
    drop-type: drop, xon: 0, xoff: 0
    Statistics:
        Ucast pkts sent over the port           : 4893502892
        Ucast bytes sent over the port          : 8465016605054
        Mcast pkts sent over the port           : 0
        Mcast bytes sent over the port          : 0
        Ucast pkts dropped                      : 0
        Ucast bytes dropped                     : 0
        Mcast pkts dropped                      : 0
        Mcast bytes dropped                     : 0

    Pkts dropped by RX thresholds               : 0
    Bytes dropped by RX thresholds              : 0

N3064# show hardware internal bcm-usd info port-info
+------+-----------------------------------------------------------------------+
|      |                    U                         S                        |
|      |                    N                         P                        |
|      |                    T                         A                        |
|      |                    A       F                 N                        |
|      |             L   U  G       R                 N                        |
|      |             I   N  G       A                 I                        |
|      |             N   T  E  D    M                                        L |
|      | F           K   A  D  I    E                 T                   O  I |
|      | R  A   P    S   G     S                  L   R               M   P  N |
|      | O  S   O    C   G  P  A    M             O   E   P   P       A   E  K |
|      | N  I   R    A   E  R  R    A             O   E   A   A   A   X   R    |
|      | T  C   T    N   D  I  D    X     L    D  P       U   U   U          S |
|      |                    O             3    U      S   S   S   T   S   S  T |
|      | P  P   I    M   V  R  M    S          P  B   T   E   E   O   P   P  A |
|      | O  O   N    O   L  I  O    I     M    L  A   A           N   E   E  T |
| INTF | R  R   T    D   A  T  D    Z     T    E  C   T   T   R   E   E   E  U |
| NAME | T  T   F    E   N  Y  E    E     U    X  K   E   X   X   G   D   D  S |
+------ -- -- ----- -- ---- - --- ----- ----- -- --- --- --- --- --- --- --- --+
Eth1/1   1  7 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/2   2  8 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/3   3  5 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/4   4  6 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/5   5 11 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/6   6 12 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
Eth1/7   7  9 SGMII sw 4094 0 non  9234     0 fd dis fwd dis dis ena  1G  1G up
```

3.7.1.4 CoPP

CoPP is used to control the rate at which packets are allowed to reach the switch's CPU. DC36 testbed uses default CoPP for both Nexus 3048 and Nexus 3064.

```
N3064# sh policy-map type control-plane expand name copp-system-policy

  policy-map type control-plane copp-system-policy
    class copp-s-selfIp
      police pps 500
    class copp-s-default
      police pps 400
    class copp-s-l2switched
      police pps 200
    class copp-s-ping
      police pps 100
    class copp-s-l3destmiss
      police pps 100
```

```
    class copp-s-glean
      police pps 500
    class copp-s-l3mtufail
      police pps 100
    class copp-s-ttl1
      police pps 100
    class copp-s-ipmcmiss
      police pps 400
    class copp-s-l3slowpath
      police pps 100
    class copp-s-dhcpreq
      police pps 300
    class copp-s-dhcpresp
      police pps 300
    class copp-s-dai
      police pps 300
    class copp-s-igmp
      police pps 400
    class copp-s-routingProto2
      police pps 1300
    class copp-s-v6routingProto2
      police pps 1300
    class copp-s-eigrp
      police pps 200
    class copp-s-pimreg
      police pps 200
    class copp-s-pimautorp
      police pps 200
    class copp-s-routingProto1
      police pps 1000
    class copp-s-arp
      police pps 200
    class copp-s-ptp
      police pps 1000
    class copp-s-bfd
      police pps 350
    class copp-s-bpdu
      police pps 12000
    class copp-icmp
      police pps 200
    class copp-telnet
      police pps 500
    class copp-ssh
      police pps 500
    class copp-snmp
      police pps 500
    class copp-ntp
      police pps 100
    class copp-tacacsradius
      police pps 400
    class copp-stftp
      police pps 400
```

3.7.1.5 PFC

Priority Flow Control (PFC), also referred to as Class-based Flow Control, is a mechanism that prevents frame loss that can be caused by congestion. PFC functions on a per class-of-service (COS) basis- only traffic flows with certain classes of service can be flow controlled while other classes are allowed to operate normally. By default, PFC is set to Auto on all ports.

On DC36 testbed, spine and leaf switches (Nexus 3048/Nexus 3064) use PFC auto mode by default while class-of-service (COS) is configured to match a value of 3.

Attaching QOS policies to any connected port in Auto mode causes all other Auto ports to go to operational ON (CSCul28008).

By default, memory management unit buffer-reservation will allow at most 2 ports to be enabled for PFC. In order to allow more interfaces to be PFC enabled, it is necessary to increase the hardware memory management unit buffer-reservation size with the command *hardware profile pfc mmu buffer-reservation <Percentage of shared pool buffers to be reserved>(*CSCul41772/CSCul28008).

Configuration and Verification:

```
class-map type qos match-all TESTQ
  match cos 3
policy-map type qos TESTP
  class TESTQ
    set qos-group 3
class-map type network-qos TESTQ
  match qos-group 3
policy-map type network-qos TESTP
  class type network-qos TESTQ
    mtu 9216
    pause no-drop
  class type network-qos class-default
    mtu 9216
system qos
  service-policy type network-qos TESTP

hardware profile pfc mmu buffer-reservation 80

interface port-channel61
  service-policy type qos input TESTP


N3064# sh interface priority-flow-control
============================================================
Port               Mode Oper(VL bmap)  RxPPP      TxPPP
============================================================


Ethernet1/1        Auto On  (8)       0          0
Ethernet1/2        Auto On  (8)       0          0
Ethernet1/3        Auto On  (8)       0          0
Ethernet1/4        Auto On  (8)       0          0
Ethernet1/5        Auto On  (8)       0          0
Ethernet1/6        Auto On  (8)       0          0
Ethernet1/7        Auto On  (8)       0          0
Ethernet1/8        Auto On  (8)       0          0
Ethernet1/9        Auto On  (8)       0          0
Ethernet1/10       Auto On  (8)       0          0
Ethernet1/11       Auto On  (8)       0          0
Ethernet1/12       Auto On  (8)       0          0
Ethernet1/13       Auto On  (8)       0          0
Ethernet1/14       Auto On  (8)       0          0
Ethernet1/15       Auto On  (8)       0          0
Ethernet1/16       Auto On  (8)       0          0
Ethernet1/31       Auto Off           0          0
Ethernet1/32       Auto Off           0          0
Ethernet1/33       Auto Off           0          0
Ethernet1/34       Auto Off           0          0
Ethernet1/41       Auto Off           0          0
Ethernet1/42       Auto Off           0          0
Ethernet1/45       Auto On  (8)       0          0
```

```
Ethernet1/47       Auto Off          0          0
Ethernet1/48       Auto Off          0          0
Ethernet1/49       Auto On   (8)     0          0
```

3.7.1.6 ECMP for IPv4 and IPv6

3.7.1.6.1 ECMP hash-offset

To avoid ECMP polarization in a multi-tier ECMP topology, a different ECMP hash-offset should be configured on each tier. By default, Nexus 3048/Nexus 3064 will set *hardware ecmp hash-offset* to 0. To prevent ECMP polarization in the DC36 testbed, ECMP hash-offset is configured to 1 and 2 for spine and leaf layers respectively.

Configuration:

```
hardware ecmp hash-offset 1
```

3.7.1.6.2 ECMP for IPv6

IPv6 ECMP hardware multipath programming is not updated upon a link flap of an ECMP/next-hop interface. In order to work around this issue, *ipv6 nd na glean* must be configured for all IPv6 ECMP interfaces (CSCul51491/ CSCtz1117).

Configuration and Verification:

```
N3064# sh run int e1/1

!Command: show running-config interface Ethernet1/1
!Time: Tue Feb 18 13:05:00 2014

version 6.0(2)U2(1)

interface Ethernet1/1
  no switchport
  mtu 9216
  no ip redirects
  ip address 36.101.11.1/24
  ipv6 address 2001:36:101:11:1::1/64
  ipv6 nd na glean
  ip ospf cost 20
  ip pim sparse-mode

N3064# sh ipv6 rout 2001:136:101:11::/64
IPv6 Routing Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]

2001:136:101:11::/64, ubest/mbest: 16/0
    *via 2001:36:106:11:1::1, Eth1/31, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:21:2::2, Eth1/32, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:31:3::3, Eth1/33, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:41:4::4, Eth1/34, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:51:5::5, Eth1/1, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:52:5::5, Eth1/2, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:53:5::5, Eth1/3, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:54:5::5, Eth1/4, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:55:5::5, Eth1/5, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:56:5::5, Eth1/6, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:57:5::5, Eth1/7, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:58:5::5, Eth1/8, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:61:6::6, Po61, [20/0], 5d17h, bgp-36105, external, tag 36
```

```
    *via 2001:36:106:62:6::6, Po62, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:63:6::6, Po63, [20/0], 5d17h, bgp-36105, external, tag 36
    *via 2001:36:106:64:6::6, Po64, [20/0], 5d17h, bgp-36105, external, tag 36

N3064# sh forwarding ipv6 route 2001:136:101:11::/64

IPv6 routes for table default/base

2001:136:101:11::/64
           2001:36:106:61:6::6, port-channel61
           2001:36:106:62:6::6, port-channel62
           2001:36:106:63:6::6, port-channel63
           2001:36:106:64:6::6, port-channel64
           2001:36:106:51:5::5, Ethernet1/1
           2001:36:106:52:5::5, Ethernet1/2
           2001:36:106:53:5::5, Ethernet1/3
           2001:36:106:54:5::5, Ethernet1/4
           2001:36:106:55:5::5, Ethernet1/5
           2001:36:106:56:5::5, Ethernet1/6
           2001:36:106:57:5::5, Ethernet1/7
           2001:36:106:58:5::5, Ethernet1/8
           2001:36:106:11:1::1, Ethernet1/31
           2001:36:106:21:2::2, Ethernet1/32
           2001:36:106:31:3::3, Ethernet1/33
           2001:36:106:41:4::4, Ethernet1/34

N3064# sh system internal forwarding ipv6 route 2001:136:101:11::/64 | inc Dev
Dev: 1   2001:136:101:11::/64, Index: 0xa792018c
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x11
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x12
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x13
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x4
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x5
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x6
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x7
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x8
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0x9
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xa
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xb
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xc
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xd
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xe
Dev: 1         "                   "        Adj Index: 0x30d42     Egress Lif: 0xf
```

3.7.1.7 Debugging on the Broadcom Shell

Nexus 3000 offers a very powerful tool that allows an easy access to the Broadcom shell. This allows to access to a big variety of commands hence enhancing the debug capabilities of the chipset. These commands should be used with caution as they are backdoors to program the hardware and bypass NX-OS.

To access the Broadcom Shell:

```
DC33-102# test hardware internal bcm-usd bcm-diag-shell
Available Unit Numbers: 0
bcm-shell.0> help
Help: Type help "command" for detailed command usage
Help: Upper case letters signify minimal match

Commands common to all modes:
        ?                   Display list of commands
        ASSert              Assert
        BackGround          Execute a command in the background.
        BCM                 Set shell mode to BCM.
        BCMX                Set shell mode to BCMX.
```

```
        break               place to hang a breakpoint
        CASE                Execute command based on string match
        CD                  Change current working directory
        cint                Enter the C interpreter
        CONFig              Configure Management interface
        CONSole             Control console options
<TRUNCATED>
```

3.7.2 Image Upgrade and Downgrade

The Nexus 3048 and Nexus 3064 switches on DC36 make use of “install all” to upgrade/downgrade software images whenever possible, but upgrade will be disruptive anyway.

```
N3064# install all kickstart bootflash:n3000-uk9-kickstart.6.0.2.U2.0.8.bin system bootflash:n3000-
uk9.6.0.2.U2.0.8.bin

Verifying image bootflash:/n3000-uk9-kickstart.6.0.2.U2.0.8.bin for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/n3000-uk9-kickstart.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n3000-uk9.6.0.2.U2.0.8.bin.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive          none



Images will be upgraded according to following table:
Module             Image         Running-Version             New-Version  Upg-Required
------  ----------------  ----------------------  ----------------------  ------------
     1            system             6.0(2)U2(1)             6.0(2)U2(1)            no
     1         kickstart             6.0(2)U2(1)             6.0(2)U2(1)            no
     1              bios      v2.5.0(06/27/2013)      v2.5.0(06/27/2013)            no
     1         power-seq                    v4.1                    v4.1            no


Additional info for this installation:
--------------------------------------

Service "bfd" : BFD feature is enabled. Upgrade will be disruptive!!!
```

```
Do you want to continue with the installation (y/n)?  [n]y

Switch will be reloaded for disruptive upgrade.

Install is in progress, please wait.
Performing runtime checks.
SUCCESS
Setting boot variables.
SUCCESS
Performing configuration copy.
SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

3.7.3 Routing Design Overview

3.7.3.1 Unicast Routing Design

3.7.3.1.1 BGP Routing Design

The network is split into three layers: core, spine and leaf. The layers are logically connected to each other through eBGP, as shown in Figure 35. The N7K core layer in BGP AS 3 is shared with other DC3 networks (DC31, DC32, and DC33). The spine layer runs OSPF to provide inter-switch connectivity to support iBGP sessions. The leaf layer is divided into multiple BGP ASes. This BGP logical design is easier to configure, maintain and debug than full mesh ibgp, route reflector, or confederations; the core can consolidate these as private ASes if there is a need to advertise to other BGP exchanges.

The spine layer is eBGP connected to the ASes configured at the Leaf layer over both IPv4 and IPv6 address families (eBGP dual stack). The spine routers also inject the default route down to the leaf ASes for both IPv4 and IPv6 address families (default-originate). ECMP is enabled on both IPv4 and IPv6 address families (maximum-path 64) across the DC36 network.

The leaf layer represents different top of rack topologies that can be deployed. AS 36101 employs two Nexus 3048 in a vPC topology, using HSRP for gateway redundancy for nodes. AS 36103 employs a routed top of rack with N3048. AS 36105 employs two Nexus 3064 in a vPC topology, using HSRP for gateway redundancy for nodes. AS 36104 is used as a test tool rather than network under test. The Catalyst 6500 is divided into multiple VRFs, with each VRF representing an extra ToR in the network. The goal is to test increasing number of ToR supported by the spine layer.

Figure 35 DC36 BGP Logical Design

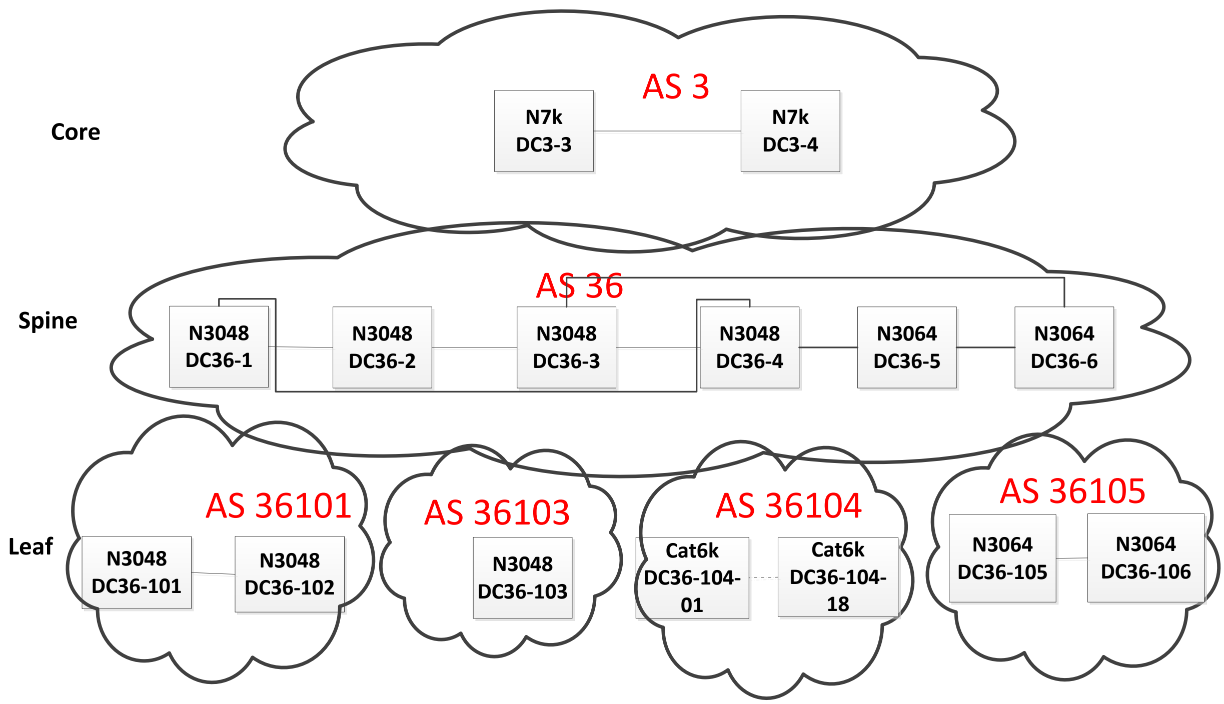

BGP peer templates are used to simplify configuration.

DC36 BGP Spine Configuration:

```
router bgp 36
  router-id 40.36.0.1
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
    network 36.101.11.0/24
…
    network 40.36.254.1/32
    maximum-paths 64
  address-family ipv6 unicast
    network 2001:1:40:36::1:0:1/128
…
    network 2001:36:114:9::/64
    maximum-paths 64
  template peer BGPLEAF
    bfd
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      default-originate
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      default-originate
      next-hop-self
      soft-reconfiguration inbound
  template peer BGPSPINE
    bfd
    remote-as 36
    password 3 a667d47acc18ea6b
```

```
    update-source loopback0
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  neighbor 36.101.11.101 remote-as 36101
    inherit peer BGPLEAF
  …
  neighbor 36.114.18.104 remote-as 36104
    inherit peer BGPLEAF
  neighbor 40.36.0.2
    inherit peer BGPSPINE
  neighbor 40.36.0.3
    inherit peer BGPSPINE
  neighbor 40.36.0.4
    inherit peer BGPSPINE
  neighbor 40.36.0.5
    inherit peer BGPSPINE
  neighbor 40.36.0.6
    inherit peer BGPSPINE
  neighbor 40.36.31.15 remote-as 3
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  neighbor 40.36.41.17 remote-as 3
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
```

DC36 BGP Leaf Configuration:

```
router bgp 36101
  graceful-restart-helper
  log-neighbor-changes
  address-family ipv4 unicast
    network 36.101.11.0/24
    …
    network 136.101.110.0/24
    maximum-paths 64
  address-family ipv6 unicast
    network 2001:136:101:100::/64
    …
    network 2001:36:101:61::/64
    maximum-paths 64
  template peer BGPLEAF
    bfd
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound
    address-family ipv6 unicast
      next-hop-self
      soft-reconfiguration inbound
  template peer BGPSPINE
    bfd
    remote-as 36
    password 3 a667d47acc18ea6b
    address-family ipv4 unicast
      soft-reconfiguration inbound
    address-family ipv6 unicast
      soft-reconfiguration inbound
  neighbor 36.101.11.1
    inherit peer BGPSPINE
  …
```

```
  neighbor 36.101.61.6
    inherit peer BGPSPINE
  neighbor 136.101.1.3 remote-as 36101
    inherit peer BGPLEAF
```

3.7.3.1.1.1 BGP Router-Id

To establish BGP sessions between peers, BGP must have a router ID, which is sent to BGP peers in the OPEN message when a BGP session is established. On DC36, NVT has configured a loopback interface IP address as the BGP router-id. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

To Verify the BGP Router-ID:

```
DC36-1# sh ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 59144, local router ID is 40.36.0.1
```

3.7.3.1.1.2 BGP Address Family

BGP address family for IPv4 and Ipv6 have been configured to achieve BGP peering, load-balancing, default route injection.

To Verify the BGP Address Family:

```
DC36-1# sh ip bgp all summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 40.36.0.1, local AS number 36
BGP table version is 59144, IPv4 Unicast config peers 45, capable peers 45
584 network entries and 5045 paths using 295044 bytes of memory
BGP attribute entries [12/1632], BGP AS path entries [5/30]
BGP community entries [0/0], BGP clusterlist entries [0/0]
5001 received paths for inbound soft reconfiguration
5001 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
36.101.11.101   4 36101   22858   25822    59144    0    0     1w2d 137
36.101.12.101   4 36101   22858   25819    59144    0    0     1w2d 137
36.101.13.101   4 36101   22859   25816    59144    0    0     1w2d 137
36.101.14.101   4 36101   22857   25811    59144    0    0     1w2d 137


BGP summary information for VRF default, address family IPv6 Unicast
BGP router identifier 40.36.0.1, local AS number 36
BGP table version is 32935, IPv6 Unicast config peers 45, capable peers 27
477 network entries and 4375 paths using 259936 bytes of memory
BGP attribute entries [10/1360], BGP AS path entries [4/24]
BGP community entries [0/0], BGP clusterlist entries [0/0]
4332 received paths for inbound soft reconfiguration
4332 identical, 0 modified, 0 filtered received paths using 0 bytes

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
```

```
36.101.11.101   4 36101   22858   25822   32935    0    0     1w2d 137
36.101.12.101   4 36101   22858   25819   32935    0    0     1w2d 137
36.101.13.101   4 36101   22859   25816   32935    0    0     1w2d 137
```

3.7.3.1.1.3 BGP Load Sharing and ECMP

DC36 has configured the maximum-paths that BGP adds to the route table for equal-cost multipath load balancing as 64 for both spine and leaf peers for IPv4/IPv6 address families.

3.7.3.1.1.4 BGP Authentication

DC36 has configured MD5 Authentication for BGP sessions.

To Verify the BGP Authentication:

```
DC36-1# sh ip bgp neighbors 36.114.18.104
BGP neighbor is 36.114.18.104,  remote AS 36104, ebgp link,  Peer index 45
  Inherits peer configuration from peer-template BGPLEAF
  BGP version 4, remote router ID 36.0.0.104
  BGP state = Established, up for 1w2d
  Peer is directly attached, interface Ethernet1/48
  BFD live-detection is configured and enabled, state is Invalid
  TCP MD5 authentication is enabled
```

3.7.3.1.1.5 BGP Update-Source

DC36 has configured BGP update-source to establish a BGP multi-hop sessions. DC36 has multi-hop sessions only on the iBGP peering between the spine switches.

To Verify the BGP Update-Source:

```
DC36-1# sh ip bgp neighbors 40.36.0.5
BGP neighbor is 40.36.0.5,  remote AS 36, ibgp link,  Peer index 4
  Inherits peer configuration from peer-template BGPSPINE
  BGP version 4, remote router ID 40.36.0.5
  BGP state = Established, up for 2w2d
  Using loopback0 as update source for this peer

DC36-1# sh ipv6 bgp neighbors 40.36.0.5
BGP neighbor is 40.36.0.5,  remote AS 36, ibgp link,  Peer index 4
  Inherits peer configuration from peer-template BGPSPINE
  BGP version 4, remote router ID 40.36.0.5
  BGP state = Established, up for 5d18h
  Using loopback0 as update source for this peer
```

3.7.3.1.1.6 BGP Default Route

The BGP default route is advertised from the spine peers to the leaf peers for both Ipv4 and Ipv6 address families.

