Nexus Validation Test

Phase 3.2

1. Introduction

This is an addendum to the NVT Phase 3.0 report. Please reference the NVT Phase 3 report for network topology and feature descriptions (located at the following link):

http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/nvt/phase3/NVT-Phase_3.pdf

Please see below for additional test results from NVT Phase 3.2 which includes new software releases and network topologies.

2. Network Hardware and Software Version Details

2.1 DC31/DC34

Platform	Model No.	Software Version
N6004	N6K-C6004-96Q-SUP	7.0(2)N1(1)
N6001	N6K-C6001-64P-SUP	7.0(2)N1(1)
N3000	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N3548	N3K-C3548P-10G-SUP	6.0(2)A3(1)
N7010	N7K-SUP2E	6.2(8a)

NOTE: Although the NVT test cycle started, NVT did not complete the whole test cycle before the release of 7.0(2)N1(1). Please reference the section below, entitled "Caveats for DC31/DC34 (Nexus 6000)", for more information on the issues found during 7.0(2)N1(1). NVT testing will continue using the next maintenance release. The NVT test cases and results for DC31/DC34 will be published as a new addendum when the NVT test cycle will be completed.

Figure 1 DC31/DC34 Topologies

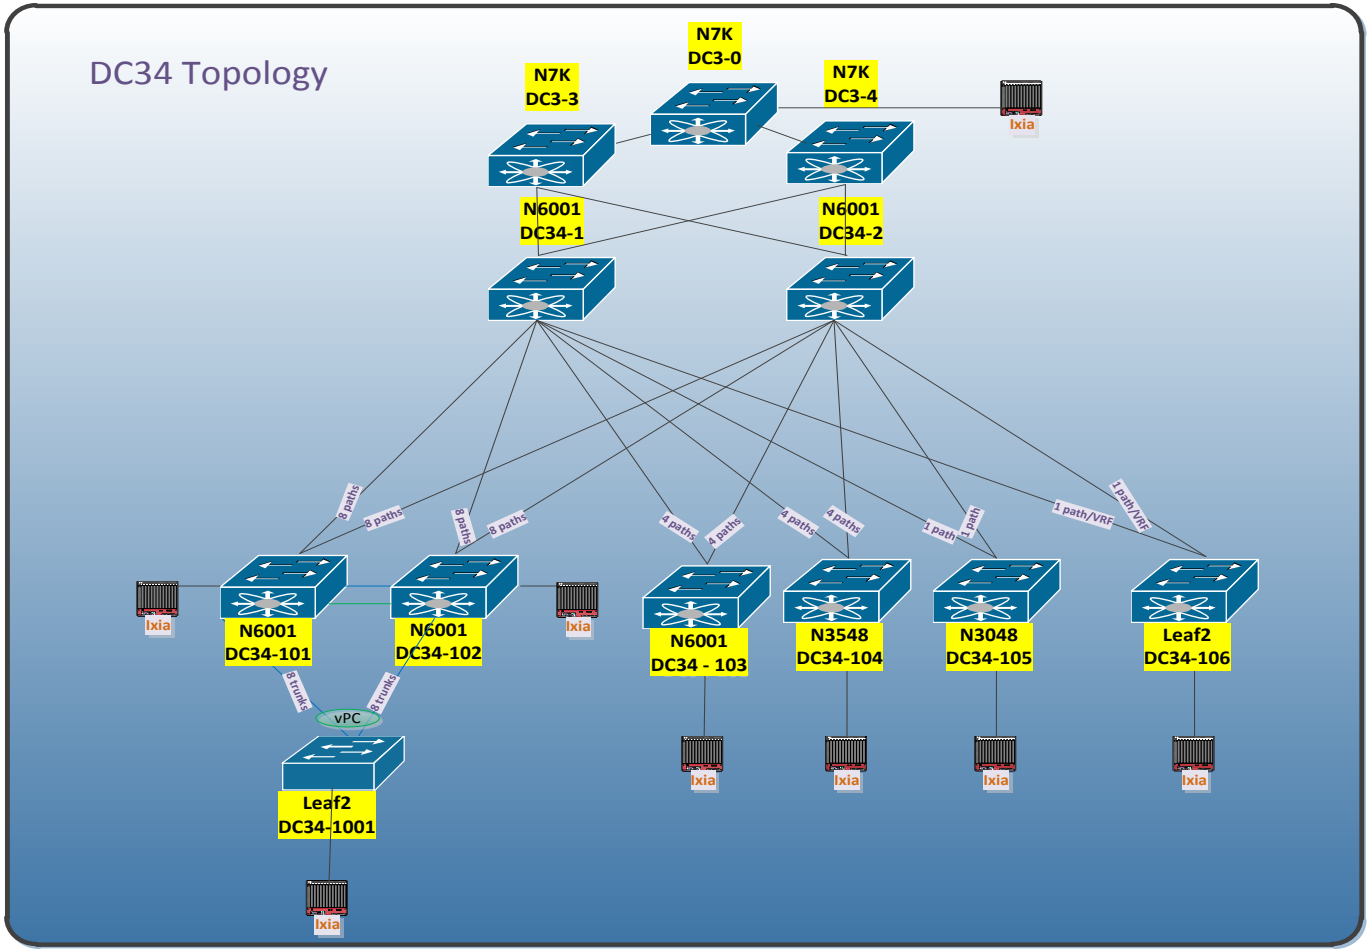

In NVT Phase 3.2, NVT has added a new Nexus 6000 POD (DC34). DC34 is very similar to the DC31 physical topology except the spine switches are Nexus 6001. DC31 has eBGP connections between the spine and the leaf layers. Instead, DC34 has OSPF configured between the spine and the leaf layers. The spine layer is configured as Area 0 and different leaves are configured in different areas.

2.2 DC32

During Phase 3.2, NVT re-executed the test cases from Phase 3.0 to verify IPv4 unicast and multicast multipath traffic on the Nexus 3548, Nexus 3048, and Nexus 7000 platforms. NVT has also augmented the topology with new Nexus 3548 switches in order to add NAT test cases.

Figure 2 DC32 Topology

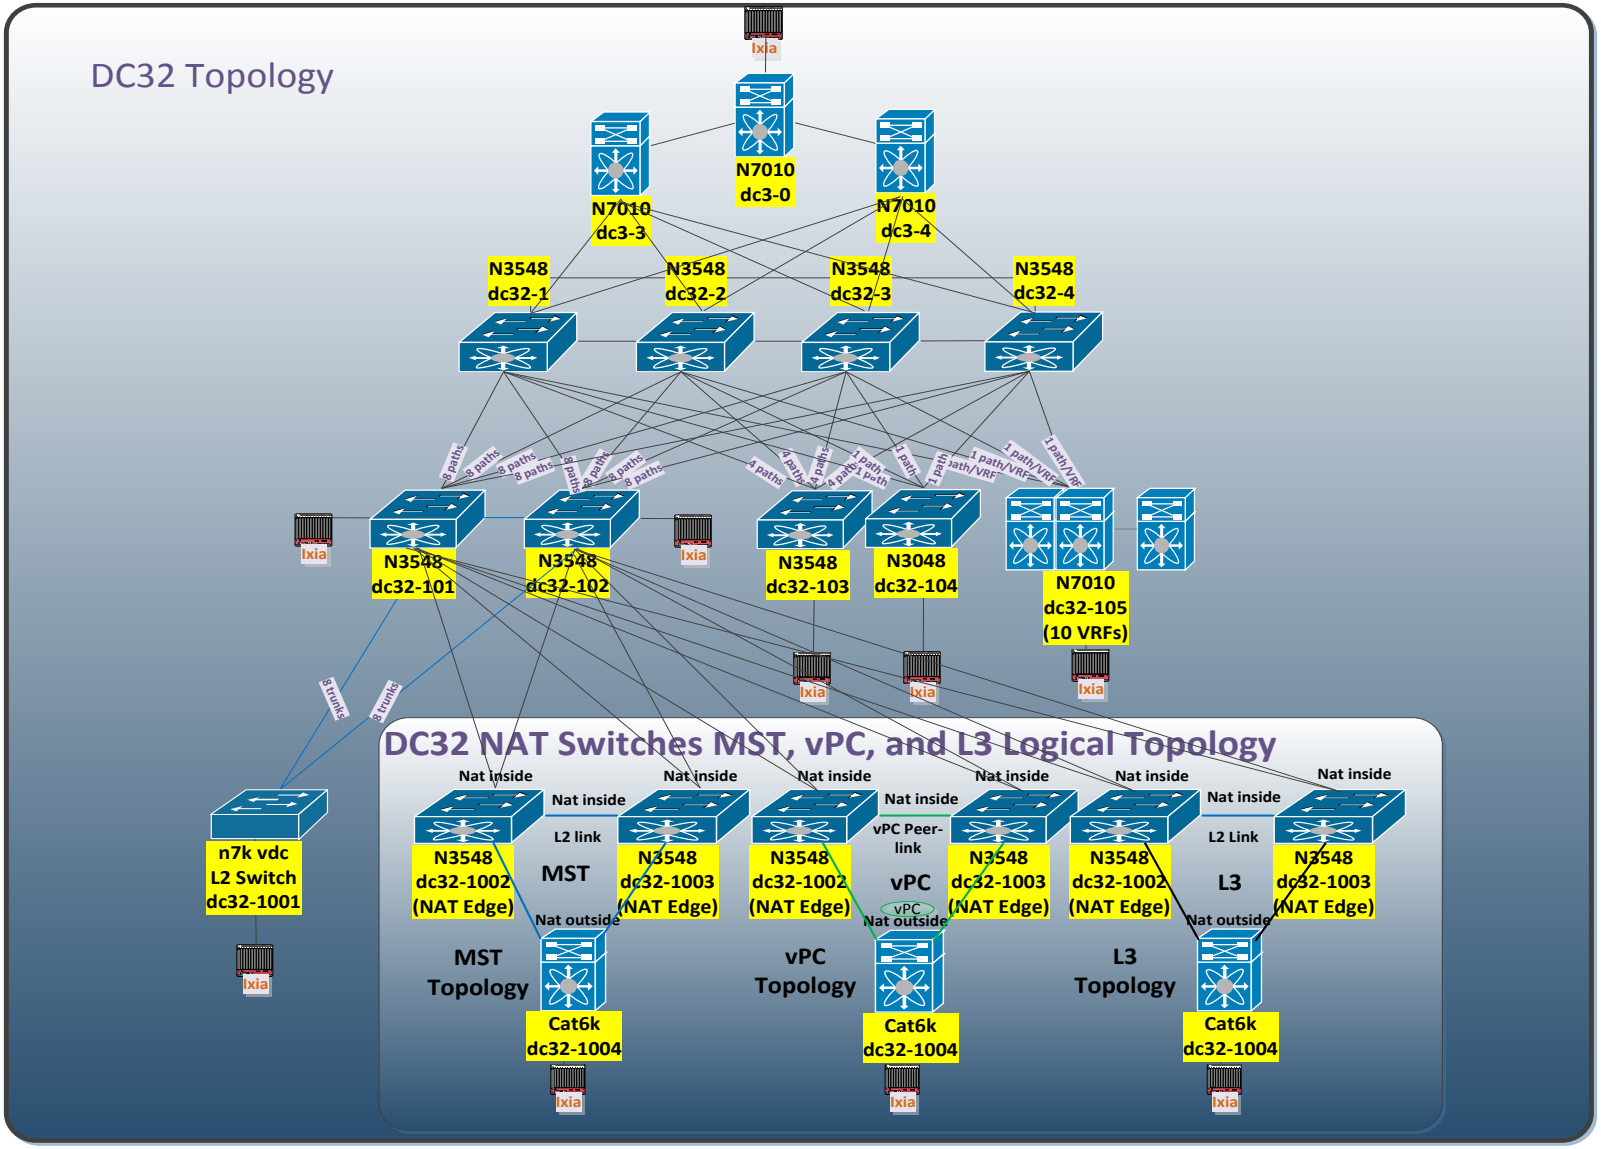

Platform	Model No.	Software Version
N3548	N3K-C3548P-10G-SUP	6.0(2)A3(1)
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N7010	N7K-SUP2E	6.2(8a)
CAT6K	VS-S720-10G	15.1(1)SY1

2.2.1 NAT

NAT feature has been enabled on the Nexus 3548 NAT edge switches DC32-1002 and DC32-1003. These two switches have been segmented into three logical switches for outside NAT interfaces while sharing the same inside NAT interfaces. The three topologies tested for NAT are Multiple Spanning Tree (MST), vPC and L3 routed access. Furthermore, HSRP has been deployed on the MST and vPC topologies. Please refer to Figure 2 DC32 Topology for more details.