To Verify the BGP Default Route:

```
DC36-1# sh ip bgp neighbors 36.101.11.101 | beg "For address family"
  For address family: IPv4 Unicast
```

```
  BGP table version 1907, neighbor version 1907
  137 accepted paths consume 7124 bytes of memory
  341 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 00:01:25 after session start

  For address family: IPv6 Unicast
  BGP table version 1690, neighbor version 1690
  137 accepted paths consume 7124 bytes of memory
  338 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.101.11.1
  Default information originate, default sent
  Last End-of-RIB received 00:01:25 after session start

  Local host: 36.101.11.1, Local port: 40115
  Foreign host: 36.101.11.101, Foreign port: 179
  fd = 49
```

3.7.3.1.1.7 BGP Next-Hop-Self

BGP next-hop-self is configured for BGP sessions between the spine switches for both IPv4 and IPv6 address families.

To Verify the BGP Next-Hop-Self:

```
DC36-1# sh ip bgp neighbors 36.114.18.104 | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 59144, neighbor version 59144
  8 accepted paths consume 416 bytes of memory
  480 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.114.18.1
  Default information originate, default sent

  For address family: IPv6 Unicast
  BGP table version 32935, neighbor version 0
  0 accepted paths consume 0 bytes of memory
  0 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.114.18.1
  Default information originate, default not sent
```

3.7.3.1.1.8 BGP Soft-Reconfiguration

BGP Soft reset is recommended because it allows routing tables to be reconfigured and activated without clearing the BGP session. Soft reset is done on a per-neighbor basis.

```
DC36-1# sh ip bgp neighbors 36.114.18.104 | beg "For address family"
  For address family: IPv4 Unicast
  BGP table version 59144, neighbor version 59144
  8 accepted paths consume 416 bytes of memory
  480 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.114.18.1
  Default information originate, default sent

  For address family: IPv6 Unicast
```

```
  BGP table version 32935, neighbor version 0
  0 accepted paths consume 0 bytes of memory
  0 sent paths
  Inbound soft reconfiguration allowed
  Nexthop always set to local peering address, 36.114.18.1
  Default information originate, default not sent
```

3.7.3.1.2 OSPF Routing Design

OSPF/OSPFv3 is used as the IGP to provide reachability for establishing iBGP peering at the spine layer only. The OSPF/OSPFv3 process is enabled only on directly connected interfaces and the Loopback interface. All the OSPF enabled interfaces are in Area 0.0.0.0. Each OSPF network type is set to point-to-point to decrease OSPF neighbor setup latency. In order to improve OSPF convergence, SPF and LSA timers are throttled to (100 200 5000 and 50 100 300) respectively.

DC36 OSPF/OSPFv3 Configuration:

```
feature ospf
router ospf 36
  router-id 40.36.0.4
  log-adjacency-changes
  timers throttle spf 100 200 5000
  timers throttle lsa 50 100 300
  auto-cost reference-bandwidth 100000

interface loopback0
  ip router ospf 36 area 0.0.0.0

interface port-channel1
  ip ospf network point-to-point
  ip router ospf 36 area 0.0.0.0

interface port-channel2
  ip ospf network point-to-point
  ip router ospf 36 area 0.0.0.0

interface port-channel5
  ip ospf network point-to-point
  ip router ospf 36 area 0.0.0.0


feature ospfv3
router ospfv3 36
  router-id 40.36.0.4
  log-adjacency-changes
  auto-cost reference-bandwidth 100000

interface loopback0
  ipv6 router ospfv3 36 area 0.0.0.0

interface port-channel1
  ipv6 router ospfv3 36 area 0.0.0.0

interface port-channel2
  ipv6 router ospfv3 36 area 0.0.0.0
```

3.7.3.1.3 Unicast Forwarding Verification

On DC36 test bed Nexus 3048 and Nexus 3064, routing is performed using hardware forwarding engines. The following sequence of commands illustrates verification of the programming of a host on a directly connected subnet on the N3048/N3064.

Below are the Commands used to look at the Number of Routes in the Forwarding Table, the Host Table, and LPM:

```
DC36-101# show forwarding ipv4 route summary

IPv4 routes for table default/base

Cumulative route updates: 3355
Cumulative route inserts: 7977
Cumulative route deletes: 1888
Total number of routes: 2767
Total number of paths : 16990

Number of routes per mask-length:
  /0  : 1         /8  : 1         /24 : 556       /32 : 2209

DC36-101# show hardware profile status
Total LPM Entries = 8191.
Total Host Entries = 16384.
Reserved LPM Entries = 1024.
Max Host4/Host6 Limit Entries (shared)=  8192/4096*
Max Mcast Limit Entries = 4096.
Used LPM Entries (Total) = 1013.
Used IPv4 LPM Entries =  559.
Used IPv6 LPM Entries =  454.
Used IPv6 LPM_128 Entries =  18.
Used Host Entries in LPM (Total) = 2213.
Used Host4 Entries in LPM = 2213.
Used Host6 Entries in LPM = 0.
Used Mcast Entries = 0.
Used Mcast OIFL Entries = 1.
Used Host Entries in Host (Total) = 388.
Used Host4 Entries in Host = 0.
Used Host6 Entries in Host = 388.
Max ECMP Table Entries = 64.
Used ECMP Table Entries = 2.
MFIB prefer-source-tree = Disabled/0/0.

*Unicast Host Table is in shared mode b/n v4 & v6...
```

This Command is Showing Directly Connected Subnet on Vlan 11: 136.101.11.2/24 and Ethernet1/1:

```
DC36-101# sh run int vlan 11

!Command: show running-config interface Vlan11
!Time: Fri Feb 14 11:17:51 2014

version 6.0(2)U2(1)

interface Vlan11
  no shutdown
  mtu 9216
  no ip redirects
  ip address 136.101.11.2/24
  ipv6 address 2001:136:101:11::2/64
  ip ospf passive-interface
  ip router ospf 36101 area 0.0.141.5
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
```

```
    preempt delay minimum 120
    priority 101
    ip 136.101.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
  ip 2001:136:101:11::1

  DC36-101# sh run int e1/1

  !Command: show running-config interface Ethernet1/1
  !Time: Fri Feb 14 14:03:27 2014

  version 6.0(2)U2(1)

  interface Ethernet1/1
    no switchport
    mtu 9216
    logging event port link-status
    no ip redirects
    ip address 36.101.11.101/24
    ipv6 address 2001:36:101:11:101::101/64
    ipv6 nd na glean
    ip ospf cost 20
    ip router ospf 36101 area 0.0.141.5
    ip pim sparse-mode
```

The Host 136.101.11.51 has been Learned via ARP on this Subnet:

```
DC36-101# sh ip arp 136.101.11.51

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface
136.101.11.51   00:09:32  0088.650b.3300  Vlan11
```

"show ip route" Shows Directly Connected Host as /32 Routes:

```
DC36-101# sh ip route 136.101.11.51/32
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

136.101.11.51/32, ubest/mbest: 1/0, attached
    *via 136.101.11.51, Vlan11, [250/0], 00:13:44, am
```

"show ip fib route" Shows Directly Connected Host as /32 Routes in FIB Table:

```
DC36-101# sh ip fib route 136.101.11.51/32

IPv4 routes for table default/base

------------------+------------------+----------------------+-----------------
Prefix            | Next-hop         | Interface            | Labels
------------------+------------------+----------------------+-----------------
136.101.11.51/32     136.101.11.51      Vlan11
```

“sh forwarding ipv4 route” Shows Directly Connected Host as /32 routes in Forward table:

```
DC36-101# sh forwarding ipv4 route 136.101.11.51/32

IPv4 routes for table default/base

------------------+------------------+---------------------+-----------------
Prefix            | Next-hop         | Interface           | Labels
------------------+------------------+---------------------+-----------------
136.101.11.0/24      Attached           Vlan11
```

Directly Connected Host Entries are Programmed as Adjacencies for Programming in the FIB Table:

```
DC36-101# sh ip adjacency 136.101.11.51/32

Flags: # - Adjacencies Throttled for Glean
       G - Adjacencies of vPC peer with G/W bit

IP Adjacency Table for VRF default
Total number of entries: 1
Address         MAC Address     Pref Source     Interface
136.101.11.51   0085.650b.3300  50   arp        Vlan11
```

Display Adjacency Index for this Route in Hardware Table:

```
DC36-101# sh system internal forwarding ipv4 route 136.101.11.51 module 1

Routes for table default/base

----+---------------------+----------+----------+-----------
Dev | Prefix              | PfxIndex | AdjIndex | LIF
----+---------------------+----------+----------+-----------
 1    136.101.11.51/32      0xaad3b4b0    0x187cf    0x7f
```

Display DMAC Entry Programmed in Adjacency Table:

```
DC36-101# sh system internal forwarding adjacency entry 0x187cf
Device: 1   Index: 0x187cf   dmac: 0085.650b.3300 smac: 4403.a7a3.bdfc                  e-lif: 0x7f
```

Find the PO Interface on which this MAC Address is Learnt:

```
DC36-101# sh mac address-table address 0085.650b.3300
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0085.650b.3300    dynamic   14450      F    F  Po11
```

Display PO11 Member Interface Information:

```
DC36-101# sh port-channel summary | in Po11
11    Po11(SU)    Eth      LACP      Eth1/39(P)
```

The Same Commands are Used to Troubleshoot LPM Table on N3048/N3064:

```
DC36-101# sh ip route 40.36.250.1/24
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
```

```
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

40.36.250.0/24, ubest/mbest: 34/0
    *via 36.101.11.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.12.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.13.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.14.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.15.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.16.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.17.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.18.1, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.21.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.22.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.23.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.24.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.25.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.26.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.27.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.28.2, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.31.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.32.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.33.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.34.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.35.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.36.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.37.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.38.3, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.41.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.42.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.43.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.44.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.45.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.46.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.47.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.48.4, [20/0], 4d04h, bgp-36101, external, tag 36
    *via 36.101.51.5, [20/0], 4d04h, bgp-36101, external, tag 36
*via 36.101.61.6, [20/0], 4d04h, bgp-36101, external, tag 36

DC36-101# sh ip fib route 40.36.250.1/24

IPv4 routes for table default/base

------------------+------------------+---------------------+-----------------
Prefix            | Next-hop         | Interface           | Labels
------------------+------------------+---------------------+-----------------
*40.36.250.0/24     36.101.11.1        Ethernet1/1
                    36.101.12.1        Ethernet1/2
                    36.101.13.1        Ethernet1/3
                    36.101.14.1        Ethernet1/4
                    36.101.15.1        Ethernet1/5
                    36.101.16.1        Ethernet1/6
                    36.101.17.1        Ethernet1/7
                    36.101.18.1        Ethernet1/8
                    36.101.21.2        Ethernet1/9
                    36.101.22.2        Ethernet1/10
                    36.101.23.2        Ethernet1/11
                    36.101.24.2        Ethernet1/12
                    36.101.25.2        Ethernet1/13
                    36.101.26.2        Ethernet1/14
                    36.101.27.2        Ethernet1/15
                    36.101.28.2        Ethernet1/16
                    36.101.31.3        Ethernet1/17
                    36.101.32.3        Ethernet1/18
                    36.101.33.3        Ethernet1/19
                    36.101.34.3        Ethernet1/20
                    36.101.35.3        Ethernet1/21
```

```
                  36.101.36.3       Ethernet1/22
                  36.101.37.3       Ethernet1/23
                  36.101.38.3       Ethernet1/24
                  36.101.41.4       Ethernet1/25
                  36.101.42.4       Ethernet1/26
                  36.101.43.4       Ethernet1/27
                  36.101.44.4       Ethernet1/28
                  36.101.45.4       Ethernet1/29
                  36.101.46.4       Ethernet1/30
                  36.101.47.4       Ethernet1/31
                  36.101.48.4       Ethernet1/32
                  36.101.51.5       Ethernet1/51
                  36.101.61.6       Ethernet1/52

DC36-101# sh forward ip route 40.36.250.1/24

IPv4 routes for table default/base

------------------+-----------------+--------------------+-----------------
Prefix            | Next-hop        | Interface          | Labels
------------------+-----------------+--------------------+-----------------
*40.36.250.0/24     36.101.11.1       Ethernet1/1
                    36.101.12.1       Ethernet1/2
                    36.101.13.1       Ethernet1/3
                    36.101.14.1       Ethernet1/4
                    36.101.15.1       Ethernet1/5
                    36.101.16.1       Ethernet1/6
                    36.101.17.1       Ethernet1/7
                    36.101.18.1       Ethernet1/8
                    36.101.21.2       Ethernet1/9
                    36.101.22.2       Ethernet1/10
                    36.101.23.2       Ethernet1/11
                    36.101.24.2       Ethernet1/12
                    36.101.25.2       Ethernet1/13
                    36.101.26.2       Ethernet1/14
                    36.101.27.2       Ethernet1/15
                    36.101.28.2       Ethernet1/16
                    36.101.31.3       Ethernet1/17
                    36.101.32.3       Ethernet1/18
                    36.101.33.3       Ethernet1/19
                    36.101.34.3       Ethernet1/20
                    36.101.35.3       Ethernet1/21
                    36.101.36.3       Ethernet1/22
                    36.101.37.3       Ethernet1/23
                    36.101.38.3       Ethernet1/24
                    36.101.41.4       Ethernet1/25
                    36.101.42.4       Ethernet1/26
                    36.101.43.4       Ethernet1/27
                    36.101.44.4       Ethernet1/28
                    36.101.45.4       Ethernet1/29
                    36.101.46.4       Ethernet1/30
                    36.101.47.4       Ethernet1/31
                    36.101.48.4       Ethernet1/32
                    36.101.51.5       Ethernet1/51
                    36.101.61.6       Ethernet1/52

DC36-101# sh system internal forward ip route 40.36.250.1/24

Routes for table default/base

----+--------------------+----------+----------+-----------
Dev | Prefix             | PfxIndex | AdjIndex | LIF
----+--------------------+----------+----------+-----------
 1    40.36.250.0/24       0xa78f9cec   0x30d40    0x9
 1         "                  "       0x30d40    0xa
 1         "                  "       0x30d40    0x4
 1         "                  "       0x30d40    0x8
 1         "                  "       0x30d40    0xb
```

```
 1          "                    "       0x30d40     0x12
 1          "                    "       0x30d40     0xf
 1          "                    "       0x30d40     0x14
 1          "                    "       0x30d40     0x5
 1          "                    "       0x30d40     0x16
 1          "                    "       0x30d40     0x11
 1          "                    "       0x30d40     0xc
 1          "                    "       0x30d40     0x6
 1          "                    "       0x30d40     0xd
 1          "                    "       0x30d40     0x7
 1          "                    "       0x30d40     0x1f
 1          "                    "       0x30d40     0x22
 1          "                    "       0x30d40     0x19
 1          "                    "       0x30d40     0x20
 1          "                    "       0x30d40     0x1b
 1          "                    "       0x30d40     0x1a
 1          "                    "       0x30d40     0x15
 1          "                    "       0x30d40     0x17
 1          "                    "       0x30d40     0xe
 1          "                    "       0x30d40     0x10
 1          "                    "       0x30d40     0x13
 1          "                    "       0x30d40     0x18
 1          "                    "       0x30d40     0x23
 1          "                    "       0x30d40     0x21
 1          "                    "       0x30d40     0x1e
 1          "                    "       0x30d40     0x24
 1          "                    "       0x30d40     0x1c
 1          "                    "       0x30d40     0x1d
 1          "                    "       0x30d40     0x25
```

The LPM Table Content can be Accessed by Broadcom Shell Command:

```
N3064K# test hardware internal bcm-usd bcm-diag-shell
bcm-shell.0> l3 l3table show
Unit 0, free L3 table entries: 14250
Entry VRF IP address       Mac Address           INTF MOD PORT    CLASS HIT

bcm-shell.0> l3 defib show
<TRUNCATED>
3403  1        33.101.47.0/24       00:00:00:00:00:00 102180    0     0     0    0 n
3404  1        33.101.48.0/24       00:00:00:00:00:00 102180    0     0     0    0 n
3404  1        133.108.5.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3405  1        133.101.39.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3405  1        133.101.25.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3406  1        133.106.3.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3406  1        133.109.6.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3407  1        133.101.108.0/24     00:00:00:00:00:00 100003    0     0     0    0 n
3407  1        133.101.57.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3408  1        133.112.9.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3408  1        133.110.7.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3409  1        133.101.43.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3409  1        133.101.75.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3410  1        133.111.8.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3410  1        133.107.4.0/24       00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3411  1        133.101.61.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3411  1        133.101.49.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3412  1        133.116.13.0/24      00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)

3412  1        133.115.12.0/24      00:00:00:00:00:00 200000    0     0     0    0 y      (ECMP)
```

```
3413  1        133.101.12.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
3413  1        133.101.93.0/24      00:00:00:00:00:00 100003    0     0     0    0 n
<TRUNCATED>

bcm-shell.0> l3 egress show
<TRUNCATED>
102612  00:00:2f:65:af:cb   67   67     1t   0        -1   no   no
102613  00:00:2f:64:fa:6c  119  119     1t   0        -1   no   no
102614  00:00:2f:64:fa:66  119  119     1t   0        -1   no   no
102615  00:00:2f:65:af:bb   67   67     1t   0        -1   no   no
<TRUNCATED>

bcm-shell.0> ipmc table show
<TRUNCATED>
133.101.42.44   230.33.0.8        97   -1 1  -1 0    1     0 y
133.101.42.41   230.33.0.8        97   -1 1  -1 0    1     0 y
133.101.11.45   230.33.0.5       127   -1 1  -1 0    1     0 y
133.101.12.43   230.33.0.4        76   -1 1  -1 0    1     0 y
133.101.51.43   230.33.0.8        85   -1 1  -1 0    1     0 y
133.101.52.45   230.33.0.8       132   -1 1  -1 0    1     0 y
<TRUNCATED>
```

3.7.3.2 Multicast Routing Design

Multicast feature is not enabled on DC36.

3.7.4 Layer-2/ Layer-3 Leaf/Access Layer Network Design Overview

3.7.4.1 vPC

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco NX-OS switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

On the DC36 test bed, vPC is configured between two Nexus 3048 switches and two Nexus 3064 switches.

vPC Peer Configurations:

N3000 1:

```
DC36-101# sh run vpc

version 6.0(2)U2(1)
feature vpc

vpc domain 601
  peer-keepalive destination 1.1.1.1 source 1.1.1.2
vrf vpc-keepalive
  peer-gateway
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel200
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/33
```

N3000 2:

```
DC36-102# sh run vpc

version 6.0(2)U2(1)
feature vpc

vpc domain 601
  peer-keepalive destination 1.1.1.2 source 1.1.1.1
vrf vpc-keepalive
  peer-gateway
  auto-recovery
  ip arp synchronize

! vpc peer-link config
interface port-channel200
  switchport mode trunk
  switchport trunk allowed vlan 1,11-110
  spanning-tree port type network
  vpc peer-link

! vpc peer-link member config
interface Ethernet1/33
```

<pre> switchport mode trunk
  switchport trunk allowed vlan 1,10-110
  channel-group 200 mode active
 	
! vpc peer-keepalive config
interface Ethernet1/35
    no switchport
  vrf member vpc-keepalive
  ip address 1.1.1.1/24
 	
! vpc member port-channel config
interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  spanning-tree port type edge trunk
  vpc 11
 	
! vpc member port config
interface Ethernet1/39
  switchport mode trunk
  switchport trunk allowed vlan 11-20
  channel-group 11 mode active</pre>
	<pre> switchport mode trunk
	  switchport trunk allowed vlan 1,10-110
	  channel-group 200 mode active
	
	! vpc peer-keepalive config
	interface Ethernet1/35
	  no switchport
	  vrf member vpc-keepalive
	  ip address 1.1.1.2/24
	
	! vpc member port-channel config
	interface port-channel11
	  switchport mode trunk
	  switchport trunk allowed vlan 11-20
	  spanning-tree port type edge trunk
	  vpc 11
	
	! vpc member port config
	interface Ethernet1/39
	  switchport mode trunk
	  switchport trunk allowed vlan 11-20
	  channel-group 11 mode active</pre>

Display vPC Status:

```
DC36-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.7.4.1.1 LACP

DC36 makes use of LACP mode active/active for vPC peer link and vPC leg link aggregation.

Display Port Channels and Link Aggregation Protocol Information:

```
DC36-102# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
11    Po11(SU)    Eth      LACP      Eth1/39(P)
21    Po21(SU)    Eth      LACP      Eth1/40(P)
31    Po31(SU)    Eth      LACP      Eth1/41(P)
41    Po41(SU)    Eth      LACP      Eth1/42(P)
51    Po51(SU)    Eth      LACP      Eth1/43(P)
61    Po61(SU)    Eth      LACP      Eth1/44(P)
71    Po71(SU)    Eth      LACP      Eth1/45(P)
81    Po81(SU)    Eth      LACP      Eth1/46(P)
91    Po91(SU)    Eth      LACP      Eth1/47(P)
101   Po101(SU)   Eth      LACP      Eth1/48(P)
200   Po200(SU)   Eth      LACP      Eth1/50(P)    Eth1/51(P)
DC33-102# show lacp interface ethernet 1/39
Interface Ethernet1/39 is up
  Channel group is 11 port channel is Po11
  PDUs sent: 44463
  PDUs rcvd: 48063
  Markers sent: 0
  Markers rcvd: 0
  Marker response sent: 0
  Marker response rcvd: 0
  Unknown packets rcvd: 0
  Illegal packets rcvd: 0
Lag Id: [ [(7f9b, 0-23-4-ee-bf-2d, 800b, 8000, 127), (8000, 0-1e-f6-e7-6c-0, b,
8000, 228)] ]
Operational as aggregated link since Wed Jan 29 11:38:49 2014

Local Port: Eth1/39   MAC Address= 0-23-4-ee-bf-2d
  System Identifier=0x8000,0-23-4-ee-bf-2d
  Port Identifier=0x8000,0x127
  Operational key=32779
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
  Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=(Ac-1:To-1:Ag-1:Sy-0:Co-0:Di-0:De-0:Ex-0)
Actor Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
Neighbor: 0x228
  MAC Address= 0-1e-f6-e7-6c-0
  System Identifier=0x8000,  Port Identifier=0x8000,0x228
  Operational key=11
  LACP_Activity=active
  LACP_Timeout=Long Timeout (30s)
  Synchronization=IN_SYNC
  Collecting=true
  Distributing=true
Partner Admin State=(Ac-0:To-1:Ag-0:Sy-0:Co-0:Di-0:De-0:Ex-0)
Partner Oper State=(Ac-1:To-0:Ag-1:Sy-1:Co-1:Di-1:De-0:Ex-0)
```

3.7.4.1.2 VLAN Trunking

DC36 makes use of VLAN trunking to provide security and segregation. Cisco devices make use of some VLANs for internal use. These VLANs must not be used externally by the network.

Display vlan information for Nexus 3000:

```
DC36-102# sh vlan internal usage

VLANs                  DESCRIPTION
-------------------    -----------------
3968-4031              Multicast
4032-4035              Online Diagnostic
4036-4039              ERSPAN
4042                   Satellite
3968-4047,4094         Current
DC33-102# show vlan id 11

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
11   VLAN0011                         active    Po11, Po200, Eth1/33, Eth1/34
                                                Eth1/39, Eth1/49, Eth1/50
                                                Eth1/51

VLAN Type  Vlan-mode
---- ----- ----------
11   enet  CE

Primary  Secondary  Type             Ports
-------  ---------  ---------------  -------------------------------------------
```

3.7.4.1.3 Spanning Tree

vPC technology helps build a loop free topology by leveraging port-channels from access devices to the vPC domain. A port-channel is seen as a logical link from the spanning tree's standpoint, so a vPC domain with vPC-attached access devices forms a star topology at Layer 2 (there are no STP blocked ports in this type of topology). In this case, STP is used as a fail-safe mechanism to protect against any network loops.

DC36 makes use of Rapid-PVST which is the default spanning tree protocol on NX-OS. For networks with larger logical port counts, MST is recommended.

Display Spanning Tree Information:

```
DC36-102# sh spanning-tree vlan 11

VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    8203
             Address     4403.a7a3.bdfc
             Cost        2
             Port        4295 (port-channel200)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8203   (priority 8192 sys-id-ext 11)
             Address     b0fa.eb5f.dafc
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po11             Desg FWD 1         128.4106 (vPC) P2p Peer(STP)
Po200            Root FWD 2         128.4295 (vPC peer-link) Network P2p
Eth1/49          Desg FWD 2         128.177  Edge P2p
```

```
DC36-102# sh spanning-tree summary totals
Switch is in rapid-pvst mode
Root bridge for: none
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short
STP-Lite                                 is enabled

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
101 vlans                     0         0        0        302        302
```

Display L2 Table-VLAN and L2 Table-STG Table Information from Broadcom Shell:

```
bcm-shell.0> vlan show 11
vlan 11 ports cpu,ge38,xe0-xe2 (0x0000000000000000000000000000000000000000000000000000e00800000001),
untagged none (0x0000000000000000000000000000000000000000000000000000000000000000) MCAST_FLOOD_UNKNOWN
bcm-shell.0> dump vlan 11
VLAN.ipipe0[11]:
<VP_GROUP_BITMAP=0,VLAN_PROFILE_PTR=3,VLAN_CLASS_ID=0x1c,VIRTUAL_PORT_EN=0,VALID=1,UUC_TRILL_NETWORK_RECEI
VERS_PRESENT=0,UUC_IDX=0,UMC_TRILL_NETWORK_RECEIVERS_PRESENT=0,UMC_IDX=0,TRILL_TRANSIT_IGMP_MLD_PAYLOAD_TO
_CPU=1,TRILL_RBRIDGE_NICKNAME_INDEX=0,TRILL_DOMAIN_NONUC_REPL_INDEX=0,TRILL_ACCESS_RECEIVERS_PRESENT=0,STG
=0x63,SRC_PVLAN_PORT_TYPE=0,SERVICE_CTR_IDX=0x64,PORT_BITMAP_W2=0,PORT_BITMAP_W1=0xe0080,PORT_BITMAP_W0=1,
PORT_BITMAP=0xe00800000001,L2_ENTRY_KEY_TYPE=0,ING_PORT_BITMAP_W2=0,ING_PORT_BITMAP_W1=0xe0080,ING_PORT_B
ITMAP_W0=1,ING_PORT_BITMAP=0xe00800000001,HIGIG_TRUNK_OVERRIDE_PROFILE_PTR=0,FID_ID=0xb,EVEN_PARITY_1=0,E
VEN_PARITY_0=0,ENABLE_IGMP_MLD_SNOOPING=0,BC_TRILL_NETWORK_RECEIVERS_PRESENT=0,BC_IDX=0>
bcm-shell.0> stg stp 63
STG 63:
    Block: ge32-ge33,ge38-ge40,ge42-ge47,xe3
  Forward: ge0-ge31,ge34-ge37,ge41,xe0-xe2
```

Display L2 Table-L2UserEntry:

```
DC36-102(config)# show hardware internal bcm-usd info tables l2 l2-user-entry all slot-num 0 | exclude
0000.0000.0000
Slot number 0

                    [ L2_USER_ENTRY (all info) - B549 TABLE ]

+----+----+--+-+-+----+---+----+----+-+-----------------+-+----+--------------+
|    |  P |  | | |             |    | |                 | |    |              |
|    | ER |  |D| |             |D   | |                 | |    |              |
|    | VO |  |O| |             |S   | |                 | |    |              |
|    | ET |  |N| |             |T   | |                 | |    |              |
|    | NO |  |T| |             |    | |                 | |    |              |
|    |  C |C | | |             |D   | |                 |K|    |              |
|    | PO |L |L| |             |I   | |                 |E|    |              |
|    |VAL |A |R|T|             |S   | |                 |Y|    |              |
|    |AR B|S |N|R|    TRUNK    |C   | |                 |T|    |              |
|    |LIPP|S | |U|=============|ACR |P|                 |Y|    |              |
|    |ITKD|  |S|N|     MOD PORT|RPPL|R|                 |P|    |              |
|ADDR|DYTU|ID|A|K|TGID ID   NUM |DUE3|I|MASK             |E|VLAN|  MAC ADDRESS |
+----+----+--+-+-+----+---+----+----+-+-----------------+-+----+--------------+
    0 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
    1 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0100.0ccc.ccc
    2 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0100.0ccc.ccc
    3 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
    4 1000  1 0 0      16    0 0100  0 1000ffffffffffff 0    0  0180.c200.000
```

3.7.4.1.4 Configuration Parameters Consistency

After the vPC feature is enabled and the vPC peer-link on both peer devices is configured, Cisco Fabric Services messages provide a copy of the local vPC peer device configuration to the remote vPC peer device. The systems then determine whether any of the crucial configuration parameters differ on the two devices.

When a Type 1 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 1 consistency check failure, all vPC member ports are set to down state.
- For a vPC interface configuration Type 1 consistency check failure, the misconfigured vPC is set to down state.

When a Type 2 consistency check failure is detected, the following actions are taken:

- For a global configuration Type 2 consistency check failure, all vPC member ports remain in up state and vPC systems trigger protective actions.
- For a vPC interface configuration Type 2 consistency check failure, the misconfigured vPC remains in up state. However, depending on the discrepancy type, vPC systems will trigger protective actions. The most typical misconfiguration deals with the allowed VLANs in the vPC interface trunking configuration. In this case, vPC systems will disable the vPC interface VLANs that do not match on both sides.

Display vPC Consistency Parameters:

```
DC36-102# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- ----------------------
QoS                         2     ([], [], [], [], [],   ([], [], [], [], [],
                                  [], [], [])            [], [], [])
Network QoS (MTU)           2     (9216, 0, 0, 0, 0, 0,  (9216, 0, 0, 0, 0, 0,
                                  0, 0)                  0, 0)
Network Qos (Pause)         2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Network Qos (WRED)          2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Network Qos (ECN)           2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
                                  F)                     F)
Output Queuing (Bandwidth)  2     (100, 0, 0, 0, 0, 0,   (100, 0, 0, 0, 0, 0,
                                  0, 0)                  0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F, F,  (F, F, F, F, F, F, F,
Priority)                         F)                     F)
STP Mode                    1     Rapid-PVST             Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
IGMP Snooping Group-Limit   2     8000                   8000
Interface-vlan admin up     2     1,11-110               1,11-110
Interface-vlan routing      2     1,11-110               1,11-110
capability
Allowed VLANs               -     1,11-110               1,11-110
Local suspended VLANs       -     -                      -
```

```
DC33-102# show vpc consistency-parameters interface port-channel 11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Shut Lan                    1     No                     No
STP Port Type               1     Edge Trunk Port        Edge Trunk Port
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-bf-2d, 800b, 0-23-4-ee-bf-2d, 800b,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-1e-f6-e7-6c-0, b, 0, 0-1e-f6-e7-6c-0, b, 0,
                                   0)]                    0)]
mode                        1     active                 active
Speed                       1     1000 Mb/s              1000 Mb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
Admin port mode             1
vPC card type               1     Empty                  Empty
Allowed VLANs               -     11-20                  11-20
Local suspended VLANs       -     -                      -
```

3.7.4.1.5 vPC Role Priority

There are two defined vPC roles: primary and secondary. The vPC role defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDUs) and responds to Address Resolution Protocol (ARP).
In case of a tie (same role priority value defined on both peer devices), the lowest system MAC will dictate the primary peer device.

Display vPC Role, System-MAC, System-Priority:

```
DC36-102# show vpc role

vPC Role status
----------------------------------------------------
vPC role                        : primary, operational secondary
Dual Active Detection Status    : 0
vPC system-mac                  : 00:23:04:ee:bf:2d
vPC system-priority             : 32667
vPC local system-mac            : b0:fa:eb:5f:da:fc
vPC local role-priority         : 201
```

3.7.4.1.6 vPC Peer-Link

The vPC peer-link is a standard 802.1Q trunk that performs the following actions:

- Carry vPC and non-vPC VLANs.
- Carry Cisco Fabric Services (CFS) messages that are tagged with CoS=4 for reliable communication CoS=4 for reliable communication.
- Carry flooded traffic between the vPC peer devices.
- Carry STP BPDUs, HSRP hello messages, and IGMP updates.

When the vPC peer-link fails and the vPC peer-keepalive link is still up, the vPC secondary peer device performs the following operations:

- Suspends its vPC member ports
- Shuts down the SVI associated to the vPC VLAN

Display vPC Peer-link Information:

```
DC36-102# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po200  up     1,11-110

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.7.4.1.7 vPC Peer-Keepalive Link

The vPC peer-keepalive link is a Layer 3 link that joins one vPC peer device to the other vPC peer device and carries a periodic heartbeat between those devices. It is used at the boot up of the vPC systems to guarantee that both peer devices are up before forming the vPC domain. It is also used when the vPC peer-link fails, in which case, the vPC peer-keepalive link is leveraged to detect split brain scenario (both vPC peer devices are active-active).

Default Values for VPC Peer-Keepalive Links:

Timer	Default value
Keepalive interval	1 seconds
Keepalive hold timeout (on vPC peer-link loss)	3 seconds
Keepalive timeout	5 seconds

Display vPC Peer-Keepalive Information:

```
DC36-102# sh vpc peer-keepalive

vPC keep-alive status             : peer is alive
--Peer is alive for             : (1334740) seconds, (86) msec
--Send status                   : Success
--Last send at                  : 2014.02.13 22:23:09 342 ms
--Sent on interface             : Eth1/35
--Receive status                : Success
--Last receive at               : 2014.02.13 22:23:08 889 ms
--Received on interface         : Eth1/35
--Last update from peer         : (0) seconds, (453) msec

vPC Keep-alive parameters
--Destination                   : 1.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc-keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

3.7.4.1.8 vPC Member Link

As suggested by the name, a vPC member port is a port-channel member of a vPC. A port-channel defined as a vPC member port always contains the keywords *vpc <vpc id>.*

A vPC only supports Layer 2 port-channels. The port-channel can be configured in access or trunk switchport mode. Any VLAN allowed on the vPC member port is by definition called a vPC VLAN. Whenever a vPC VLAN is defined on a vPC member port, it must also be defined on the vPC peer-link. Not defining a vPC VLAN on the vPC peer-link will cause the VLAN to be suspended.

The configuration of the vPC member port must match on both the vPC peer devices. If there is an inconsistency, a VLAN or the entire port channel may be suspended (depending on Type-1 or Type-2 consistency check for the vPC member port). For instance, a MTU mismatch will suspend the vPC member port.

Display vPC Member Port-channel Information:

```
DC36-102# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
```

```
-------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ -------------------------------------------------
1    Po200  up     1,11-110

vPC status
--------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
DC36-102# show vpc consistency-parameters interface port-channel 11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- ----------------------
Shut Lan                    1     No                     No
STP Port Type               1     Edge Trunk Port        Edge Trunk Port
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-bf-2d, 800b, 0-23-4-ee-bf-2d, 800b,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-1e-f6-e7-6c-0, b, 0, 0-1e-f6-e7-6c-0, b, 0,
                                   0)]                    0)]
mode                        1     active                 active
Speed                       1     1000 Mb/s              1000 Mb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
MTU                         1     1500                   1500
Admin port mode             1
vPC card type               1     Empty                  Empty
Allowed VLANs               -     11-20                  11-20
Local suspended VLANs       -     -                      -
```

3.7.4.1.9 vPC ARP Synchronization

The vPC ARP Synchronization feature improves the convergence time for Layer 3 flows (North to South traffic). When the vPC peer-link fails and subsequently recovers, vPC ARP Synchronization performs an ARP bulk synchronization over Cisco Fabric Services (CFS) from the vPC primary peer device to the vPC secondary peer device.

Displays vPC ARP Synchronization Information:

```
DC36-102# sh ip arp sync-entries

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Address         Age       MAC Address     Interface
136.101.52.51   00:00:25  0088.6534.3300  Vlan52
136.101.52.52   00:00:25  0088.6534.3301  Vlan52
136.101.52.53   00:00:25  0088.6534.3302  Vlan52
136.101.52.54   00:00:25  0088.6534.3303  Vlan52
…
```

3.7.4.1.10 vPC Delay Restore

After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from the access to the core until the Layer 3 connectivity is reestablished.

The vPC Delay Restore feature delays the vPCs leg bringup on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on the vPC leg. The result provides a graceful restoration along with zero packet loss during the recovery phase (traffic still gets diverted to the alive vPC peer device).

This feature is enabled by default with a vPC restoration default timer of 30 seconds, which DC36 maintains in the testbed.

3.7.4.1.11 vPC Auto-Recovery

vPC auto-recovery feature was designed to address 2 enhancements to vPC.

- To provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure (vPC auto-recovery feature).
- To handle a specific case where both vPC peer devices reload but only one comes back to life (vPC auto-recovery reload-delay feature).

The switch which unsuspends its vPC role with vPC auto-recovery continues to remain primary even after peer-link is on. The other peer takes the role of secondary and suspends its own vPC until a consistency check is complete. Therefore, to avoid this situation from occurring erroneously, auto-recovery reload-delay-timer should be configured to be long enough for the system to fully complete its bootup sequence.

Helpful Commands for vPC Object Tracking:

Show vpc brief	Displays Auto-recovery status

Configuration Check:

```
DC36-102# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 301
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 10
Peer Gateway                      : Enabled
Peer gateway excluded VLANs     : -
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

vPC Peer-link status
```

```
-------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ -------------------------------------------------
1    Po200  up     1,11-110

vPC status
--------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-20
21     Po21        up     success     success                    21-30
31     Po31        up     success     success                    31-40
41     Po41        up     success     success                    41-50
51     Po51        up     success     success                    51-60
61     Po61        up     success     success                    61-70
71     Po71        up     success     success                    71-80
81     Po81        up     success     success                    81-90
91     Po91        up     success     success                    91-100
101    Po101       up     success     success                    101-110
```

3.7.4.1.12 HSRP/HSRPv6 Active/Active with vPC

HSRP in the context of vPC has been improved from a functional and implementation standpoint to take full benefits of the L2 dual-active peer devices nature offered by vPC technology. HSRP operates in active-active mode from a data plane standpoint, as opposed to classical active/standby implementation with a STP based network. No additional configuration is required. As soon as a vPC domain is configured and interface VLAN with an associated HSRP group is activated, HSRP will behave by default in active/active mode (on the data plane side).

From a control plane standpoint, active-standby mode still applies for HSRP in context of vPC; the active HSRP instance responds to ARP request. ARP response will contain the HSRP vMAC which is the same on both vPC peer devices. The standby HSRP vPC peer device just relays the ARP request to active HSRP peer device through the vPC peer-link.

HSRPv4&v6 Configurations:

N3000 1:

```
interface Vlan11
  no shutdown
  mtu 9216
  no ip redirects
  ip address 133.101.11.2/24
  ipv6 address 2001:133:101:11::2/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
    ip 133.101.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
    ip 2001:133:101:11::1
```

N3000 2:

```
interface Vlan11
  no shutdown
  mtu 9216
  no ip redirects
  ip address 133.101.11.3/24
  ipv6 address 2001:133:101:11::3/64
  ip pim sparse-mode
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 133.101.11.1
  hsrp 101 ipv6
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 2001:133:101:11::1
```

Helpful Commands for HSRP Active/Active with vPC:

Show hsrp brief	Displays hsrp status
Show mac address-table vlan <vlan id>	Displays mac addresses including HSRP vMAC;

	check for G-flag on vMAC for active/active HSRP

Configuration Check:

```
DC36-102# sh hsrp brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr      Standby addr     Group addr
Vlan11      1   99   P Standby  133.101.11.2     local            133.101.11.1
  (conf)
Vlan11      101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local            fe80
::5:73ff:fea0:65 (impl auto EUI64)
<TRUNCATED>
Vlan109     1   99   P Standby  133.101.109.2    local            133.101.109.1
  (conf)
Vlan109     101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local            fe80
::5:73ff:fea0:65 (impl auto EUI64)
Vlan110     1   99   P Standby  133.101.110.2    local            133.101.110.1
  (conf)
Vlan110     101 99   P Standby  fe80::4603:a7ff:fea3:bdfc  local            fe80
::5:73ff:fea0:65 (impl auto EUI64)
```

3.7.4.2 L2 TCAM Tables

Nexus 3000/3548 platforms display MAC age as “seconds since first seen.” This behavior differs from the Nexus 5000, 6000 and 7000 platforms which are displayed as “seconds since last seen” and should be taken into account when reading the table (CSCun37434).

When topology change notifications or MAC address clears are initiated on the Nexus 3000 the ARP address table also gets flushed (CSCun32115). As a result, the ARP table will be re-learned.

List of Useful Commands for the TCAM Table:

Display L2 Table-L2EntryTable:

```
DC36-102(config)# sh mac address-table vlan 11
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since first seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 11       0000.0028.2a0c    dynamic   40         F    F  Po11
* 11       0085.650b.2900    dynamic   40         F    F  Po11
* 11       0085.650b.2901    dynamic   2320       F    F  Po11
* 11       0085.650b.2902    dynamic   40         F    F  Po11
<TRUNCATED>
* 11       0085.650b.33c5    dynamic   2320       F    F  Po11
* 11       0085.650b.33c6    dynamic   40         F    F  Po11
* 11       0085.650b.33c7    dynamic   2320       F    F  Po11
  11       0100.5e21.0001    igmp      0          F    F  Po11 Po200
<TRUNCATED>
  11       0100.5e21.0009    igmp      0          F    F  Po11 Po200
  11       0100.5e21.000a    igmp      0          F    F  Po11 Po200
```

Display L2 Table-L2EntryTable:

```
DC36-102(config)# show platform fwm info hw-stm
VLAN   MAC Address      Port    PC
------+----------------+-------+---------
32     00:85:65:20:33:aa 2      Y[Po-1355349985
12     00:85:65:0c:33:58 0      Y[Po-1355354805
11     00:85:65:0b:33:22 0      Y[Po-1355354805
```

```
22    00:85:65:16:33:89 1        Y[Po-1355349995
42    00:85:65:2a:33:56 3        Y[Po-1355349975
<TRUNCATED>
```

Display L2 Table-L2EntryTable from Broadcom Shell:

```
bcm-shell.0> l2 show
<TRUNCATED>
mac=00:85:65:0b:33:50 vlan=11 GPORT=0x0 Trunk=0 Hit
mac=00:85:65:29:33:22 vlan=41 GPORT=0x0 Trunk=3 Hit
mac=00:85:65:34:33:89 vlan=52 GPORT=0x0 Trunk=4 Hit
mac=00:85:65:1f:33:01 vlan=31 GPORT=0x0 Trunk=2
mac=01:00:5e:21:00:03 vlan=41 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2070
mac=00:00:2f:64:fa:86 vlan=21 GPORT=0x0 Trunk=1
mac=00:85:65:0b:33:7c vlan=11 GPORT=0x0 Trunk=0
mac=00:85:65:2a:33:08 vlan=42 GPORT=0x0 Trunk=3
mac=01:00:5e:21:00:04 vlan=32 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2064
mac=00:85:65:0c:33:06 vlan=12 GPORT=0x0 Trunk=0 Hit
mac=01:00:5e:21:00:04 vlan=11 GPORT=0x0 modid=0 port=0/cpu0 Static CPU MCast=2072
mac=00:85:65:0b:33:46 vlan=11 GPORT=0x0 Trunk=0 Hit
mac=00:85:65:2a:33:32 vlan=42 GPORT=0x0 Trunk=3
<TRUNCATED>
```

4. NVT Test Methodology

4.1 Test Cycle

The test cycle consists of the following steps:

1. Network configuration and verification.
2. Software and firmware upgrade and downgrade.
3. Trigger network disruptions.
4. Stress platform control-plane.
5. Check for CPU usage anomalies and memory leaks.

4.2 Network Disruption Test Cases

The following sections describe the test disruptions and the verification criteria:

- System Level
- Core Layer
- Spine Layer
- Leaf Layer

System Level:

Disruption	Verification
Image upgrade and rollback with ISSU	Hitless upgrade/rollback for all configured features with parallel enhancement

Core Layer:

Disruption	Verification
Router Link Failure/Recovery between Core and Edge	• IGP and PIM reconvergence (control-plane & data plane)
Member of Port-channel Failure/Recovery between Core and Edge	• Traffic load-sharing for port-channels • LACP interoperability • Unidirectional Link Detection (UDLD)
Clear IGP Neighbors/Process at Core	Stress test for control-plane recovery
Clear IPv4 Unicast Routes at Core	Stress test for control-plane recovery
Clear IPv4 Multicast Routes at Core	Stress test for control-plane recovery
Core Switch System Failure/Recovery	• IGP and PIM reconvergence (control-plane & data plane) • PIM Rendezvous Point redundancy & Back-up verification • VDC failure does not impact other VDCs
Core Switch Power Redundancy	Partial Power loss causes no impact to control/data plane
Core Switch Supervisor High-Availability	• NSF, GR, in-chassis and on peers

	• NSF interoperability
Core Switch Fabric High-Availability	Fabric module failure causes no impact to control/data plane
Line Card OIR at Core Switch	• Hitless operation for non-affected ports • Traffic load-sharing for distributed port-channels • IGP and PIM reconvergence (control-plane & data plane) • LACP interoperability for distributed port-channels • Unidirectional Link Detection (UDLD)

Aggregation or Spine Layer:

Disruption	**Verification**
Router Link Failure/Recovery between Aggregation and Core	• IGP and PIM reconvergence (control-plane & data plane)
Member of Port-channel Failure/Recovery between Aggregation and Core	• Traffic load-sharing for port-channels • LACP interoperability • Unidirectional Link Detection (UDLD)
Layer 2 Trunk Link Failure/Recovery between Aggregation and Access	• STP reconvergence • IGMP reprogramming with snooping • MAC address re-learning • Security ACL & FNF reprogramming • No FHRP impact • No ARP/ND impact • vPC functionality
FabricPath Core Link Failure/Recovery	• MAC address re-learning • No FHRP impact • No ARP/ND impact • FabricPath Functionality • vPC+ functionality
Member of Port-channel Failure/Recovery between Aggregation and Access	• Traffic load-sharing for port-channels • LACP interoperability • Unidirectional Link Detection (UDLD)
Clear IGP Neighbors/Process at Aggregation	Stress test for control-plane recovery
Clear IPv4 Unicast Routes at Aggregation	Stress test for control-plane recovery
Clear IPv4 Multicast Routes at Aggregation	Stress test for control-plane recovery
Aggregation Switch System Failure/Recovery	• STP reconvergence • IGP and PIM reconvergence (control-plane & data plane) • PIM Rendezvous Point redundancy & Back-up verification

	• PIM DR/BDR functionality • IGMP Snooping & Querier functionality • VDC failure does not impact other VDCs • Security ACL & FNF reprogramming • FHRP redundancy • MAC address learning • ARP/ND re-learning • vPC/vPC+ functionality • FabricPath functionality
Aggregation Switch Power Redundancy	Partial Power loss causes no impact to control/data plane
Aggregation Switch Supervisor High-Availability	• NSF, GR, in-chassis and on peers • NSF and GR interoperability • No impact to vPC peering status
Aggregation Switch Fabric High-Availability	Fabric module failure causes no impact to control/data plane
Line Card OIR at Aggregation Switch	• Hitless operation for non-affected ports • Traffic load-sharing for distributed port-channels • IGP and PIM reconvergence (control-plane & data plane) • LACP interoperability for distributed port-channels • Unidirectional Link Detection (UDLD)
vPC/vPC+ peer-link/keep-alive Failure/Recovery	vPC functionality and peering status
vPC/vPC+ Leg Failure/Recovery	• No impact to STP overlay • IGMP reprogramming with snooping • MAC address re-learning • Security ACL & FNF reprogramming • No FHRP impact • No ARP/ND impact
vPC/vPC+ Leg member Failure/Recovery	• Traffic load-sharing for port-channels • LACP interoperability • Unidirectional Link Detection (UDLD)

Leaf or Access/ToR Layer:

Disruption	**Verification**
Access/ToR Switch System Failure/Recovery	• STP reconvergence • IGMP snooping reprogramming • MAC address re-learning • No impact to other vPC/vPC+ • FabricPath functionality

Access/End-host Layer:

Disruption	Verification
Member of Port-channel Failure/Recovery between FI and upstream switches	• Verify FI uplink static pinning works as expected • Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning • LACP interoperability • Verify DHCP functionalities
Port-channel Failure/Recovery between FI and upstream switches	• Verify FI uplink static pinning works as expected • Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning • LACP interoperability • Verify DHCP functionalities
Port-channel Failure/Recovery between FI and IOM	• Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning
Cluster Link Failure/Recovery between FIs	• Traffic load-sharing for link members • Traffic load-sharing within the FI cluster • Recovery of system functionalities
Member of Cluster Link Failure/Recovery between FIs	• Traffic load-sharing for link members • Traffic load-sharing within the FI cluster • Recovery of system functionalities
Fabric Interconnect System Failure/Recovery	• Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning • vPC functionality/FabricPath functionality • LACP interoperability
Blade OIR	• Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities
NIC Bonding	• Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning
Service Profile Operations	• Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities
Software Upgrade/Downgrade	• Traffic load-sharing for port-channels

	• Traffic load-sharing within the FI cluster • Recovery of system functionalities
VMware® vMotion™	• Traffic load-sharing for port-channels • Traffic load-sharing within the FI cluster • Recovery of system functionalities • MAC address learning • Verify DHCP functionalities

Sample Test Case:

Sample Test Case	
Title	Link failure between aggregation and core layers
Description	Verify network control and data plane recovery after link flap
Test Setup	• Reference topology • Reference network configuration setup test case • Reference test plan for control and data plane setup matrices
Procedure	**1.** Fail one of the links between the aggregation and core layers. **2.** Recover the above link. **3.** Repeat the same test at least 5 iterations to ensure consistent behavior for the devices and network. **4.** Repeat the above procedures for the other links between the aggregation and core layers.
Pass/Fail Criteria	• During the link failure, traffic should drop in proportion to the number of links and paths affected, and the traffic should be able to reconverge within the expected time . • Ensure that the unicast and multicast routing protocols have detected peer failure in order to start network reconvergence within the expected time. • Verify the convergence pattern is as expected. • Verify the CPU usage pattern is as expected. • Verify the memory usage is as expected. • Verify the route tables for both unicast and multicast routing are updated correctly on all switches in the network. Ensure that only affected switches show change in the forwarding tables. • Verify the hardware forwarding entries, line card programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast routing are updated correctly on all switches in the network. • Verify Layer 2 forwarding tables on aggregation and access switches. They should not be affected by this failure.

4.3 Automation Methodology

NVT's test methodology employs a heavy emphasis on automation testing. The core functions of NVT's automation framework are to provide intelligence and repeatability of test case execution, to generate reporting and provide representation of test progress status in an easy to understand manner. Automated tests have the benefit of ensuring the reliability and repeatability of test runs.

4.3.1 Automated Test Cases

Figure 36 Automation Framework Overview

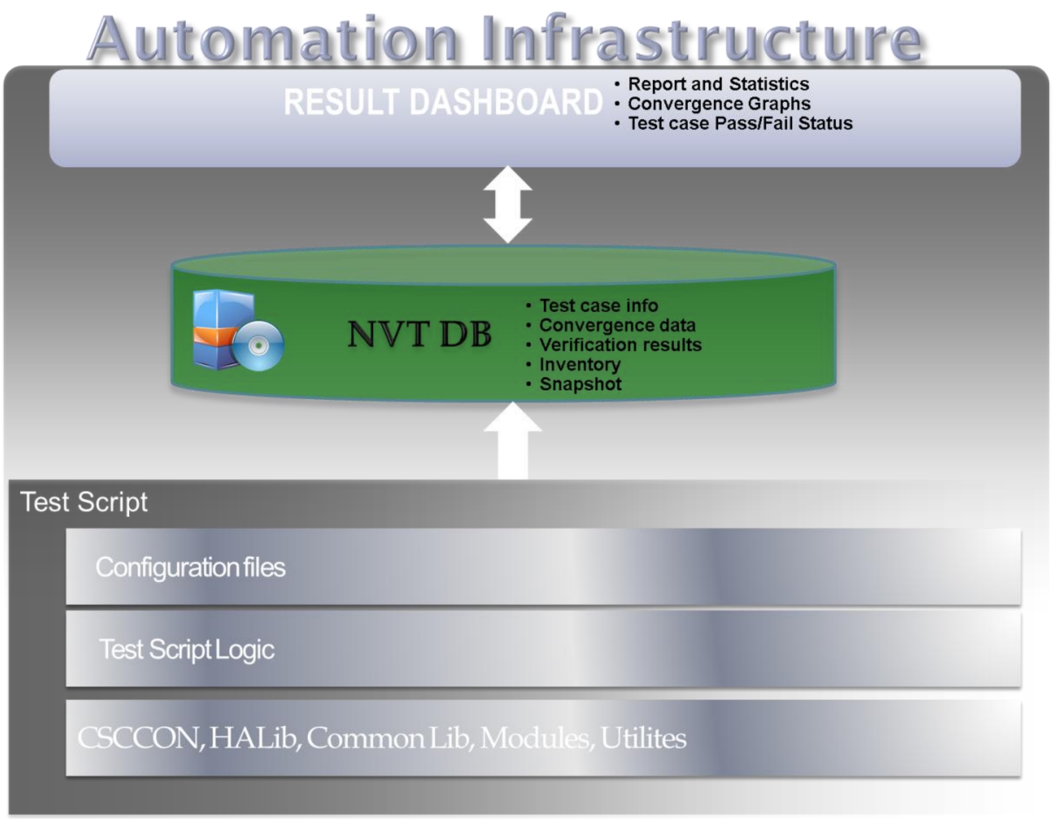

The automation framework has three key components:

- Test script suites - Test scripts are used to perform network failure executions and data collection.
- NVT database – The NVT database is a repository of test results.
- Results Dashboard - The results dashboard is responsible for presenting data to engineers for result analysis, calculating test case execution statistics and generating result reports.

Test scripts are composed of several key components:

- Initialization
- Failure execution
- Verification

- Data collection
- Pass/Fail analysis

During initialization, all necessary information about the device under test is collected before failure execution. The script connects to the device under test (DUT) and retrieves neighbor in formation (via CDP or LLDP), port-channel information and environment information. Detailed hardware, software and EPLD information is gathered from each device under test in the network. This includes hardware model types and software versions of all components in the device. This data is stored in the NVT's result repository. Running configuration of each DUT in the network is also collected and stored in the repository. Collecting inventory and running configuration information at the beginning of e ach test execution help ensure that the environment can be recreated which is an important factor for repeatability of test cases. In addition, test cases are executed with the same set of configuration files, sequence of failures and script timers which provide repeatability. The automation framework's repeatability helps demonstrate network reliability. Moreover, if a network issue is found while running the script, the reproducibility of the network issue can be increased by re -running the script with the same configuration.

The following is a list of test triggers supported by NVT automation today:

- Supervisor Switchover (applies to HA systems)
- Software Upgrade/Downgrade (depending on platform will be disruptive or ISSU/D)
- VDC suspend/unsuspend (applies to platform that support Virtual Device Contexts)
- Link Shut
- Link No-Shut
- Software Module power on
- Software Module power off
- Xbar power on
- Xbar power off
- Clear OSPF process
- Clear BGP process
- Clear Ip Routes
- Clear Ip Mroutes
- Clear PIM neighbors
- Reload Switch

The data collection script connects to a third party traffic generation tool (Ixia IxNetwork). Using the traffic tool's API, the script collects convergence data including aggregated port statistics and per -flow packet loss duration. Once data is collected, the results are analyzed in graphical format for convergence pattern and behavior. The graphs have a one second resolution and hence for sub -second convergence raw data is used. Traffic convergence results are represented graphically as shown be low:

Figure 37 Sample link disruption convergence graph

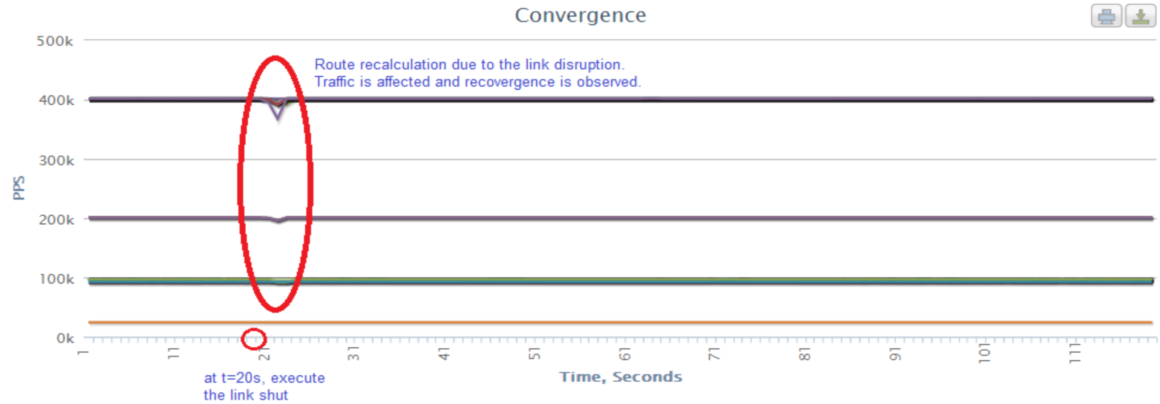

In the case of per-flow statistics, a post-processing script further analyzes the packet loss duration for each flow to determine the worst packet loss duration for the test run. The same failure type (ie same link on the same DUT) are categorized together to determine the mean, standard deviation, min and max over all the iterations of the same type. If the standard deviation is large, it calls for further investigation. In general, standard deviation is useful for determining if the DUT is capable of handling data plane and control plane convergence in a consistent manner.

Example of Packet Loss Duration Report

RX Port Name	Traffic Item Name	dut	failure type	failure property	Mean (ms)	Std Deviation (ms)	Min (ms)	Max (ms)
DC36-1002 vPC T9/46 IXIA 97-6/12 U	DC36-IPv4-Unicast-Meshed	DC36-101	LinkShut	DC36-101+e1!1	133.333	47.140	100.000	200.000

The above displays a packet loss report when one of the ECMP routed interfaces is shut on the leaf switch. It represents a particular traffic item or ixia receive port pair and its mean, standard deviation, min, and max for packet loss duration over three iterations.

4.3.2 Automated Test Case Verifications

Specific modules are developed to perform verification checks. For example, startup and running configuration consistency is verified before and after a switchover. Numerous verification criterias are predefined for each test case and the details can be found in section 6. The status (pass/fail) for a test case is determined from the outcome of all verifications for that test case. If any verification fails, then the overall status of the test case is marked as failed.

Example of Automated Verifications

Verification Description	CLI
Compare baseline running config and startup with existing config.	show startup-config/show running-config
Verify SSH works through the management network.	(verify by connecting via SSH)
Verify protocol (bgp/ospf/eigrp/isis/pim) neighbors before failure operation remain the same after failure operation recovers.	show ip <protocol> neighbors
Verify UDLD neighbors before failure operation remain the same after failure operation recovers.	show udld neighbors
Verify RSA key does not change on device. Verify RSA keys before failure operation remain the same after failure operation recovers.	show crypto key mypubkey rsa | excl generated
Verify CDP/LLDP neighbors before failure operation remain the same after failure operation recovers. Verify that CDP/LLDP peer is removed for disrupted link (if applicable).	show cdp neighbors/show lldp neighbors
Verify BFD neighbors before failure operation remain the same after failure operation recovers.	show bfd neighbors
Check for err-disabled interfaces.	sh interfaces status err-disabled
Verify any core dumps.	show core
Verify traffic deadflows (traffic flows not reaching destination).	(via Ixia API's)
Verify traffic tx/rx port rate before failure operation remains same after failure operation recovers.	(via Ixia API's)
Verify packet loss within accepted range.	(via Ixia API's)

A detailed report generated from the results of NVT test runs can be found in section 6.

4.4 Host/Server Configuration

This section describes the host/server configuration across the NVT testbed. It also describes the unicast and multicast traffic configuration.

4.4.1 DC1 Traffic Generator Configuration

Unicast Host(s) configuration:

- Block 1: Nexus 7000 vPC block consists of [20 VLANs x 120 hosts] spread across Layer 2 ToR devices and FEX.
- Block 2: Nexus 7000/5000 FabricPath block consists of [20 VLANs x 350 hosts] spread across Layer 2 devices attached to the leaf and spine of the FabricPath network.
- Blocks 3-7: These blocks consist of [20 VLANs x 100 hosts] in each block.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

Multicast Source Configuration:

- Block 1: Nexus 7000 vPC block consists of 100 hosts spread across Layer 2 ToR devices sourcing multicast traffic for groups with local RP.
- Block 2: Nexus 7000/5000 FabricPath block consists of 350 hosts spread across Layer 2 devices attached to the leaf and spine of the FabricPath network sourcing multicast traffic for groups with local RP.
- Blocks 3-7: These blocks consist of 100 hosts in each block sourcing multicast traffic for groups with local RP.

Multicast Receiver Configuration:

- Block 1: Nexus 7000 vPC block consists of [20 VLANs x 2hosts] spread across Layer 2 ToR devices and FEX joining all multicast groups sourced in the network.
- Block 2: Nexus 7000/5000 FabricPath block consists of [20 VLANs x 7 host] spread across Layer 2 devices attached to the leaf and spine of the FabricPath network joining all multicast groups sourced in the network.
- Blocks 3-7: These blocks consist of [20 VLANs x 2 hosts] in each block joining all multicast groups sourced in the network.

4.4.2 DC2 Traffic Generator Configuration

Unicast Hosts configuration:

- Block 1: Nexus 7000 vPC block consists of [20 VLANs x 120 hosts] spread across Layer 2 ToR devices and FEX.
- Block 2: Nexus 7000/5000 FabricPath block consists of [20 VLANs x 350 hosts] spread across Layer 2 devices attached to the leaf and spine of the FabricPath network.
- Blocks 3-7: These blocks consist of [20 VLANs x 100 hosts] in each block.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

Multicast Source Configuration:

- Block 1: Nexus 7000 vPC block consists of 100 hosts spread across Layer 2 ToR devices sourcing multicast traffic for groups with local RP.
- Block 2: Nexus 7000/5000 FabricPath block consists of 350 hosts spread across Layer 2 devices attached to the leaf and spine of the FabricPath network sourcing multicast traffic for groups with local RP.
- Blocks 3-7: These blocks consist of 100 hosts in each block sourcing multicast traffic for groups with local RP.

Multicast Receiver Configuration:

- Block 1: Nexus 7000 vPC block consists of [20 VLANs x 2 hosts] spread across Layer 2 ToR devices and FEX joining all multicast groups sourced in the network.
- Block 2: Nexus 7000/5000 FabricPath block consists of [20 VLANs x 7 hosts] spread across Layer 2 devices attached to the leaf and spine of the FabricPath network joining all multicast groups sourced in the network.
- Blocks 3-7: These blocks consist of [20 VLANs x 2 hosts] in each block joining all multicast groups sourced in the network.

4.4.3 DC3 Core Traffic Generator Configuration

Unicast Hosts configuration:

- Core: The Nexus 7000 core switch has been configured [1 VLAN x 100 hosts for each POD] to send and receive North-South IPv4 and IPv6 bidirectional unicast traffic.
- Leaf/Access: Each leaf consists of [10 VLANs x 200 hosts] to send and receive North-South IPv4 and IPv6 bidirectional unicast traffic.

Multicast Source Configuration:

- Core: The Nexus 7000 core switch has been configured with 1 VLAN x 2 hosts sourcing multicast traffic for 200 groups to send North-South IPv4 multicast traffic.

Multicast Receiver Configuration:

- Leaf/Access: Each leaf consists of 10 VLANs with 5 hosts x VLAN joining all the multicast groups sourced at the core.

4.4.4 DC31 Traffic Generator Configuration

Unicast Hosts configuration:

- Nexus 6000 Leafs: vPC leaf switches consist of 10 VLANs x 50 IPv4 hosts configured on each orphan port plus 50 IPv4 hosts configured at the vPC connected access switch, for a total of 150 hosts x VLAN.
- Nexus 6000 Leaf: Standalone leaf switch consists of 10 VLANs x 50 IPv4 hosts.
- Nexus 3548 Leaf: Standalone leaf switch consists of 10 VLANs x 50 IPv4 hosts.
- Nexus 3000 Leaf: Standalone leaf switch consists of 10 VLANs x 50 IPv4 hosts.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 50 IPv4 host.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

Multicast Source Configuration:

- Nexus 6000 Leafs: vPC leaf switches consist of 10 VLANs x 5 IPv4 hosts configured on each orphan port plus 5 IPv4 hosts configured at the vPC connected access switch. Each of the 15 hosts is sourcing multicast traffic for 10 groups.
- Nexus 6000 Leaf: Standalone leaf switch consists of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Nexus 3548 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Nexus 3000 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host sourcing multicast traffic for 10 groups.

Multicast Receiver Configuration:

- Nexus 6000 Leafs: vPC leaf switches consist of 10 VLANs x 5 IPv4 hosts configured on each orphan port plus 5 IPv4 hosts configured at the vPC connected access switch. Each of the 15 hosts is joining all the 10 multicast groups sourced from each leaf.
- Nexus 6000 Leaf: Standalone leaf switch consists of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.
- Nexus 3548 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.
- Nexus 3000 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host joining all the 10 multicast groups sourced from each leaf.

4.4.5 DC32 Traffic Generator Configuration

Unicast Hosts configuration:

- Nexus 3548 Leafs: Both classical STP access and standalone leaf switches consist of 10 VLANs x 200 IPv4 hosts.
- Nexus 3000 Leaf: Standalone leaf switch consist of 10 VLANs x 200 IPv4 hosts.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 200 IPv4 host.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

Multicast Source Configuration:

- Nexus 3548 Leafs: Both classical STP access and standalone leaf switches consist of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Nexus 3000 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host sourcing multicast traffic for 10 groups.

Multicast Receiver Configuration:

- Nexus 3548 Leafs: Both classical STP access and standalone leaf switches consist of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.

- Nexus 3000 Leaf: Standalone leaf switch consist of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.
- Nexus 7000 Leaf: The Nexus 7000 switch is configured with 10 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host joining all the 10 multicast groups sourced from each leaf.

4.4.6 DC33 Traffic Generator Configuration

Unicast Hosts configuration:

- Nexus 3000 Leafs: Both vPC and standalone leaf switches consist of 10 VLANs x 200 IPv4 hosts and 2 VLANs x 100 IPv6 hosts.
- Catalyst 6500 Leaf: The Catalyst switch is configured with 18 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

Multicast Source Configuration:

- Nexus 3000 Leafs: Both vPC and standalone leaf switches consist of 10 VLANs x 5 IPv4 hosts sourcing multicast traffic for 10 groups.
- Catalyst 6500 Leaf: The Catalyst switch is configured with 18 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host sourcing multicast traffic for 10 groups.

Multicast Receiver Configuration:

- Nexus 3000 Leafs: Both vPC and standalone leaf switches consist of 10 VLANs x 5 IPv4 hosts joining all the 10 multicast groups sourced from each leaf.
- Catalyst 6500 Leaf: The Catalyst switch is configured with 18 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host joining all the 10 multicast groups sourced from each leaf.

4.4.7 DC36 Traffic Generator Configuration

Unicast Hosts configuration:

- Nexus 3000 Leafs: Both vPC and standalone leaf switches consist of 10 VLANs x 200 IPv4 hosts and 5 VLANs x 200 IPv6 hosts.
- Catalyst 6500 Leaf: The Catalyst switch is configured with 18 VRFs. Each VRF consists of 1 VLAN x 1 IPv4 host.

Each host in the network sends traffic to all other hosts to form a fully meshed traffic configuration.

4.5 UCS Traffic Generation Tool

Current traffic generation software solutions are designed to provide users with a fully pre-defined system to generate network traffic and track flow statistics. The problem with this approach is that several pertinent layers within this system are not accounted for, thus reducing the flexibility and variability of user-defined configurations. Some solutions also lack an aggregation of detailed, traffic flow statistics and are difficult to manage with a scaled number of hosts. This is not optimal for large-scale networks composed of hundreds to thousands of different hosts.

The UCS Traffic Generation Tool resides on the Application Layer to simulate a "real world" application. This provides the user with a centrally-managed software solution that is scalable, portable, and provides the necessary functionality of generating and accurately measuring network traffic. The tool validates the various UCS hardware and software components during network disruptions by collecting and analyzing detailed, traffic flow statistics between hosts. These statistics are used by the NVT team to help determine the functionality and scalability of fully integrated data centers comprised of Cisco's Nexus switches and UCS blade servers.

The Traffic Generation tool generates UDP traffic across a user configured network by having a user send commands to one or more managers, which then appropriately forwards these commands to various virtual machines responsible for communicating with one another. The generated traffic can vary in several conditions, such as its rate, duration, packet size and host's outgoing interface.

The system is separated into three main components to be deployed on separate machines (whether physical or virtual): the traffic server, the VM Daemon, and the Reporter. These components will allow the user to generate network traffic across various network configurations and track the aggregated network traffic statistics.

As traffic is generated, the user receives periodic updates about the number of packets sent, number of packets received, number of packets dropped, and any out-of-order packets.

Figure 38 High Level Overview of the Traffic Generation Tool's Components and their Communication Flows

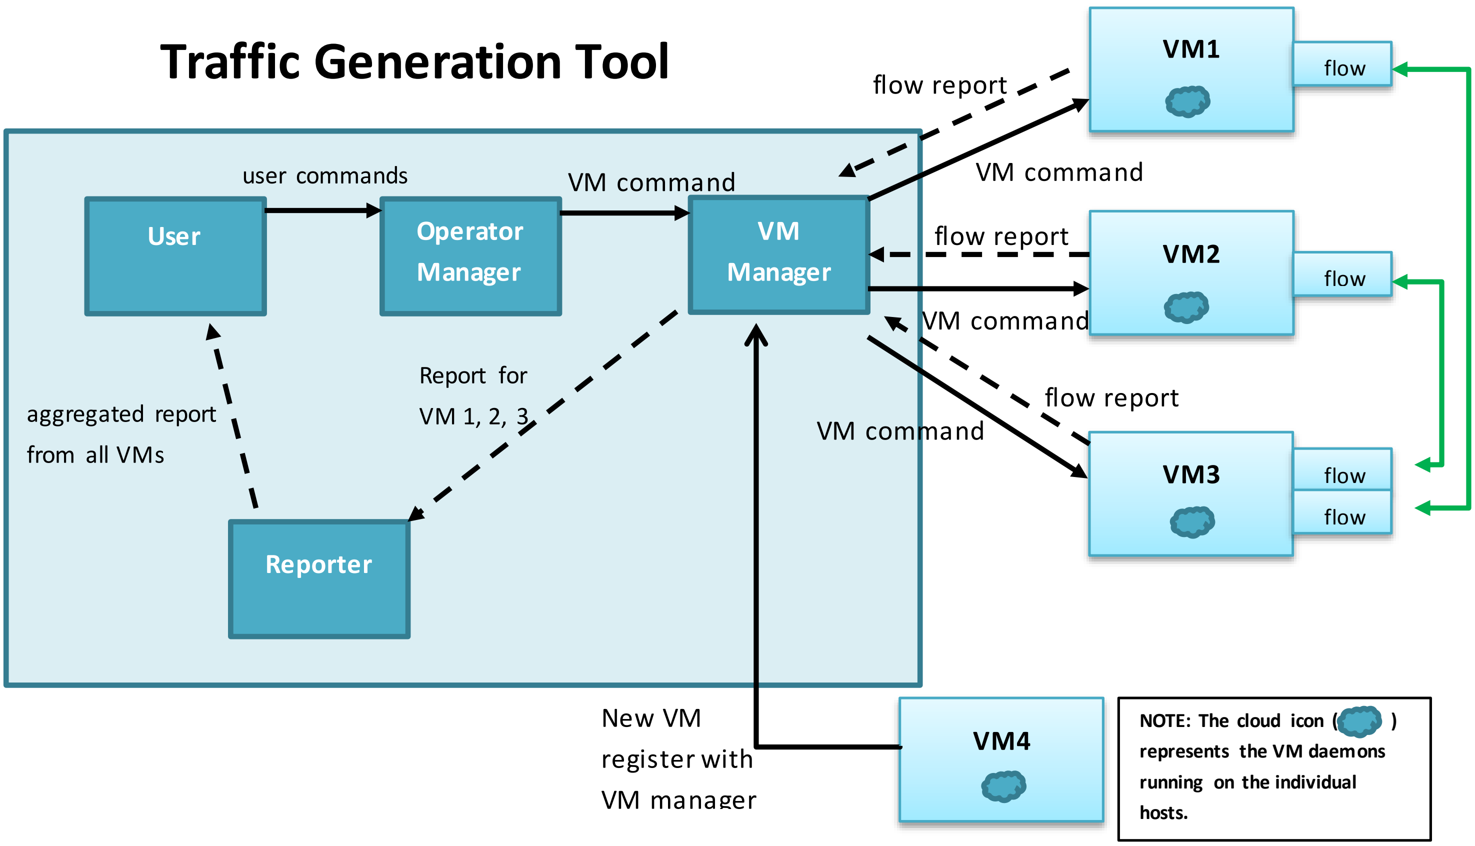

Each line represents communication between the components. The solid black lines represent commands from the user. The solid green lines represent connections between the VM Daemons. The dashed lines represent reported data being sent from the VM Daemons back to the user for analyzing. First, the data is sent from each daemon back to its respective manager. Then, the manager forwards that data to the Reporter for further processing, presentation to the user.

5. NVT Findings/Conclusion/Recommendations

Assigned/New	→	*Still working on fixes and may be seen in CCO image*
Unreproducible	→	*Not seen in CCO image, may be have fixed by other code fixes.*
Verified/Resolved	→	*Fixed in CCO image*
Closed	→	*System limitation and behavior will remain the same*

5.1 Caveats for DC1 and DC2

CSCuh90209/ CSCul48388
Symptom: ISSU gets stuck from 6.1.4.CCO to 6.2.5.55
Conditions: After initiating ISSU, ISSU gets stuck at the point where "cmp" version from system image of 6.2.5.55 is being extracted
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCui61039
Symptom: N7700: XBAR ASIC interrupt errors when XBAR is inserted
Conditions: An N7706 chassis is powered up without any spines. Once the spines are inserted and LC's come up with traffic, then for each subsequently inserted spine, xbar asic interrupt errors are seen on the console
Workaround: None
Severity: Moderate
Status: Assigned
Platform Seen: N7700
Resolved Releases:
Applicable Releases: 6.2(6)

CSCuj56624