To enable NAT on Nexus 3548 switches:

```
feature nat
```

2.2.1.1 Inside NAT interfaces

IP NAT inside has been configured on multiple ECMP interfaces. On the Nexus 3548, NAT configuration impacts ECMP functionality on interfaces with NAT configuration. Therefore, NAT unicast traffic will traverse one path until a disruption occurs on this link; in which case, traffic will switchover to the redundant NAT configured interface. IP NAT inside has also been configured on the SVI interface between NAT Edge peers (DC32-1002 and DC32-1003).

Configuration:

```
interface Ethernet1/1
  ip nat inside

interface Ethernet1/2
```

```
  ip nat inside

interface Vlan5
  ip nat inside
```

2.2.1.2 Outside NAT interfaces

Outside NAT interfaces are segmented into three logical switches with NAT unicast traffic over Multiple Spanning Tree (MST), vPC, and L3 routed interfaces. MST has been configured on SVIs 101-110 with HSRP configured on DC32-1002 and DC32-1003. VPC has been configured over SVIs 121-130 with HSRP configured on DC32-1002 and DC32-1003. While L3 routed interfaces have been configured on DC32-1002 and DC32-1003 for outside NAT interfaces with SVIs 141-150 configured on DC32-1004. Please refer to Figure 2 DC32 Topology for more details.

Configuration:

```
interface Vlan101
  description MST SVI
  no shutdown
  mtu 9216
  no ip redirects
  ip address 50.32.101.2/24
  ip pim sparse-mode
  ip nat outside
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 99
    ip 50.32.101.1
…

interface Vlan130
  description vPC SVI
  no shutdown
  mtu 9216
  no ip redirects
  ip address 50.32.130.2/24
  ip pim sparse-mode
  ip nat outside
  hsrp version 2
  hsrp 1
    authentication md5 key-string cisco
    preempt delay minimum 120
    priority 101
```

```
    ip 50.32.130.1

interface Ethernet1/5
  description dedicated NAT outside routed interface
  no switchport
  mtu 9216
  ip address 50.32.2.2/24
  ip pim sparse-mode
  ip nat outside
```

2.2.1.3 Static NAT for an Inside Source Address Translation

Static NAT for inside source address translation is configured on the Nexus 3548 NAT edge peers. 600 static NAT translations have been configured on the edge peers.

Configuration:

```
ip nat inside source static 132.103.103.50 50.103.103.50
ip nat inside source static 132.103.103.51 50.103.103.51
…
ip nat inside source static 132.103.105.249 50.103.105.249
```

NAT Translations:

Unicast	Inside SRC DC32-103	NAT Translated Source	Vlan	Remote Destination	Local Src Flows <-> Global Destination
MST	132.103.101-103.(50-250)	50.103.101-103.x	101-103	50.32.101-103.(50-250)	132.103.101-103.(50-250) <-> 50.32.101-103.(50-250)
vPC	132.103.121-123.(50-250)	50.103.121-123.x	121-123	50.32.121-123.(50-250)	132.103.121-123.(50-250) <-> 50.32.121-123.(50-250)
L3	132.103.141-143.(50-250)	50.103.141-143.x	141-143	50.32.141-143.(50-250)	132.103.141-143.(50-250) <-> 50.32.141-143.(50-250)

2.3 DC33

During Phase 3.2, NVT re-executed the test cases from Phase 3.0 to verify IPv4 unicast and multicast multipath traffic on the Nexus 3048.

Platform	Model No.	Software Version
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N7010	N7K-SUP2E	6.2(8a)

2.3.1 Recommendations for DC33 (Nexus 3000)

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 PIM RP will forward traffic on both shared and source trees during the switchover to the source tree. However, the Nexus 30xx platforms will not drop the traffic received on the (*,G) when the (S,G) is also present with a different RPF interface (Please refer to CSCuj89158 for more details).

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 are present on the switch.

```
hardware profile multicast prefer-source-tree eternity
```

Advantages:

As seen in the convergence graphs below, packet duplication is eliminated for certain test cases such as local and remote link failure and recovery. The graphs shown below are for link failure and recovery between a spine switch and one of the leaf switches. Figure 4 & Figure 6 show network convergence for 6.0(2)U3(1) without *hardware profile multicast prefer-source-tree eternity* command whearas Figure 5 & Figure 7 show convergence with *hardware profile multicast prefer-source-tree eternity* command configured.

Legend:

EW: East West traffic which stays within spine & leaf switches.
NS: North South traffic from upstream core switches towards leaf switches.