Symptom: OIL is not programmed in MFDM
Conditions: This may be seen in a multicast environment after a device reload.
Workaround: Issuing either of the below commands will fix this issue:
#clear ip mroute <multicast group ip> - on DR for a particular group
#clear ip mroute * - on DR for all groups
Severity: Severe
Status: Resolved
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCuj79031/ CSCuj95182
Symptom: n7k-sup2: /var/tmp location filled by diag_port_lb.6158 file
Conditions: On N7k loaded with 6.2.5.33_S1, these messages are seen: "N7K %$ VDC-1 %$ %SYSMGR-2-TMP_DIR_FULL: System temporary directory usage is unexpectedly high at 100% ". This issue is because of diag_port_lb file filling up /var/tmp location.
Workaround: None
Severity: Moderate
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCuj92558
Symptom: In a vpc+ setup running f2 cards as part of both vpc peer reload ,CFS errors are seen: 'sw-226-54 %$ VDC-1 %$ %L2FM-2-L2FM_CFS_SEND_FAILED: cfs send failed, num 2'
Conditions: l2fm is trying to send data over peer-link event before peer-link is declared up, which is causing the failure
Workaround: None
Severity: Moderate
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCuj95402
Symptom: ethpm cores on VDC reload on 6.2.5.33_S1
Conditions: N7k with sup1 has 3 VDC's, two VDC's are in FabricPath. After doing a reload of a FabricPath VDC, the VDC failed to come online and ethpm cored.
Workaround: Not reproducible in the final images
Severity: Severe
Status: Unreproducible
Platform Seen: N7000
Resolved Releases:
Applicable Releases: 6.2(6)

CSCuj97300/ CSCul01126
Symptom: aclqos cores seen with M-1 module failure after a switch reboot
Conditions: aclqos crash seen on M1 module after switch is reloaded with 6.2(5.38)S0
Workaround: None
Severity: Moderate
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)

Applicable Releases:

CSCul06388
Symptom: ipqosmgr crashed while doing ISSU from 6.1.x to 6.2.6
Conditions: After doing ISSU from 6.1.x to 6.2.6, ipqosmgr core is seen on N7K
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul16225
Symptom: When switches, one N7706 and one N7710 when running 6.2.5.45.S1 have diag failures on all modules
Conditions: Diags fail on modules with error: %DIAG_PORT_LB-2-REWRITE_ENGINE_LOOPBACK_TEST_FAIL: Module:2 Test:RewriteEngine Loopback failed 10 consecutive times. Faulty module:Module 5 Error:Loopback test failed. Packets possibly lost on the switch SUP fabric
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N7700
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul18616
Symptom: Memory leaks observed in 'mtm' process on M1 module during MIB walks
Conditions: Memory leaks detected in ‘mtm’ process during MIB walk of CiscoProcessMIB and CiscoCBQosMIB
Workaround: Not reproducible in the final images
Severity: Minor
Status: Unreproducible
Platform Seen: N7000
Resolved Releases: None
Applicable Releases: 6.2(6)

CSCul20672/ CSCul81685
Symptom: ISSD Fails from 6.2.5.65.S2 to 6.2.2a with service vdc_mgr error.
Conditions: ISSD of 6.2.6 --> 6.2.2/6.2.2a - if "f3" shows up in either "limit-resource module-type" or "system module-type", then ISSD will abort with error: VDC_MGR has detected a potential issue and blocked upgrade (0x413C0017)(vdc: 1). System detected f3 in switchwide VDC mode("system module - type"), which is not supported in the version you are downgrading to. Please remove f3 from the relevant config before the downgrade"

Workaround: None
Severity: Moderate
Status: Resolved
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul26450
Symptom: rpm core seen during 'copy r s vdc-all', config copy is aborted
Conditions: After setting the boot string and doing a 'copy r s vdc-all' on N7700, rpm core is seen. Config copy is aborted after the core:
%SYSMGR-2-SERVICE_CRASHED: Service "rpm" (PID 7647) hasn't caught signal 6 (core will be saved).
%SYSMGR-2-CFGWRITE_ABORTED: Configuration copy aborted.
Workaround: None
Severity: Moderate
Status: Resolved
Platform Seen: N7700
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul28020
Symptom: "plugin" core is seen after "copy r s" is done on 6.2.5.48.S0 - N7K
Conditions: plugin core was seen on N7K, running version 6.2.5.48_S0. The core was seen after these series of steps: (1) Loading 6.2.5.48_S0 (previously running 6.2.5.33_S2) and doing a couple of system switchovers. (2) After 2nd switchover a "copy r s" was done (3) 'plugin' cored
Workaround: None
Severity: Severe
Status: Unreproducible
Platform Seen: N7000
Resolved Releases:
Applicable Releases:

CSCul30416
Symptom: ISSD Failure: Workaround suggested by NX-OS not working
Conditions: After initiating ISSD from 6.2.5.48 (S0) to 6.2.2.S42, pre-upgrade check fails with error which in-turn aborts the ISSD: Return code 0x41A10008 (Config check failure). Service "pltfm_config" in vdc 1: 'rate-limiter otv and/or netflow is configured for module <mod>'.This is not supported in the target version. Please issue the 'no hardware rate-limiter command to remove the module rate-limiters'
Workaround: Need to disable netflow and otv at hardware level. Command: N7K(config)# no hardware rate-limiter layer-2 netflow disable module x
Severity: Moderate
Status: Closed
Platform Seen: N7000
Resolved Releases: None

Applicable Releases: 6.2(6)

CSCul34953/ CSCul36654
Symptom: Packet loss will be seen after ISSU from 6.1.4/6.1.4a to 6.2.5.52.S0 on N7K
Conditions: After doing ISSU from 6.1.4/6.1.4a to 6.2.5.52.S0 image, ping between directly connected interfaces and also MGMT interface doesn't work due to which there is traffic loss.
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul44583
Symptom: On N5000 vpc peers there are 4 fex's on each peer. Fex downstream links from each of them have vpc connections to host (a cat6K switch). When either of the peer-switch is reloaded, some of the fex downstream links do not come once switch is back online.
Conditions: When one of the vpc peer's is reloaded, only some of the fex downstream ports comes up in VPC (to fanout switch), rest remain in I (Individual) state and at times in D (down) state. This issue is seen when host ports have scaled vlans (128 downstream fex host ports, each configured with 1000 vlans). If the vlans are scaled down – to 250 or 50 vlans per host port, this issue is not seen.
Workaround: Flap the fex downstream ports. Once the links are flapped, they come up in VPC.
Severity: Moderate
Status: New
Platform Seen: N5000
Resolved Releases:
Applicable Releases: 5.2(1)N1(4)

CSCul44598
Symptom: Intermittent traffic loss for hosts with spt-threshold infinity configured in a network which also has Sparse Mode hosts
Conditions: This issue is seen when all the following conditions are met:

- the last hop router with spt-threshold infinity and the Sparse Mode host have the common intermediate router
- common intermediate router is in the shared tree path for both the hosts and also in the (S,G,rpt) prune path from the Sparse Mode host while it sends joins to the source tree

Workaround: Make shared tree and source tree the same path for the Sparse Mode host or have spt-threshold infinity hosts only
Severity: Severe
Status: Assigned
Platform Seen: N7000
Resolved Releases:
Applicable Releases: 6.2(6)

CSCul66808
Symptom: isis_FabricPath cores while doing ISSD from 6.2.5.60.S2 to 6.2.2
Conditions: ISSD was done on N7K from 6.2.5.60_S2 to CCO 6.2.2 image (sup2). N7K has 2 vdc's in FabricPath. isis_FabricPath cored on these vdc's after system switchover was done.
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCul88464
Symptom: ISSU aborts occasionally with timeout error
Conditions: Occasionally while testing ISSU from 5.2.9 - CCO image to 6.2.5.65.S2/6.2.5.60.S2 image, ISSU aborts with timeout error, however on re-issue of ISSU command, it runs smooth and ISSU completes successfully
Workaround: Re-issue the ISSU command “install all kickstart <kickstart_image> system <system_image>”
Severity: Minor
Status: New
Platform Seen: N7000
Resolved Releases:
Applicable Releases: 6.2(6)

CSCul98066
Symptom: Standby SUP fails to come online with correct image during ISSU.
Conditions: ISSU to image 6.2.6.S1 from 5.2.9/6.1.4 fails because standby SUP fails to come online with 6.2.6.S1 after reload, returning error: Install has failed. Return code 0x40930040 (standby supervisor booted up with unexpected version)
Workaround: None
Severity: Severe
Status: Duplicate of CSCul47945
Platform Seen: N7000
Resolved Releases: 6.2(6)
Applicable Releases:

CSCum58738/ CSCul18399
Symptom: On reloading N7K vdc, netstack and syslogd core is seen on the switch
Conditions: N7k has a vdc which is a vpc secondary peer. On doing a 'reload vdc', vdc state is moved to 'suspend in progress', when a netstack and syslogd cores are seen. The vdc remains in failed state.
Workaround: None
Severity: Severe

Status: Verified
Platform Seen: N7000
Resolved Releases: 6.2(6a)
Applicable Releases:

CSCum80838
Symptom: ISSU aborts with failure of supervisor from 6.2.6a to 6.2.2a
Conditions: While performing ISSD from 6.2.6a to 6.2.2a, after the switchover, the standby sup fails to boot up with 6.2.2a. The sup remains in fail state.
Workaround: Reload the switch for standby sup to come up.
Severity: Severe
Status: Unreproducible
Platform Seen: N7000
Resolved Releases:
Applicable Releases: 6.2(6a)

5.2 Caveats for DC31 (Nexus 6000)

CSCub68098
Symptom: "1%KERN-3-SYSTEM_MSG: packet sendmsg: packet size 9250 > MTU 9230" message is seen
Conditions: When the system MTU is set to greater than 9192 this message could be seen as the internal header on the nexus 6000 is 24 bytes and the MTU can hence become grea ter than 9216
Workaround: Set system MTU to less than 9192
Severity: Severe
Status: Closed
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2) 6.0(2)N2(1)

CSCul56319
Symptom: (S,G) states not created on source DR with spt-threshold config
Conditions: When spt-threshold infinity is configured, the source DR does not create (S,G) states. This also results in no PIM registers being sent.
Workaround: None
Severity: Severe
Status: New
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2) 6.0(2)N2(1)

CSCul84598

Symptom: Source DR dropping pim register stop due to "no state"
Conditions: On starting multicast data traffic the source sends PIM register to the RP; but when it receives the register stop the message is discarded with the following message “pim: [4137] (default-base) No state for (131.30.11.12/32, 230.31.0.2/32), message discarded”
Workaround: None
Severity: Severe
Status: New
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2)

CSCum16110
Symptom: OIF on mroute not removed when interface is remotely shut
Conditions: When interface is remotely shut, OIF on mroute is not removed from the OIF list even it's down. This does NOT happen on local shut and for port-channels. This causes a problem as traffic gets forwarded immediately after link up.
Workaround: Configure routed port-channels instead of individual routed links.
Severity: Severe
Status: New
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2)

CSCun06145
Symptom: Incoming PIM Join not processed on link recovery
Conditions: Upon recovery from a link failure the immediate PIM Join may not be processed. A traffic drop of up to 60 seconds could be expected.
Workaround: None
Severity: Severe
Status: New
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2)

CSCun31570
Symptom: BGP next-hop-self not shown for IPv6 address family
Conditions: When BGP next-hop-self is configured under the IPv6 address family, it does not get reflected in the ‘show ip bgp neighbors <neighbor>’ command. However, it can be confirmed that the command works by issuing ‘show ip bgp neighbors <neighbor> received-routes’ command on the peer.
Workaround: None
Severity: Cosmetic
Status: New
Platform Seen: N6000
Resolved Releases:

Applicable Releases: 6.0(2)N2(2)

5.3 Caveats for DC32 (Nexus 3548)

CSCul27903
Symptom: Nexus 3548 PIM prune not sent upon link recovery
Conditions: Upon the recovery (no shut) of the RPF interface the PIM prune message is not sent on the original incoming interface causing temporary multicast packet duplication.
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCuj81917
Symptom: Packets sent to router mac are not reaching router CPU.
Conditions: "peer-gateway" config under vpc config.
Workaround: Do not configure "peer-gateway" under vpc config.
Severity: Severe
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCul88331
Symptom: Nexus 3548 sends a copy of IPv4 unicast packets with IP protocol number 103 to CPU.
Conditions: When Nexus 3548 receives PIM packets, it also copies the IPv4 packets with proto 103 to CPU, consequently creating unwanted duplication.
Workaround: None
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3548
Resolved Releases: 6.0(2)A1(1d)
Applicable Releases: 6.0(2)A1(1c)

CSCum13379
Symptom: Service "snmpd" will crash repeatedly upon MIB walk
Conditions: Intense mibwalk processing may trigger SNMPD to crash on the Nexus 3548.
Workaround: None
Severity: Severe
Status: Unreproducible
Platform Seen: Nexus 3548

Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCum63413
Symptom: Nexus 3548 non RPF multicast traffic might not get dropped
Conditions: Non RPF multicast traffic received over the shared tree is forwarded even in the presence of the related source tree entry causing packet duplication.
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCul56932
Symptom: Nexus 3548 PIM register message sent without creating (S,G)
Conditions: PIM Register will be sent even without creating the related (S,G) causing subsequent PIM Register Stop messages to be dropped by the FHR
Workaround: None
Severity: Moderate
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCuj56903
Symptom: Ipfib crash causes switch to reload when bgp ipv6 address family configurations are present .
Conditions: ipv6 configuration under bgp can cause ipfib crash, need to limit cli.
Workaround: Do not configure ipv6 under bgp.
Severity: Moderate
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCuj95690
Symptom: Inconsistent IPv4 host ECMP configuration in the running-config
Conditions: In running-config, while "hardware profile unicast enable-host-ecmp" shows that ecmp for hosts is enabled, "no hardware profile unicast enable-host-ecmp ipv4" is also mistakenly shown.
Workaround: None
Severity: Moderate
Status: New
Platform Seen: Nexus 3548

Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCul27880
Symptom: Nexus 3548 OIF not removed upon interface failure.
Conditions: Upon the failure of an individual routed interface, the OIF is not removed from the associated mroute entries
Workaround: Configure routed port-channels instead of individual routed links.
Severity: Moderate
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCun31859
Symptom: Clearing mac table causes ip arp table to flush
Conditions: Issuing "clear mac address table dynamic" command causes the ip arp table to flush and entries to be relearned.
Workaround: None
Severity: Severe
Status: New
Platform Seen: N3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

CSCun37474
Symptom: Nexus 3548 show mac address-table age inconsistent with other Nexus platforms
Conditions: In a Nexus 3548, the show mac address-table refers to the age as "seconds since first seen" whereas in the other NxOS platforms the age refers to "seconds since last seen"
Workaround: None
Severity: Moderate
Status: New
Platform Seen: N3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c)

5.4 Caveats for DC33 (Nexus 3000)

CSCuj58599
Symptom: Nexus 3000 License lost upon image change with system reload
Conditions: Sometimes after upgrading the image through system reload the license can get lost
Workaround: None
Severity: Moderate

Status: Unreproducible
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCuj92589
Symptom: Nexus 3000 can crash with an assert (FWM-2-FWM_ASSERT_FAILURE: ../platform/nuova/forwarding-sw/server/fwm_mcec.c.1227 assertion '0' stack)
Conditions: In a vPC system, after recovering the keepalive link and the vPC peer-link on the secondary vPC peer, the primary vPC peer crashes with an assert
Workaround: None
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3000
Resolved Releases: 6.0(2)U2(1)
Applicable Releases: 6.0(2)U1(3)

CSCuj64147
Symptom: Nexus 3000 configuration loss upon image upgrade
Conditions: Changing the image from 6.0(2)U1(3) to 6.0(2)U2(1) might cause the loss of the configuration on the last four interfaces since these are named differently in 6.0(2)U2(1) (ethernet1/49 - 52 become ethernet1/49/1-4).
Workaround: Last 4 interfaces must be reconfigured
Severity: Severe
Status: Assigned
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U2(1)

CSCuj64562
Symptom: Nexus 3000 passes the wrong interface type to its neighbors
Conditions: After system upgrade some interfaces pass the wrong CDP information to their CDP neighbors
Workaround: None
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3000
Resolved Releases: 6.0(2)U2(1)
Applicable Releases: 6.0(2)U1(3)

CSCuj74966/ CSCul30735
Symptom: Nexus 3000 might become unusable
Conditions: A system upgrade with “install all” might lead to a complete unrecoverable switch failure

Workaround: None
Severity: Severe
Status: Closed
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCub70536/CSCui63140/CSCuj89158
Symptom: Nexus 3000 non RPF multicast traffic might not get dropped
Conditions: Non RPF multicast traffic received over the shared tree is forwarded even in the presence of the related source tree entry causing packet duplication.
Workaround: employ "hardware profile multicast prefer-source-tree eternity". However, the usage of this CLI will impact multicast traffic convergence timing.
Severity: Severe
Status: Closed
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCuj67358
Symptom: Nexus 3000 SNMPwalk ifHCOutMulticastPkts counters get cleared upon clear count
Conditions: Clear counter through the CLIs will cause ifHCOutMulticastPkts counter to reset
Workaround: None
Severity: Moderate
Status: Assigned
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCuj67375
Symptom: Nexus 3000 might lose portion of the configuration (it happened only once)
Conditions: Upon switch failure portion of the configuration gets lost
Workaround: System needs a password recovery procedure
Severity: Severe
Status: Assigned
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCuj67375
Symptom: Nexus 3000 some CDP neighbors are missing
Conditions: After switch reload, some of the CDP neighbors are missing from the CDP table indefinitely.
Workaround: None

Severity: Moderate
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul14373
Symptom: Nexus 3000 can experience some temporary multicast packet duplication every minute
Conditions: In a vPC system, when the number of uplink paths of the vPC peers differs from one another, periodic duplication is seen on the receivers connected downstream to the vPC peers even if "ip pim pre-build-spt" option enabled on the vPC peers.
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul08871
Symptom: Nexus 3000 temporary packet duplication upon spine router failure
Conditions: Spine router reload might cause temporary packet duplication even if "hardware profile multicast prefer-source-tree eternity" is configured on the Nexus 3000
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul46510
Symptom: Nexus 3000 "show routing hash" fails to return a value
Conditions: CLI fails to return the proper value
Workaround: None
Severity: Moderate
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul28254
Symptom: Nexus 3000 PIM prune not sent upon link recovery
Conditions: Upon the recovery (no shut) of the RPF interface the PIM prune message is not sent on the original incoming interface causing temporary multicast packet duplication.

Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul28087
Symptom: Nexus 3000 OIF not removed upon interface failure.
Conditions: Upon the failure of an individual routed interface, the OIF is not removed from the associated mroute entries
Workaround: Configure routed port-channels instead of individual routed links.
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul39829
Symptom: Nexus 3000 PIM register is fully polarized
Conditions: The PIM register message is polarized when deployed with Multicast with Multipath
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul45536
Symptom: Nexus 3000 IGMP Static OIF configuration loss (it happened only once)
Conditions: Upon system reload the static IGMP OIF command might get removed from the configuration
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul46458
Symptom: Nexus 3000 “show port-channel load-balance” return wrong information
Conditions: After changing the default value for that CLI, the running configuration will still return the default value

Workaround: None
Severity: Minor
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCum01506
Symptom: Nexus 3000 Static OIF over vPC does not work properly on a RP-on-a-stick topology
Conditions: For vPC, the multicast RP should be equidistant from both vPC peers. If this condition is not met static-oif on vPC peers is not supported
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCul63968
Symptom: Nexus 3000 PIM register message sent without creating (S,G)
Conditions: PIM Register will be sent even without creating the related (S,G) causing subsequent PIM Register Stop messages to be dropped by the FHR
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCun29189
Symptom: Nexus 3000 show run vpc all does not show ip arp sync configuration even if enabled
Conditions: In a vPC system, when enabling ip arp synchronization-entries, the running-config does not reflect the proper configuration on both vPC peers
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3)

CSCun37434
Symptom: Nexus 3000 show mac address-table age inconsistent with other Nexus platforms

Conditions: In a Nexus 3000, the show mac address-table refers to the age as “seconds since first seen” whereas in the other NxOS platforms the age refers to “seconds since last seen”
Workaround: None
Severity: Moderate
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U2)(1)

5.5 Caveats for DC36 (Nexus 3048/3064)

CSCul13663
Symptom: N3k ping is not working on directly connected interface
Conditions: The Nexus 3048 uplink to Nexus 3064 is an L3 directly connected interface. The interface is up and cdp neighbor and pim neighbor are up, but ping does not work
Workaround: Shut/no shut the affected interface will recover the problem
Severity: Severe
Status: Unreproducible
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul28008
Symptom: Applying qos policy to one interface causes others interface autonegotiate to ON
Conditions: Attaching qos policy to any connected port in Auto mode, causes all other Auto ports to go to ON
Workaround: None
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul28467
Symptom: Removing the PFC config from PO doesn’t remove PFC from port interface config
Conditions: Configure "priority-flow-control mode on" under Portchannel and it automatically configures PFC for the interface under that portchannel, then Remove PFC by "no priority-flow-control mode on" from that Portchannel, but it dosen't remove PFC config from the interfa ce under that portchannel
Workaround: Refer to PSS corruption recovery procedure
Severity: Severe
Status: Unreproducible
Platform Seen: N3000
Resolved Releases:

Applicable Releases: 6.0(2)U2(1)

CSCul36464/ CSCul32511
Symptom: BGP peers keep flapping since received checksum CRC error from peers
Conditions: When the switch MTU is configured to 9216, due to port max frame setting incorrectly, it cause BGP received checksum CRC error from peers and continuously flapping
Workaround: None
Severity: Severe
Status: Verified
Platform Seen: N3000
Resolved Releases: 6.0(2)U2(2z)
Applicable Releases:

CSCul38909/CSCtl82866
Symptom: Nexus3000: SPAN not capturing packets sourced by CPU
Conditions: Tx Control packets are not monitored in SPAN session
Workaround: None
Severity: Enhancement
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul41772
Symptom: Applied qos policy to 2 members of PO, only one member interface is autonegotiated to ON
Conditions: The issue happens if the total PFC buffer requirement exceeds default or configured MMU reservation for PFC.
Workaround: Change MMU buffers by the command "hardware profile pfc mmu buffer-reservation xx"
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul42485
Symptom: "debug ip packet detail" output of chksum is truncated
Conditions: Turn on "debug ip packet detail" on N3k, the output of chksum is truncated which is made the chksum incorrect
Workaround: None
Severity: Minor
Status: New

Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul46656
Symptom: "show interface priority-flow-control" should also show portchannel status
Conditions: Currently N3k "show interface priority-flow-control " only show port status, not portchannel status
Workaround: None
Severity: Enhancement
Status: Assigned
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul51491/CSCtz1117
Symptom: IPv6 ECMP HW programming fails on shut/no shut of interface
Conditions: Configure IPv6 BGP transport over IPV4 and ECMP for IPv6 routes, then doing shut/no shut of ECMP/next-hop interface, ECMP/HW multipath programming is not updated for this interface. U6RIB/U6FIB is updated correctly but fails in HW
Workaround: Configure *"ipv6 nd na glean"* on all ECMP interfaces and this will cause the neighbor table to be updated immediately on shut/no shut of interface.
Severity: Moderate
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1) 6.0(2)U1(1)

CSCul69815/ CSCuj68430
Symptom: Seen %URIB-3-RNH_LOOP_ERROR when "clear ip route *"or "clear ipv6 route *" or shut/no shut range link
Conditions: When "clear ip route *"or "clear ipv6 route *" or shut/no shut range link, seen %URIB-3-RNH_LOOP_ERROR, it may cause convergence delay in ::/0 route re-programming in HW
Workaround: None
Severity: Moderate
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1) 6.0(2)U2(2z)

CSCul79204
Symptom: Seen "sd wrap: unknown syslog level:19 " when switch bootup
Conditions: When switch bootup, seen "sd wrap: unknown syslog level:19 "

Workaround: None
Severity: Minor
Status: New
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul81364
Symptom: Seen %NETSTACK-3-TCP_MD5_AUTH_FAILURE: netstack [3360] when restart bgp

Conditions: When "clear ip bgp * " and then follow by "restart bgp xxx" on leaf switches, seen the error message "%NETSTACK-3-TCP_MD5_AUTH_FAILURE:netstack [3360] MD5_DIGEST_MISSING:Dropping packets from src:36.106.41.106.179,dst:36.106.41.4.24363" on the BGP peer spine switch
Workaround: None
Severity: Moderate
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul81414
Symptom: Seen "Received unknown MTS message" when "restart bgp xxx"
Conditions: When "clear ip bgp *" and then follow by "restart bgp xxx", seen the following error on local switch "%BGP-4-MTSUNKOPC: bgp-36103 [8115] Received unknown MTS message on bgp-36103 queue, opc 4852"
Workaround: None
Severity: Moderate
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul87439
Symptom: Switch install wrong IPv6 route entry when IPv6 BGP peer/IPv4 transport
Conditions: Bring up IPv6 peer over IPv4 transports with NH as local ip, both BGP ipv4 and ipv6 routes are learn by the N3k switch correctly, but IPv6 route entries are installed wrong