Figure 1 Legend

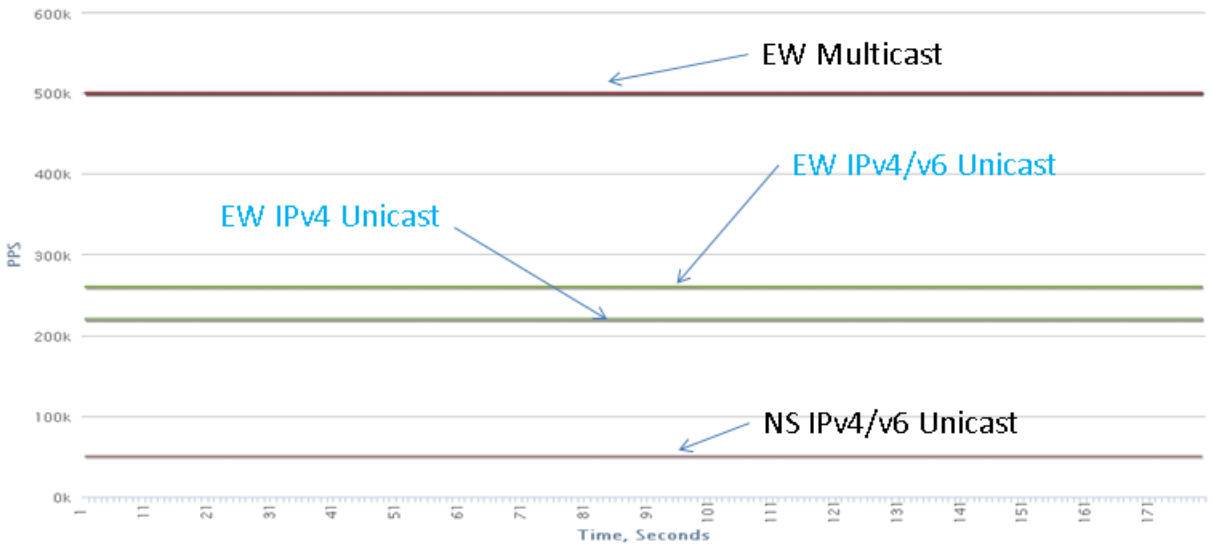

Link Shut (No Etnerity vs Eternity)

Figure 2 Software Version 6.0(2)U3(1) without Eternity command (Link shut)

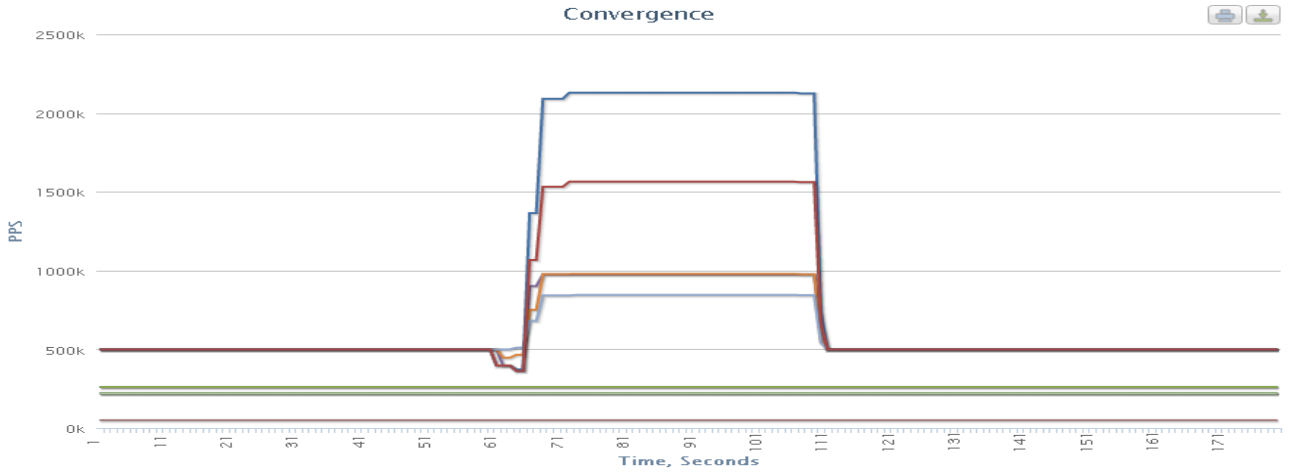

Figure 3 Software Version 6.0(2)U3(1) with Eternity command (Link Shut)