Workaround: Change IPv6 route updates over IPv4 peering using global v6 prefix as NH
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul95628
Symptom: BGP "shutdown due to no memory condition" when applied v6 route-map
Conditions: Apply the route-map with Next hop "ipv6 next-hop 2001:40:36:250:6::1" input to the neighbor having Ipv4/Ipv6 AFI enabled (at this case is IPv6 transport over ipv4), bgp session shut down due to NoMem on applying rpm with "ipv6 next-hop 2001:40:36:250:6::1"
Workaround: None
Severity: Severe
Status: Resolved
Platform Seen: N3000
Resolved Releases: 6.0(2)U2(2z)
Applicable Releases: 6.0(2)U2(1)

CSCum51358
Symptom: Packets drop due to "input discard" after BGP peer switch reload
Conditions: Reloading BGP peer switch causes all incoming packets to be dropped due to "input discard"
Workaround: Shut/No Shut the interface will resolve the packet drop.
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCum55853
Symptom: N3k MyStation TCAM corrupted after multiple link shut/no shut
Conditions: After performing multiple link shut/no shut on Nexus 3064 switch, the MyStation TCAM gets corrupted. This causes all ARP among directly connected peers to fail.
Workaround: None
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCum69086
Symptom: Error "%USER-3-SYSTEM_MSG: user delete failed for interop:userdel:..."
Conditions: After reloading the switch, the following error gets displayed for each user [interop] not defined locally which was logged in previously: "%USER-3-SYSTEM_MSG: user delete failed for interop:userdel: error removing directory /var/home/interop o such file or directory - securityd"
Workaround: None
Severity: Minor
Status: New
Platform Seen: N3000

Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCun32115
Symptom: Clearing mac table causes ip arp table to flush
Conditions: Issuing "clear mac address table dynamic" command causes the ip arp table to flush and entries to be relearned.
Workaround: None
Severity: Severe
Status: New
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCul46641
Symptom: %IPFIB-2-FIB_TCAM_RESOURCE_EXHAUSTION: FIB TCAM exhausted
Conditions: When URPF enable and If ipv6 route_128 > 128, reload switch may hit FIB TCAM Exhausted issue
Workaround: Disable URPF by the command “system urpf disable” and it changes max ipv6 route_128 to 256
Severity: Severe
Status: Closed
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1)

CSCun32115
Symptom: Clearing mac table causes ip arp table to flush
Conditions: when issuing "clear mac address table dynamic" command in cli causes the ip arp table to flush and entries need to be relearned
Workaround: None
Severity: Severe
Status: New
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(1)

References:

Cisco NX-OS Licensing Guide
Nexus 7000 Install and Upgrade Guides
Nexus 7000 Configuration Guides
Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide
Design Considerations for Classical Ethernet Integration of the Cisco Nexus 7000 M1 and F1 Modules
Cisco FabricPath Best Practices
Cisco FabricPath Design Guide: Using FabricPath with an Aggregation and Access Topology
Data Center Access Design with Cisco Nexus 5000 Series Switches and 2000 Series Fabric Extenders and Virtual PortChannels
Cisco UCS Manager Configuration Common Practices and Quick-Start Guide
Cisco VM-FEX Best Practices for VMware ESX Environment Deployment Guide
Virtual Machine Mobility with VMware VMotion and Cisco Data Center Interconnect Technologies
UCS Command References
UCS Install and Upgrade Guides
UCS Configuration and Firmware Management Guides

6. NVT Test Results

The following section contains test case results for:

- DC31
- DC32
- DC33
- DC36

For DC1 and DC2 results please refer to addendum NVT phase 2.6.

Total # of test cases – Total number of test cases
Total # of Pass – Total number of test cases that meet the passing criteria for the latest test run
Total # of Pass with Exception – Total number of test cases that meet passing criteria with exceptions for the latest test run
Total # of Fail – Total number of test cases that fail to meet the passing criteria for the latest test run
Total # of Iteration – Total number of times a test case has been executed

	Folders	Verification	Total # of test cases	Total # of Pass	Total # of Pass w/Exception	Total # of Fail	Total # of Iteration	Defect(s)
1	NVT 3.0		5436	5006	47	383	20088	
1.1	DC33		1286	991	11	284	4132	
1.1.1	Configuration		32	32	0	0	176	
1.1.1.1	Common Configuration		4	4	0	0	23	
		Verify SSH works through the management network on a dedicated vrf						
		Verify RSA key does not change on device						
		Verify MTU setting (9216)						
		Verify logging server config on switch and that logs in logging server						
		Verify CoPP						
		Verify SNMP and traps						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify licensing						

		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
1.1.1.2	Ixia Setup/Configuration		4	4	0	0	23	
		Physical cabling						
		Upgrade chassis and client software to IxOS/IxNetwork 6.30						
		Configure and verify Static IP w/Auth						
		Check arp resolve/mac address						
		Generate east-west Ucast/Mcast/L2 Traffic						
		Generate north-south Ucast/Mcast Traffic						
1.1.1.3	Interface and LACP Configs		4	4	0	0	23	
		Verify interface and lacp config.						
1.1.1.4	SVI and HSRP Configs		4	4	0	0	19	
		Verify SVI and HSRP						
1.1.1.5	SPT Configs (MST)		4	4	0	0	19	
		Verify root guard, bpdu filter, edge trunk, port fast						
		Verify QinQ for fanout						
1.1.1.6	OSPF Configs		4	4	0	0	23	
		Verify OSPF authentication						
		Verify OSPF neighbor						
1.1.1.7	BGP Configs		4	4	0	0	23	
		Configure and verify BGP to other core						
		Configure and verify eBGP to spine						
		Verify BGP neighbor						
1.1.1.8	Mcast Configs		4	4	0	0	23	
		Configure PIM						
		Configure PIM prebuild						
		Verify PIM neighbor						
		Verify RP placement and advertisement						
		Verify anycast RP with MSDP with mesh-group						
		Verify static IGMP join						

1.1.2	Spine to Core Setup		4	4	0	0	4	
1.1.2.1	Spine to Core Setup		4	4	0	0	4	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						

		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.1.3	Spine to Leaf Setup		4	4	0	0	4	
1.1.3.1	Spine to Leaf N3000 Setup		3	3	0	0	3	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						

		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.1.3.2	Spine to Leaf C6K Setup		1	1	0	0	1	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						

		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify spanning tree status on all vlans.						
		Verify vPC status and consistency parameters.						

		Verify that there are no dead flows						
1.1.4	Leaf to Spine Setup		4	4	0	0	44	
1.1.4.1	Leaf N3000 to N3K Spine		3	3	0	0	42	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						

		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.1.4.2	Leaf C6K to N3K Spine		1	1	0	0	2	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						

		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.1.5	Leaf to Hosts Ixia Setup		4	3	1	0	47	
1.1.5.1	Leaf to Hosts Ixia Setup		4	3	1	0	47	CSCul16104
		Verify spanning tree status (edge) on all vlans for the host ports.						
		Verify mac table is populated correctly.						

		Verify IGMP/MLD snooping.						
		Verify that there are no dead flows						
1.1.6	Leaf to L2 C6K Switch Setup		1	1	0	0	40	
1.1.6.1	Leaf to L2 C6K Switch Setup		1	1	0	0	40	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						

		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify spanning tree status on all vlans.						
		Verify that there are no dead flows						
1.1.7	Unicast ECMP		577	454	0	123	1620	
1.1.7.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers		4	4	0	0	12	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						

		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.2	L3 Port-channel Failure/Recovery between Spines		16	16	0	0	48	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						

		Verify SPAN is mirroring packets correctly.					
		Verify OTV traffic reconverges and optimize OSPF as needed.					
		Verify SNMP traps are sent to SNMP collector.					
		All unicast and multicast traffic should re-converge with proportionate packet loss.					
		Verify traffic destined for CoPP classes is policed as expected.					
		Verify OSPF interface status for the affected links.					
		Verify OSPF neighbor changes and authentication.					
		Verify OSPF DB/Topology consistency.					
		Verify OSPF routes and forwarding table consistency..					
		Verify OSPF multi-path load-balancing.					
		Verify HW and SW entries are properly programmed and synchronized.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify BFD peer detection and client notifications.					
		Verify frames delta does not increase.					
		Verify rx rate for all ixia ports are as expected (compared to baseline).					
		Verify packet loss duration is within expected range.					
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.					
1.1.7.3	L3 Port-channel member Failure/Recovery between Spines		16	16	0	0	49	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.					
		Verify that there are no dead flows					
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify port-channel load balancing and rbh assignment					
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.					

		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.4	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf		214	172	0	42	548	CSCul28254,CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						

		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.5	L3 Routed Port Failure/Recovery		214	169	0	45	566	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul28254,CSCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						

		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.6	L3 Port-channel Failure/Recovery between Spine and Leaf		24	0	0	24	48	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul28254,CSCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						

		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.7	L3 port-channel member Failure/Recovery between Spine and Leaf		68	56	0	12	192	CSCul28254,C SCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						

		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.7.8	Clear Neighbors		7	7	0	0	42	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						

		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.					
		Verify frames delta does not increase.					
		Verify rx rate for all ixia ports are as expected (compared to baseline).					
		Verify packet loss duration is within expected range.					
1.1.7.9	Clear Ipv4/IPv6 Unicast Routes		7	7	0	0	26
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.					
		Verify that there are no dead flows					
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		All unicast and multicast traffic should re-converge.					
		Verify OSPF IPv4/IPv6 neighbors will restart and come back correctly.					
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.					
		Verify that CDP/LLDP does not lose peer information.					
		Verify that no flooding happens after traffic convergence.					
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.					
		Verify SPAN is mirroring packets correctly.					
		Verify SNMP traps are sent to SNMP collector.					
		Verify traffic destined for CoPP classes is policed as expected.					
		Verify OSPF neighbor changes and authentication.					
		Verify OSPF DB/Topology consistency.					
		Verify OSPF routes and forwarding table consistency.					
		Verify OSPF multi-path load-balancing.					
		Verify HW and SW entries are properly programmed and synchronized.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					

		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.1.7.10	Restart process		7	7	0	0	89	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.1.8	L2 Link Failure/Recovery		16	14	0	2	32	
1.1.8.1	vPC leg failure/recovery between Leaf and ToR		4	4	0	0	8	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC leg is shut.						
		Multicast forwarder should not change.						
		Verify that there is no protocol flapping.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify mac move and any missing mac address.						
		Verify mac table is empty after link shut.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						

		Verify traffic drop based on interface counters.						
		Verify that no flooding happens after traffic convergence.						
		Verify STP port states after link disruption are in the expected forwarding mode. Verify that the STP root does not change.						
1.1.8.2	vPC leg member failure/recovery between Leaf and ToR		4	4	0	0	8	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic.						
		The maximum traffic loss for member failure multicast upstream will drop proportionate and for downstream will be either 50% disrupted or no loss depending on which vPC leg member is shut (assuming th						
		Multicast forwarder should not change.						
		Verify that there is no protocol flapping.						
		Verify port-channel load balancing and rbh assignment.						
		Verify that IGMP/MLD membership is not affected.						
1.1.8.3	vPC peer-link failure/recovery between Leaf vPC peer switches		4	2	0	2	8	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		There is no expected effects, both vPC peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.						
		Verify that on recovery, the original states will be re-established.						

1.1.8.4	vPC Peer-keepalive failure/recovery between Leaf vPC peer switches		4	4	0	0	8	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		There is no expected effects, both vPC peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.						
		Verify that on recovery, the original states will be re-established.						
1.1.9	Multicast with Multipath		628	473	3	152	2107	
1.1.9.1	First receiver on first leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.2	First receiver on first leaf - IGMP leave G1 (1)		1	1	0	0	1	

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.3	First receiver on first leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.4	First receiver on first leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.5	First receiver on first leaf - IGMP join G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.6	Second receiver on first leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.7	Second receiver on first leaf - IGMP leave G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.8	Second receiver on first leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.9	Second receiver on first leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.10	Second receiver on first leaf - IGMP join G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.11	All remaining 8 receivers on first leaf - IGMP join G1		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.12	Leave on first leaf - last most recently joined 8 receivers G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.13	First receiver on second leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.14	First receiver on second leaf - IGMP leave G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.15	First receiver on second leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.16	First receiver on second leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.17	First receiver on second leaf - IGMP join G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.18	Second receiver on first leaf - IGMP leave G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.19	Second receiver on first leaf - IGMP join G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.20	Second receiver on first leaf - IGMP silent leave G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					
		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)					
1.1.9.2 1	First receiver on first leaf - IGMP leave G1 (3)		1	1	0	0	1
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					
		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.22	First receiver on first leaf - IGMP join G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.23	First receiver on first leaf - IGMP silent leave G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.24	First receiver on first leaf - IGMP join G1 (5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.25	Second receiver on first leaf - IGMP join G1 (5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.26	Second receiver on second leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.27	All remaining 8 receivers on second leaf - IGMP join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.28	RPF Failure/Recovery between leaf and spine		3	0	0	3	3	CSCul28254 CSCul28087,CSCul28254 CSCul27808
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						

1.1.9.29	Progressive RPF Failure/Recovery between leaf and spine		3	0	0	3	3	CSCul28254 CSCul28087
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.30	Stop all receivers G1		2	2	0	0	2	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						

		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.3 1	Start one source from first leaf for G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.3 2	Stop one source from first leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.33	Start 5 sources from first leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.34	Stop 5 sources from first leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.3 5	Start 5 sources from first leaf on different vlans for G1		1	0	0	1	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.3 6	Stop 5 sources from first leaf on different vlans for G1		1	0	0	1	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						

		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.37	Start one source from first leaf for G1-10		1	0	0	1	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.38	Stop one source from first leaf for G1-10		1	1	0	0	1	

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.39	Start one source from second leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.40	Stop one source from second leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.41	Start 5 sources from second leaf on same vlan for G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.4 2	Stop 5 sources from second leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.4 3	Start 5 sources from second leaf on different vlans for G1		1	0	0	1	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.44	Stop 5 sources from second leaf on different vlans for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.45	Start one source from second leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.4 6	Stop one source from second leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.4 7	Start one source from first leaf for G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.48	Start 5 sources from second leaf on same vlan for G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.49	RPF Failure/Recovery between first leaf and elected RP		2	0	0	2	10	CSCul39647
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.50	RPF Failure/Recovery between second leaf (DR) and elected RP		4	0	0	4	4	CSCul39647
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.5 1	Start all sources		3	0	0	3	3	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.5 2	Start all igmp joins from all hosts		3	0	3	0	3	CSCul14373
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.1.9.53	L3 Port-channel Failure/Recovery between Core and Distribution Layers		4	4	0	0	12	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						

		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.54	L3 Port-channel Failure/Recovery between Spines		16	16	0	0	64	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						

		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.5 5	L3 Port-channel member Failure/Recovery between Spines		16	16	0	0	65	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						

		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.5 6	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf		214	169	0	45	747	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul28254,CSCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.5 7	L3 Routed Port Failure/Recovery		214	169	0	45	747	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul28254,C SCul39647
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						

		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.58	L3 Port-channel Failure/Recovery between Spine and Leaf		24	0	0	24	100	CSCul28254,CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						

		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.59	L3 port-channel member Failure/Recovery between Spine and Leaf		68	56	0	12	250	CSCuj89158 CSCul14373 CSCul28087 CSCul39647 CSCum21940, CSCul28254,CSCul39647

		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.1.9.60	RP,DR Failure		7	0	0	7	49	CSCuj58599 CSCuj64147 CSCuj67375 CSCuj58981 CSCul08871 CSCul45536
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify BGP neighbors status and authentication.						
		Verify BGP table and routing table consistency in accordance to the NEXT-HOP attribute settings.						
		Verify BGP multi-path load-balancing.						
		Verify proper BGP policy routing and filtering based on prefix, AS-PATH, LOCAL_PREFERENCE attributes.						
		Verify the conditional injection of the default route from BGP into the IGP.						
		Verify BGP recursive lookup scenario.						
		Verify BGP reconvergence (control-plane & data-plane).						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.1.10	Software Upgrade and Downgrade		7	0	7	0	7	
1.1.10.1	Software Upgrade and Downgrade		7	0	7	0	7	CSCuj74966 CSCul30735
		Verify if ISSU image compatibility for non-disruptive upgrade/downgrade						

		Verify ISSU-ISSD happens as expected. OSPF graceful restart, PIM triggered Joins should work as expected.						
		Compare startup/running configuration on Active Sup and Standby Sup before and after ISSU-ISSD.						
		Verify STP port states during and after ISSU-ISSD.						
		Verify FHRP peers status during and after ISSU-ISSD.						
		Verify CDP/LLDP status after ISSU-ISSD.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the distribution switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after ISSU-ISSD.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly during and after ISSU-ISSD.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbors status and authentication.						
		Verify BGP table and routing table consistency in accordance to the NEXT-HOP attribute settings.						
		Verify proper BGP policy routing and filtering based on prefix, AS-PATH, LOCAL_PREFERENCE attributes.						
		Verify the conditional injection of the default route from BGP into the IGP.						
		Verify BGP recursive lookup scenario.						
		Verify BGP reconvergence for control-plane.						
		Verify OSPF interface status.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency.						

		Verify HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify PIM neighbor status.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify BFD peer should not flap during and after ISSU-ISSD.						
		No traffic loss is expected.						
		If ISSU is disruptive, verify that all unicast/multicast traffic reconverges.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.1.11	Reload and Power Cycle Switch		8	1	0	7	50	
1.1.11.1	Reload Spine		4	0	0	4	28	CSCuj58599 CSCuj64147 CSCuj67375 CSCuj58981 CSCul08871 CSCul45536
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						

		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.1.11.2	Reload Leaf		4	1	0	3	22	CSCuj58599 CSCuj64147 CSCuj67375 CSCuj58981 CSCul08871 CSCul45536
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						

		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.1.12	Leaf to Hosts Setup		1	1	0	0	1	
1.1.12.1	Leaf to N7K Switch Setup		1	1	0	0	1	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						

		Verify SNMP traps to monitor network events					
		Verify UDLD neighbors and UDLD aggressive mode					
		Verify LACP for link aggregation					
		Verify BFD peering for all possible clients with default protocol timers for the clients					
		Verify SSO/NSF and GR					
		Verify CoPP function					
		Verify CoPP counters					
		Verify hardware rate limiter					
		Verify SPAN ensuring cross-module SPAN.					
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)					
		Verify DHCP IP helper and primary/backup server					
		Verify interfaces in error					
		STP: Verify RSTP parameters and port status.					
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping					
		VACL, PACL: Verify that all the policies are properly programmed in hardware.					
		OSPF: Verify OSPFv2/OSPFv3 peering.					
		PIM: Verify PIM peering.					
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.					
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.					
		QoS: Verify QoS marking.					
		DHCP Relay Agent: Verify DHCP relay functionality.					
		BOOTP Relay Agent: Verify BOOTP relay functionality.					
		Verify spanning tree status on all vlans.					
		Verify that there are no dead flows					
1.2	DC32		2068	1989	14	65	6482	
1.2.1	Configuration		32	32	0	0	32	
1.2.1.1	Common Configuration		4	4	0	0	4	
		Verify SSH works through the management network on a dedicated					

		vrf						
		Verify RSA key does not change on device						
		Verify MTU setting (9216)						
		Verify logging server config on switch and that logs in logging server						
		Verify CoPP						
		Verify SNMP and traps						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify licensing						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
1.2.1.2	Ixia Setup/Configuration		4	4	0	0	4	
		Physical cabling						
		Upgrade chassis and client software to IxOS/IxNetwork 6.30						
		Configure and verify Static IP w/Auth						
		Check arp resolve/mac address						
		Generate east-west Ucast/Mcast/L2 Traffic						
		Generate north-south Ucast/Mcast Traffic						
1.2.1.3	Interface and LACP Configs		4	4	0	0	4	
		Verify interface and lacp config.						
1.2.1.4	SVI and HSRP Configs		4	4	0	0	4	
		Verify SVI and HSRP						
1.2.1.5	SPT Configs (MST)		4	4	0	0	4	
		Verify root guard, bpdu filter, edge trunk, port fast						
		Verify QinQ for fanout						
1.2.1.6	OSPF Configs		4	4	0	0	4	
		Verify OSPF authentication						
		Verify OSPF neighbor						
1.2.1.7	BGP Configs		4	4	0	0	4	
		Configure and verify BGP to other core						
		Configure and verify eBGP to spine						

		Verify BGP neighbor						
1.2.1.8	Mcast Configs		4	4	0	0	4	
		Configure PIM						
		Configure PIM prebuild						
		Verify PIM neighbor						
		Verify RP placement and advertisement						
		Verify anycast RP with MSDP with mesh-group						
		Verify static IGMP join						
1.2.2	Spine to Core Setup		4	4	0	0	4	
1.2.2.1	Spine to Core Setup		4	4	0	0	4	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						

		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.2.3	Spine to Leaf Setup		4	4	0	0	4	
1.2.3.1	Spine to Leaf N3500 Setup		4	4	0	0	4	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						

		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.2.4	Leaf to Spine Setup		5	5	0	0	38	
1.2.4.1	Leaf N3500 to Spine Setup		3	3	0	0	29	CSCuj56903
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						

		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						

		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.2.4.2	Leaf N3000 to Spine Setup		1	1	0	0	3	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						

		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.2.4.3	Leaf N7K to Spine Setup		1	1	0	0	6	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						

		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						

		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.2.5	Leaf to Hosts Ixia Setup		5	5	0	0	38	
1.2.5.1	Leaf to Hosts Ixia Setup		5	5	0	0	38	
		Verify spanning tree status (edge) on all vlans for the host ports.						
		Verify mac table is populated correctly.						
		Verify IGMP/MLD snooping.						
		Verify that there are no dead flows						
1.2.6	Leaf to Hosts Setup		1	1	0	0	2	
1.2.6.1	Leaf to N7K Switch Setup		1	1	0	0	2	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						

		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify spanning tree status on all vlans.						
		Verify that there are no dead flows						
1.2.7	Unicast ECMP		943	909	0	34	2927	
1.2.7.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers		16	16	0	0	48	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.2	L3 Port-channel Failure/Recovery between Spines		16	16	0	0	48	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						

		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.3	L3 Port-channel member Failure/Recovery between Spines		16	16	0	0	16	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						

		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.4	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf		256	239	0	17	986	CSCul27903,CSCum13379
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch						

		back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.5	L3 Routed Port Failure/Recovery		260	243	0	17	1014	CSCul27903,CSCum13379
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						

		Verify frames delta does not increase.					
		Verify rx rate for all ixia ports are as expected (compared to baseline).					
		Verify packet loss duration is within expected range.					
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.					
1.2.7.6	L3 Port-channel Failure/Recovery between Spine and Leaf		76	76	0	0	377	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.					
		Verify that there are no dead flows					
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.					
		Verify the L2 forwarding table should remove entries of the affected link.					
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.					
		Verify SPAN is mirroring packets correctly.					
		Verify OTV traffic reconverges and optimize OSPF as needed.					
		Verify SNMP traps are sent to SNMP collector.					
		All unicast and multicast traffic should re-converge with proportionate packet loss.					
		Verify traffic destined for CoPP classes is policed as expected.					
		Verify OSPF interface status for the affected links.					
		Verify OSPF neighbor changes and authentication.					
		Verify OSPF DB/Topology consistency.					
		Verify OSPF routes and forwarding table consistency..					
		Verify OSPF multi-path load-balancing.					
		Verify HW and SW entries are properly programmed and synchronized.					
		Verify multicast HW and SW entries are properly programmed and					

		synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.7	L3 port-channel member Failure/Recovery between Spine and Leaf		176	176	0	0	226	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						

1.2.7.8	L3 Port-channel Subinterface Failure/Recovery between Spine and Leaf		100	100	0	0	101	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.7.9	Clear Neighbors		9	9	0	0	41	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						

		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.7.10	Clear Ipv4/IPv6 Unicast Routes		9	9	0	0	41	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify OSPF IPv4/IPv6 neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency.						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.7.11	Restart process		9	9	0	0	29	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						

		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.8	L2 Link Failure/Recovery		36	32	0	4	49	

1.2.8.1	L2 Port-channel Failure/Recovery between Leaf and ToR devices		32	32	0	0	32	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify FHRP peers status does not change. Verify FHRP MAC in ARP/ND table. Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link at the access switch and re-learnt correctly on the alternative link.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		Verify that the L2 forwarding entries on all switches for nodes connected to the access layer are associated with the corresponding STP forwarding ports.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected link for non-vpc setup.and re-learnt correctly on the alternative link after query from the IGMP snooping router.						
		Verify that IGMP/MLD membership is not affected on the routers.						
		Verify ACL TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify isolated vlans remain to have complete separation from other ports within the same PVLAN but not from the promiscuous ports using proxy-arp.						
		DHCP relay configured on the aggregation switches should remain unaffected.						

		Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, FHRP services, ARP, proxy arp and IGMP.						
		Verify that IPv6 global HSRP is functional.						
		Verify that packets only traverse the fabric for known unicast/multicast destinations and flood through the fabric for unknown unicast, multicast when IGMP snooping is disabled, and broadcast.						
		All unicast and multicast traffic should re-converge with minimal packet loss.						
		Verify SNMP traps are sent to SNMP collector						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify mac move and any missing mac address.						
		Verify mac table is empty after link shut.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
		Verify traffic drop based on interface counters.						
		Verify that no flooding happens after traffic convergence.						
		Verify STP port states after link disruption are in the expected forwarding mode. Verify that the STP root does not change.						
1.2.8.2	L2 Port-channel Failure/Recovery between Leaf devices		4	0	0	4	17	CSCun31859
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify FHRP peers status does not change. Verify FHRP MAC in ARP/ND table. Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						

		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.					
		Verify the L2 forwarding table should remove entries of the affected link at the access switch and re-learnt correctly on the alternative link.					
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.					
		Verify that the L2 forwarding entries on all switches for nodes connected to the access layer are associated with the corresponding STP forwarding ports.					
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.					
		Verify IGMP/MLD snooping entries are deleted for the affected link for non-vpc setup.and re-learnt correctly on the alternative link after query from the IGMP snooping router.					
		Verify that IGMP/MLD membership is not affected on the routers.					
		Verify ACL TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.					
		Verify SPAN is mirroring packets correctly.					
		Verify isolated vlans remain to have complete separation from other ports within the same PVLAN but not from the promiscuous ports using proxy-arp.					
		DHCP relay configured on the aggregation switches should remain unaffected.					
		Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, FHRP services, ARP, proxy arp and IGMP.					
		Verify that IPv6 global HSRP is functional.					
		Verify that packets only traverse the fabric for known unicast/multicast destinations and flood through the fabric for unknown unicast, multicast when IGMP snooping is disabled, and broadcast.					
		All unicast and multicast traffic should re-converge with minimal packet loss.					
		Verify SNMP traps are sent to SNMP collector					
		Verify traffic destined for CoPP classes is policed as expected.					
		Verify frames delta does not increase.					
		Verify rx rate for all ixia ports are as expected (compared to baseline).					
		Verify packet loss duration is within expected range.					

		Verify mac move and any missing mac address.						
		Verify mac table is empty after link shut.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
		Verify traffic drop based on interface counters.						
		Verify that no flooding happens after traffic convergence.						
		Verify STP port states after link disruption are in the expected forwarding mode. Verify that the STP root does not change.						
1.2.9	Multicast with Multipath		1030	997	6	27	3370	
1.2.9.1	A01 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.2	A01 First receiver IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.3	A02 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.4	A02 First receiver Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.5	A03 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.6	A04 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.7	A04 Second receiver IGMP leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.8	A05 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.9	A05 Second receiver Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.10	A06 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.11	A07 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.12	A07 Eight receivers IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.13	A08 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.14	A08 Eight receivers IGMP Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.15	A09 RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul27880
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.16	A10 Progressive RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul27880
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.17	A11 Stop all receivers G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.18	B01 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.19	B01 First receiver IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.20	B02 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.2 1	B02 First receiver Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.2 2	B03 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.23	B04 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.24	B04 Second receiver IGMP leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.25	B05 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.26	B05 Second receiver Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.27	B06 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.28	B07 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.29	B07 Eight receivers IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.30	B08 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
1.2.9.3 1	B08 Eight receivers IGMP Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					
		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
1.2.9.3 2	B09 RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul27880,CSCul27903
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.33	B10 Progressive RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul27880,CSCul27903
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.34	B11 Stop all receivers G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.35	C01 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.36	C01 First receiver IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.37	C02 First receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.38	C02 First receiver Silent Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
1.2.9.39	C03 First receiver IGMP Join G1		1	1	0	0	1
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					
		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
1.2.9.40	C04 Second receiver IGMP Join G1		1	1	0	0	1
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.41	C04 Second receiver IGMP leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.42	C05 Second receiver IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.43	C05 Second receiver Silent Leave G1		1	1	0	0	1	
1.2.9.44	C06 Second receiver IGMP Join G1		1	1	0	0	1	
1.2.9.45	C07 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.46	C07 Eight receivers IGMP Leave G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						

		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.47	C08 Eight receivers IGMP Join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.48	C08 Eight receivers IGMP Silent Leave G1		1	1	0	0	1	
1.2.9.49	C09 RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul28087, CSCul28254
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.50	C10 Progressive RPF Failure/Recovery between leaf and spine		1	0	0	1	1	CSCul28087, CSCul28254
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.51	C11 Stop all receivers G1		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.52	D01 Start traffic (S1,G1)		1	0	0	1	1	CSCul88331
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.53	D01 Stop traffic (S1,G1)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.54	D02 Start traffic (S1-5,G1) same vlan		1	0	1	0	1	CSCul56932
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.55	D02 Stop traffic (S1-5,G1)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.56	D03 Start traffic (S1-5,G1) diff vlan		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.57	D03 Stop traffic (S1-5,G1)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.58	D04 Start traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.59	D04 Stop traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.60	D05 Start traffic (S1-5,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.61	D06 RPF Failure/Recovery between leaf and spine		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.62	D07 Progressive RPF Failure/Recovery between leaf and spine		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.63	D08 Stop all Sources		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.64	E01 Start traffic (S1,G1)		1	1	0	0	1	
1.2.9.65	E01 Stop traffic (S1,G1)		1	1	0	0	1	
1.2.9.66	E02 Start traffic (S1-5,G1) same vlan		1	0	1	0	1	CSCul56932
1.2.9.67	E02 Stop traffic (S1-5,G1)		1	1	0	0	1	
1.2.9.68	E03 Start traffic (S1-5,G1) diff vlan		1	1	0	0	1	
1.2.9.69	E03 Stop traffic (S1-5,G1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and						

		synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.70	E04 Start traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.71	E04 Stop traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.7 2	E05 Start traffic (S1-5,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.7 3	E06 RPF Failure/Recovery between leaf and spine		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.74	E07 Progressive RPF Failure/Recovery between leaf and spine		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.75	E08 Stop all Sources		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.76	F01 Start traffic (S1,G1)		1	0	1	0	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.77	F01 Stop traffic (S1,G1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.78	F02 Start traffic (S1-5,G1) same vlan		1	0	1	0	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.79	F02 Stop traffic (S1-5,G1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.80	F03 Start traffic (S1-5,G1) diff vlan		1	0	1	0	1	CSCul39829
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.81	F03 Stop traffic (S1-5,G1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.8 2	F04 Start traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.8 3	F04 Stop traffic (S1,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.84	F05 Start traffic (S1-5,G1-5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.85	F06 RPF Failure/Recovery between leaf and spine		1	0	1	0	1	CSCul39647
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.86	F07 Progressive RPF Failure/Recovery between leaf and spine		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.87	F08 Stop all Sources		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						

		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
1.2.9.8 8	L3 Port-channel Failure/Recovery between Core and Distribution Layers		4	4	0	0	12	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.89	L3 Port-channel Failure/Recovery between Spines		16	16	0	0	48	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						

		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.90	L3 Port-channel member Failure/Recovery between Spines		16	16	0	0	16	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						

		Verify interfaces in error							
		Verify any core dumps							
		Verify port-channel load balancing and rbh assignment							
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.							
		Verify LACP rebundle for port-channel after member recover.							
		The traffic should be able to re-converge within acceptable time.							
		Verify the convergence pattern is as expected.							
		Verify the route tables for both unicast and multicast are updated correctly.							
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.							
		Verify frames delta does not increase.							
		Verify rx rate for all ixia ports are as expected (compared to baseline).							
		Verify packet loss duration is within expected range.							
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.							
1.2.9.91	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf		256	244	0	12	1096	CSCul27903,CSCum63413, CSCul27880,CSCum13379
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.							
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.							
		Verify that there are no dead flows							
		Verify TB, error, crash							
		Verify interfaces in error							
		Verify any core dumps							
		Verify traffic is load balance to other ECMP paths							
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.							
		Verify LACP rebundle for port-channel after member recover.							

		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.9 2	L3 Routed Port Failure/Recovery		260	252	0	8	1115	CSCum63413, CSCul27880
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						

		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.9 3	L3 Port-channel Failure/Recovery between Spine and Leaf		76	76	0	0	512	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.94	L3 port-channel member Failure/Recovery between Spine and Leaf		176	176	0	0	226	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						

		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.9 5	L3 Port-channel Subinterface Failure/Recovery between Spine and Leaf		100	100	0	0	101	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.2.9.9 6	RP,DR Failure		8	8	0	0	19	
1.2.9.9 7	Clear Ipv4 Multicast Routes		9	9	0	0	37	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						

		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All multicast traffic should re-converge.						
		Verify periodic PIM joins are received and sent upstream after clearing.						
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.9.98	Clear PIM Routes		9	9	0	0	50	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						

		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All multicast traffic should re-converge.						
		Verify periodic PIM joins are received and sent upstream after clearing.						
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.9.99	Clear IGMP Routes/Groups		8	8	0	0	36	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						

		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All multicast traffic should re-converge.						
		Verify periodic PIM joins are received and sent upstream after clearing.						
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.2.9.100	Restart process		5	5	0	0	15	
1.2.10	Reload and Power Cycle Switch		8	0	8	0	18	
1.2.10.1	Reload Spine		4	0	4	0	8	CSCum69086, CSCum13379

		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						

		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.2.10.2	Reload Leaf		4	0	4	0	10	CSCum69086
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						

		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.3	DC31		507	465	10	32	2582	
1.3.1	Configuration		64	64	0	0	624	
1.3.1.1	Common Configuration		8	8	0	0	78	CSCub68098
		Verify SSH works through the management network on a dedicated vrf						
		Verify RSA key does not change on device						
		Verify MTU setting (9216)						
		Verify logging server config on switch and that logs in logging server						
		Verify CoPP						
		Verify SNMP and traps						

		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify licensing						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify UDLD neighbors and UDLD aggressive mode						
1.3.1.2	Ixia Setup/Configuration		8	8	0	0	78	
		Physical cabling						
		Upgrade chassis and client software to IxOS/IxNetwork 6.30						
		Configure and verify Static IP w/Auth						
		Check arp resolve/mac address						
		Generate east-west Ucast/Mcast/L2 Traffic						
		Generate north-south Ucast/Mcast Traffic						
1.3.1.3	Interface and LACP Configs		8	8	0	0	78	
		Verify interface and lacp config.						
1.3.1.4	SVI and HSRP Configs		8	8	0	0	78	CSCub68098
		Verify SVI and HSRP						
1.3.1.5	SPT Configs (MST)		8	8	0	0	78	
		Verify root guard, bpdu filter, edge trunk, port fast						
		Verify QinQ for fanout						
1.3.1.6	OSPF Configs		8	8	0	0	78	
		Verify OSPF authentication						
		Verify OSPF neighbor						
1.3.1.7	BGP Configs		8	8	0	0	78	CSCun31570
		Configure and verify BGP to other core						
		Configure and verify eBGP to spine						
		Verify BGP neighbor						
1.3.1.8	Mcast Configs		8	8	0	0	78	CSCul56319
		Configure PIM						
		Configure PIM prebuild						
		Verify PIM neighbor						

		Verify RP placement and advertisement						
		Verify anycast RP with MSDP with mesh-group						
		Verify static IGMP join						
1.3.2	Spine to Core Setup		2	2	0	0	46	
1.3.2.1	Spine to Core Setup		2	2	0	0	46	CSCub68098
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						

		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.3.3	Spine to Leaf Setup		2	2	0	0	38	
1.3.3.1	Spine to Leaf Setup		2	2	0	0	38	CSCub68098
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						

		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.3.4	Leaf to Spine Setup		6	6	0	0	22	
1.3.4.1	Leaf N6000 to N6K Spine		3	3	0	0	13	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						

		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.3.4.2	Leaf N3500 to N6K Spine		1	1	0	0	4	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						

		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						

		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.3.4.3	Leaf N3000 to N6K Spine		1	1	0	0	4	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						

		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.3.4.4	Leaf N7000 to N6K Spine		1	1	0	0	1	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						

		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						

		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.3.5	Leaf to Hosts Ixia Setup		6	6	0	0	22	
1.3.5.1	Leaf to Hosts Ixia Setup		6	6	0	0	22	
		Verify spanning tree status (edge) on all vlans for the host ports.						
		Verify mac table is populated correctly.						
		Verify IGMP/MLD snooping.						
		Verify that there are no dead flows						
1.3.6	Leaf to Hosts Setup		1	1	0	0	1	
1.3.6.1	Leaf to N7K Switch Setup		1	1	0	0	1	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						

		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify spanning tree status on all vlans.						
		Verify that there are no dead flows						
1.3.7	Software Upgrade and Downgrade		5	5	0	0	5	
1.3.7.1	Software Upgrade and Downgrade		5	5	0	0	5	
		Verify if ISSU image compatibility for non-disruptive upgrade/downgrade						
		Verify ISSU-ISSD happens as expected. OSPF graceful restart, PIM triggered Joins should work as expected.						
		Compare startup/running configuration on Active Sup and Standby Sup before and after ISSU-ISSD.						

		Verify STP port states during and after ISSU-ISSD.						
		Verify FHRP peers status during and after ISSU-ISSD.						
		Verify CDP/LLDP status after ISSU-ISSD.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the distribution switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after ISSU-ISSD.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly during and after ISSU-ISSD.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbors status and authentication.						
		Verify BGP table and routing table consistency in accordance to the NEXT-HOP attribute settings.						
		Verify proper BGP policy routing and filtering based on prefix, AS-PATH, LOCAL_PREFERENCE attributes.						
		Verify the conditional injection of the default route from BGP into the IGP.						
		Verify BGP recursive lookup scenario.						
		Verify BGP reconvergence for control-plane.						
		Verify OSPF interface status.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency.						
		Verify HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify PIM neighbor status.						
		Verify static RP mapping as the backup of auto RP.						

		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify BFD peer should not flap during and after ISSU-ISSD.						
		No traffic loss is expected.						
		If ISSU is disruptive, verify that all unicast/multicast traffic reconverges.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.3.8	Reload and Power Cycle Switch		7	2	5	0	21	
1.3.8.1	Reload Spine		2	0	2	0	6	CSCum69086
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						

		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						

		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.3.8.2	Reload Leaf		5	2	3	0	15	CSCum69086
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						

		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						

		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.3.9	Multicast with Multipath		414	377	5	32	1803	
1.3.9.1	First receiver on first leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.2	First receiver on first leaf - IGMP leave G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.3	First receiver on first leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.4	First receiver on first leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)					
1.3.9.5	First receiver on first leaf - IGMP join G1 (3)		1	1	0	0	1
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					
		Verify PIM neighbor status.					
		Verify PIM both multipath and non-multipath functionalities.					
		Verify AutoRP mapping and boundaries.					
		Verify static RP mapping as the backup of auto RP.					
		Verify MSDP neighbors and SA cache consistency.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify IGMP Snooping table					
		Verify IGMP table					
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)					
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)					
1.3.9.6	Second receiver on first leaf - IGMP join G1 (1)		1	1	0	0	1
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify CDP neighbors					

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.7	Second receiver on first leaf - IGMP leave G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.8	Second receiver on first leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.9	Second receiver on first leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.10	Second receiver on first leaf - IGMP join G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.11	All remaining 8 receivers on first leaf - IGMP join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.1 2	Leave on first leaf - last most recently joined 8 receivers G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.1 3	First receiver on second leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.14	First receiver on second leaf - IGMP leave G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.15	First receiver on second leaf - IGMP join G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.16	First receiver on second leaf - IGMP silent leave G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.17	First receiver on second leaf - IGMP join G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.18	Second receiver on first leaf - IGMP leave G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.19	Second receiver on first leaf - IGMP join G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.20	Second receiver on first leaf - IGMP silent leave G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.21	First receiver on first leaf - IGMP leave G1 (3)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.22	First receiver on first leaf - IGMP join G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.23	First receiver on first leaf - IGMP silent leave G1 (4)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.24	First receiver on first leaf - IGMP join G1 (5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.25	Second receiver on first leaf - IGMP join G1 (5)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.26	Second receiver on second leaf - IGMP join G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.27	All remaining 8 receivers on second leaf - IGMP join G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.28	RPF Failure/Recovery between leaf and spine		3	3	0	0	3	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.29	Progressive RPF Failure/Recovery between leaf and spine		3	3	0	0	3	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.30	Stop all receivers G1		2	2	0	0	2	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.31	Start one source from first leaf for G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.32	Stop one source from first leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.33	Start 5 sources from first leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.3 4	Stop 5 sources from first leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.3 5	Start 5 sources from first leaf on different vlans for G1		1	1	0	0	1	
		Verify TB, error, crash						

		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.36	Stop 5 sources from first leaf on different vlans for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.37	Start one source from first leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.38	Stop one source from first leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.39	Start one source from second leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.40	Stop one source from second leaf for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.4 1	Start 5 sources from second leaf on same vlan for G1 (1)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.4 2	Stop 5 sources from second leaf on same vlan for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.43	Start 5 sources from second leaf on different vlans for G1		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.44	Stop 5 sources from second leaf on different vlans for G1		1	1	0	0	1	

		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.45	Start one source from second leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						

		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.46	Stop one source from second leaf for G1-10		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.47	Start one source from first leaf for G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						

		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.48	Start 5 sources from second leaf on same vlan for G1 (2)		1	1	0	0	1	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.49	RPF Failure/Recovery between first leaf and elected RP		2	1	0	1	12	CSCun06145
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						

		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.50	RPF Failure/Recovery between second leaf (DR) and elected RP		4	4	0	0	4	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.51	Start all sources		7	7	0	0	7	CSCul84598
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.52	Start all igmp joins from all hosts		7	7	0	0	7	
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify IGMP Snooping table						
		Verify IGMP table						
		Verify PIM paramenters on both DR(s) and RPs ((*,G)/(S,G), iif, oif, flags, RPF interface and neighbor)						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
1.3.9.53	L3 Routed Port Failure/Recovery		164	133	0	31	965	CSCun06145,CSCum16110

		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.3.9.5 4	L3 port-channel member Failure/Recovery between Spine and Leaf		120	120	0	0	568	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						

		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.3.9.5 5	L3 Port-channel member Failure/Recovery between Spines		8	8	0	0	40	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						

		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.3.9.56	RP,DR Failure		7	2	5	0	21	CSCum69086
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify BGP neighbors status and authentication.						
		Verify BGP table and routing table consistency in accordance to the NEXT-HOP attribute settings.						
		Verify BGP multi-path load-balancing.						
		Verify proper BGP policy routing and filtering based on prefix, AS-PATH, LOCAL_PREFERENCE attributes.						
		Verify the conditional injection of the default route from BGP into the IGP.						
		Verify BGP recursive lookup scenario.						
		Verify BGP reconvergence (control-plane & data-plane).						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						

		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.3.9.57	Clear Ipv4 Multicast Routes		7	7	0	0	21	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All multicast traffic should re-converge.						
		Verify periodic PIM joins are received and sent upstream after clearing.						
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						

		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.							
		Verify SPAN is mirroring packets correctly.							
		Verify SNMP traps are sent to SNMP collector.							
		Verify traffic destined for CoPP classes is policed as expected.							
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.							
		Verify frames delta does not increase.							
		Verify rx rate for all ixia ports are as expected (compared to baseline).							
		Verify packet loss duration is within expected range.							
1.3.9.58	Clear PIM Routes		7	7	0	0	21		
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.							
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.							
		Verify that there are no dead flows							
		Verify TB, error, crash							
		Verify interfaces in error							
		Verify any core dumps							
		All multicast traffic should re-converge.							
		Verify periodic PIM joins are received and sent upstream after clearing.							
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps							
		Verify that CDP/LLDP does not lose peer information.							
		Verify that no flooding happens after traffic convergence.							
		Verify PIM neighbor status.							
		Verify PIM both multipath and non-multipath functionalities.							
		Verify AutoRP mapping.							
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.							
		Verify PIM source register and register stop.							

		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.3.9.5 9	Clear IGMP Routes/Groups		14	14	0	0	42	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All multicast traffic should re-converge.						
		Verify periodic PIM joins are received and sent upstream after clearing.						
		Verify that the multicast hardware entries are properly removed and re-installed during the mroute flaps						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						

		Verify IGMP/MLD snooping entries are deleted and re-learnt correctly after query from the IGMP snooping router.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.3.9.60	Restart process		14	14	0	0	42	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						

		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.4	DC36		1575	1561	12	2	6892	
1.4.1	Configuration		102	102	0	0	481	
1.4.1.1	Common Configuration		12	12	0	0	57	
		Verify SSH works through the management network on a dedicated vrf						
		Verify RSA key does not change on device						
		Verify MTU setting (9216)						
		Verify logging server config on switch and that logs in logging server						
		Verify CoPP						
		Verify SNMP and traps						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify licensing						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify UDLD neighbors and UDLD aggressive mode						
1.4.1.2	Ixia Setup/Configuration		12	12	0	0	57	
		Physical cabling						
		Upgrade chassis and client software to IxOS/IxNetwork 6.30						
		Configure and verify Static IP w/Auth						
		Check arp resolve/mac address						

		Generate east-west Ucast/Mcast/L2 Traffic						
		Generate north-south Ucast/Mcast Traffic						
1.4.1.3	Interface and LACP Configs		12	12	0	0	57	
		Verify interface and lacp config.						
1.4.1.4	SVI and HSRP Configs		12	12	0	0	57	
		Verify SVI and HSRP						
1.4.1.5	SPT Configs (MST)		12	12	0	0	57	
		Verify root guard, bpdu filter, edge trunk, port fast						
		Verify QinQ for fanout						
1.4.1.6	OSPF Configs		6	6	0	0	25	CSCul38909
		Verify OSPF authentication						
		Verify OSPF neighbor						
1.4.1.7	BGP Configs		12	12	0	0	57	CSCl36464 CSCul38909 CSCul42485 CSCul95628 CSCul87439,C SCl36464
		Configure and verify BGP to other core						
		Configure and verify eBGP to spine						
		Verify BGP neighbor						
1.4.1.8	Mcast Configs		12	12	0	0	57	
		Configure PIM						
		Configure PIM prebuild						
		Verify PIM neighbor						
		Verify RP placement and advertisement						
		Verify anycast RP with MSDP with mesh-group						
		Verify static IGMP join						
1.4.1.9	BFD		12	12	0	0	57	
		Verify BFD peering for all possible clients with default protocol timers for the clients						
1.4.2	Spine to Core Setup		6	6	0	0	25	

1.4.2.1	Spine to Core Setup		6	6	0	0	25	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		OSPF: Verify OSPFv2/OSPFv3 peering.						

		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.4.3	Spine to Leaf Setup		6	6	0	0	25	
1.4.3.1	Spine to Leaf Setup		6	6	0	0	25	CSCul28008 CSCul28467 CSCul41772
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						

		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		MSDP: Verify MSDP peering and SA-cache						
		Verify that there are no dead flows						
1.4.4	Leaf to Spine Setup		6	5	1	0	19	
1.4.4.1	Leaf N3048 to N3k Spine		3	2	1	0	13	CSCul13663
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						

		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.4.4.2	Leaf Cat6k to N3k Spine		1	1	0	0	2	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						

		Verify Syslog to syslog.interop.cisco.com							
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10							
		Verify DNS search list: interop.cisco.com, cisco.com							
		Verify CMP port connections to the management network.							
		Verify CDP neighbors							
		Verify SNMP agent (read community): public + interop; (private community): private + cisco							
		Verify SNMP traps to monitor network events							
		Verify UDLD neighbors and UDLD aggressive mode							
		Verify LACP for link aggregation							
		Verify BFD peering for all possible clients with default protocol timers for the clients							
		Verify SSO/NSF and GR							
		Verify CoPP function							
		Verify CoPP counters							
		Verify hardware rate limiter							
		Verify SPAN ensuring cross-module SPAN.							
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)							
		Verify DHCP IP helper and primary/backup server							
		Verify interfaces in error							
		STP: Verify RSTP parameters and port status.							
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping							
		VACL, PACL: Verify that all the policies are properly programmed in hardware.							
		OSPF: Verify OSPFv2/OSPFv3 peering.							
		PIM: Verify PIM peering.							
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.							
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.							
		QoS: Verify QoS marking.							
		DHCP Relay Agent: Verify DHCP relay functionality.							

		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.4.4.3	Leaf N3064 to N3k Spine		2	2	0	0	4	CSCul28008 CSCul28467 CSCul41772,CSCul28008 CSCul28467 CSCul41772
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						
		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						

		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						
		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify vPC status and consistency parameters.						
		Verify that there are no dead flows						
1.4.5	Leaf to Hosts Ixia Setup		6	6	0	0	19	
1.4.5.1	Leaf to Hosts Ixia Setup		6	6	0	0	19	CSCun32115
		Verify spanning tree status (edge) on all vlans for the host ports.						
		Verify mac table is populated correctly.						
		Verify IGMP/MLD snooping.						
		Verify that there are no dead flows						
1.4.6	Leaf to Hosts Setup		2	2	0	0	2	
1.4.6.1	Leaf to N7K Switch Setup		1	1	0	0	1	
		Verify SSH works through the management network on a dedicated vrf						
		Verify startup and running config						
		Verify TB, error, crash						

		Verify any core dumps						
		Verify RSA key does not change on device						
		Verify ssh on device is functional						
		Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers						
		Verify NTP/PTP and Time Zone : ntp.interop.cisco.com						
		Verify Syslog to syslog.interop.cisco.com						
		Verify DNS domain : interop.cisco.com and server : 172.28.92.9-10						
		Verify DNS search list: interop.cisco.com, cisco.com						
		Verify CMP port connections to the management network.						
		Verify CDP neighbors						
		Verify SNMP agent (read community): public + interop; (private community): private + cisco						
		Verify SNMP traps to monitor network events						
		Verify UDLD neighbors and UDLD aggressive mode						
		Verify LACP for link aggregation						
		Verify BFD peering for all possible clients with default protocol timers for the clients						
		Verify SSO/NSF and GR						
		Verify CoPP function						
		Verify CoPP counters						
		Verify hardware rate limiter						
		Verify SPAN ensuring cross-module SPAN.						
		Configure Authentication for: OSPF/OSPFv3, HSRP/HSRPv6, MSDP, Layer 2 ISIS (FabricPath, OTV)						
		Verify DHCP IP helper and primary/backup server						
		Verify interfaces in error						
		STP: Verify RSTP parameters and port status.						
		IGMP/MLD Snooping: Verify IGMP/MLD Snooping						
		VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		OSPF: Verify OSPFv2/OSPFv3 peering.						
		PIM: Verify PIM peering.						

		ARP & MAC / ND: Verify ARP and MAC addresses are properly learnt across all the forwarding engines.						
		ACL, VACL, PACL: Verify that all the policies are properly programmed in hardware.						
		QoS: Verify QoS marking.						
		DHCP Relay Agent: Verify DHCP relay functionality.						
		BOOTP Relay Agent: Verify BOOTP relay functionality.						
		Verify spanning tree status on all vlans.						
		Verify that there are no dead flows						
1.4.6.2	Leaf to Cat6k Switch Setup		1	1	0	0	1	
1.4.7	Software Upgrade and Downgrade		11	11	0	0	11	
1.4.7.1	Software Upgrade and Downgrade		11	11	0	0	11	
		Verify if ISSU image compatibility for non-disruptive upgrade/downgrade						
		Verify ISSU-ISSD happens as expected. OSPF graceful restart, PIM triggered Joins should work as expected.						
		Compare startup/running configuration on Active Sup and Standby Sup before and after ISSU-ISSD.						
		Verify STP port states during and after ISSU-ISSD.						
		Verify FHRP peers status during and after ISSU-ISSD.						
		Verify CDP/LLDP status after ISSU-ISSD.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the distribution switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after ISSU-ISSD.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly during and after ISSU-ISSD.						
		Verify SNMP traps are sent to SNMP collector.						

		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbors status and authentication.						
		Verify BGP table and routing table consistency in accordance to the NEXT-HOP attribute settings.						
		Verify proper BGP policy routing and filtering based on prefix, AS-PATH, LOCAL_PREFERENCE attributes.						
		Verify the conditional injection of the default route from BGP into the IGP.						
		Verify BGP recursive lookup scenario.						
		Verify BGP reconvergence for control-plane.						
		Verify OSPF interface status.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency.						
		Verify HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify PIM neighbor status.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized after ISSU-ISSD.						
		Verify BFD peer should not flap during and after ISSU-ISSD.						
		No traffic loss is expected.						
		If ISSU is disruptive, verify that all unicast/multicast traffic reconverges.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.4.8	Reload and Power Cycle Switch		11	1	9	1	33	
1.4.8.1	Reload Spine		6	1	5	0	18	CSCum69086, CSCul79204
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						

		Verify that there are no dead flows							
		Verify TB, error, crash							
		Verify interfaces in error							
		Verify any core dumps							
		Verify STP port states during and after reload.							
		Verify FHRP peers status during and after reload.							
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.							
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.							
		Verify FHRP MAC in ARP/ND table.							
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.							
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.							
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.							
		Verify that no flooding happens after traffic convergence.							
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.							
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.							
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.							
		Verify SPAN is mirroring packets correctly.							
		Verify SNMP traps are sent to SNMP collector.							
		All unicast and multicast traffic should re-converge.							
		Verify traffic destined for CoPP classes is policed as expected.							
		Verify OSPF interface status for the affected links.							
		Verify OSPF neighbor changes and authentication.							
		Verify OSPF DB/Topology consistency.							
		Verify OSPF routes and forwarding table consistency..							

		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						
		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.4.8.2	Reload Leaf		5	0	4	1	15	CSCum69086, CSCum51358
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						

		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify the L2 forwarding table should remove entries of the affected link at the neighbor switch.						
		Verify FHRP MAC in ARP/ND table.						
		Verify FHRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.						
		Verify that MAC's for SVI's are programmed as router/static entries on the switches where they are configured and learned as dynamic entries on the L2 peers.						
		On the aggregation switches, verify that the ARP/ND are programmed as adjacencies for L3 next hop forwarding after reload.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify IGMP/MLD snooping entries are deleted for the affected links at the access switches and re-learnt correctly on the alternative links after query from the IGMP snooping router.						
		Verify ACL/QoS TCAM is programmed correctly to share for ACL's and features that allow for sharing and verify ACL's are not sharing when not expected.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify PIM neighbor status.						
		Verify PIM both multipath and non-multipath functionalities.						

		Verify AutoRP mapping and boundaries.						
		Verify static RP mapping as the backup of auto RP.						
		Verify MSDP neighbors and SA cache consistency.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		On the multicast LHR, verify (*,G) and (S,G) creation based on SPT-threshold settings.						
		Verify PIM source register and register stop.						
		Verify GRE Tunnel re-route due to transport disruption.						
		Verify MTU fragmentation and reassembling at tunnel edge.						
		Verify BFD peer detection and client notifications.						
		The maximum traffic disruption for unicast will be half for both upstream and downstream traffic.						
		The maximum traffic loss for multicast upstream will be half and for downstream will be either 100% disrupted or no loss depending on which vPC peer switch reload.						
		Verify vPC peer status (role, peer link, keepalive link and consistency parameters)						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
1.4.9	Unicast ECMP		1425	1422	2	1	6277	
1.4.9.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers		24	24	0	0	120	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						

		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.					
		Verify SPAN is mirroring packets correctly.					
		Verify OTV traffic reconverges and optimize OSPF as needed.					
		Verify SNMP traps are sent to SNMP collector.					
		All unicast and multicast traffic should re-converge with proportionate packet loss.					
		Verify traffic destined for CoPP classes is policed as expected.					
		Verify OSPF interface status for the affected links.					
		Verify OSPF neighbor changes and authentication.					
		Verify OSPF DB/Topology consistency.					
		Verify OSPF routes and forwarding table consistency..					
		Verify OSPF multi-path load-balancing.					
		Verify HW and SW entries are properly programmed and synchronized.					
		Verify multicast HW and SW entries are properly programmed and synchronized.					
		Verify BFD peer detection and client notifications.					
		Verify frames delta does not increase.					
		Verify rx rate for all ixia ports are as expected (compared to baseline).					
		Verify packet loss duration is within expected range.					
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.					
1.4.9.2	L3 Port-channel Failure/Recovery between Spines		28	28	0	0	132	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.					
		Verify that there are no dead flows					
		Verify TB, error, crash					
		Verify interfaces in error					
		Verify any core dumps					
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.					

		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.4.9.3	L3 Port-channel member Failure/Recovery between Spines		12	12	0	0	72	
1.4.9.4	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf		449	449	0	0	2189	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						

		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.4.9.5	L3 Routed Port Failure/Recovery		598	597	0	1	2592	CSCum55853, CSCul51491
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify traffic is load balance to other ECMP paths						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						

		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.4.9.6	L3 Port-channel Failure/Recovery between Spine and Leaf		64	64	0	0	237	CSCul51491
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.						
		Verify the L2 forwarding table should remove entries of the affected link.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify OTV traffic reconverges and optimize OSPF as needed.						
		Verify SNMP traps are sent to SNMP collector.						
		All unicast and multicast traffic should re-converge with proportionate packet loss.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF interface status for the affected links.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency..						
		Verify OSPF multi-path load-balancing.						

		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut respectively.						
1.4.9.7	L3 port-channel member Failure/Recovery between Spine and Leaf		192	192	0	0	754	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify port-channel load balancing and rbh assignment						
		Verify traffic switches to high Bandwidth port-channels for both unicast and multicast when member failure and traffic will switch back when member recovers.						
		Verify LACP rebundle for port-channel after member recover.						
		The traffic should be able to re-converge within acceptable time.						
		Verify the convergence pattern is as expected.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
		Verify interface status is UP/DOWN state after linkNoShut/linkShut						

		respectively.						
1.4.9.8	ECMP hash-algorithm/hardware ecmp hash-offset change		11	11	0	0	11	
1.4.9.9	BGP AS-Path boundary conditions		2	0	2	0	2	CSCul87439
1.4.9.10	Clear Neighbors		17	17	0	0	51	
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						

		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.4.9.1 1	Clear Ipv4/IPv6 Unicast Routes		11	11	0	0	33	CSCul69815
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify OSPF IPv4/IPv6 neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify OSPF neighbor changes and authentication.						
		Verify OSPF DB/Topology consistency.						
		Verify OSPF routes and forwarding table consistency.						
		Verify OSPF multi-path load-balancing.						
		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and						

		synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						
1.4.9.12	Restart process		17	17	0	0	84	CSCul81364 CSCul81414
		Verify that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		All unicast and multicast traffic should re-converge.						
		Verify BGP neighbors will restart and come back correctly.						
		Verify that the hardware entries are properly removed and re-installed during the neighbor/process flapping.						
		Verify that CDP/LLDP does not lose peer information.						
		Verify that no flooding happens after traffic convergence.						
		Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.						
		Verify SPAN is mirroring packets correctly.						
		Verify SNMP traps are sent to SNMP collector.						
		Verify traffic destined for CoPP classes is policed as expected.						
		Verify BGP neighbor changes and authentication.						
		Verify BGP routes and forwarding table consistency.						
		Verify BGP multi-path load-balancing.						

		Verify HW and SW entries are properly programmed and synchronized.						
		Verify multicast HW and SW entries are properly programmed and synchronized.						
		Verify BFD peer detection and client notifications.						
		Verify the route tables for both unicast and multicast are updated correctly.						
		Verify the hardware entries, LC programming, fabric programming, outgoing interface, forwarding engine entries, for both unicast and multicast are updated correctly.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify packet loss duration is within expected range.						