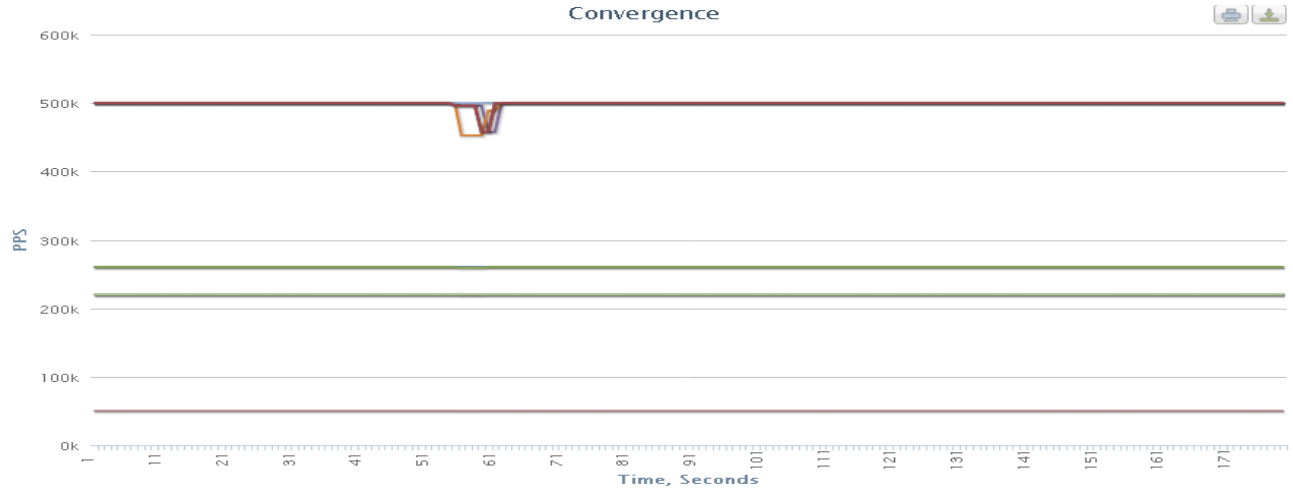

Link NO Shut (No Etnerity vs Eternity)

Figure 4 Software Version 6.0(2)U3(1) without Eternity command (Link NO shut)

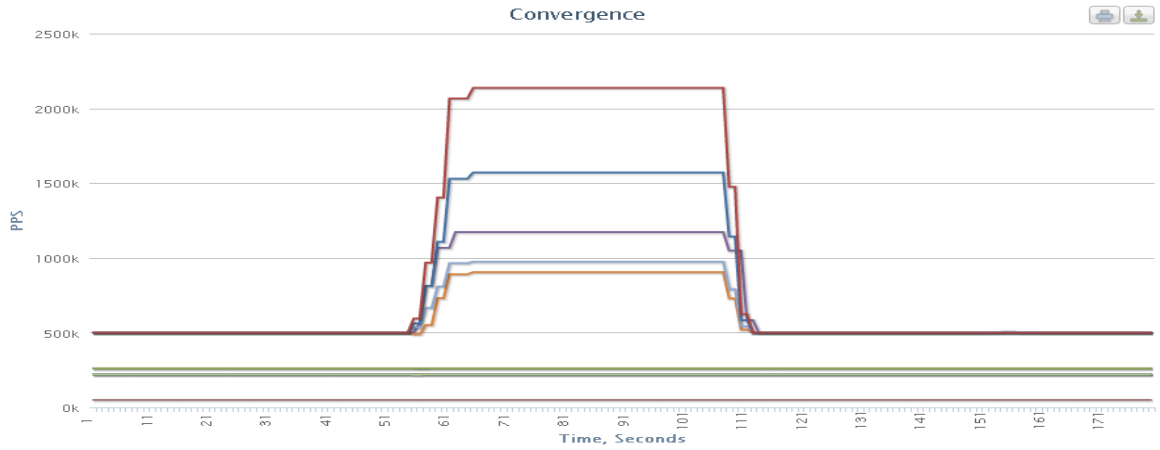

Figure 5 Software Version 6.0(2)U3(1) with Eternity command (Link NO Shut)

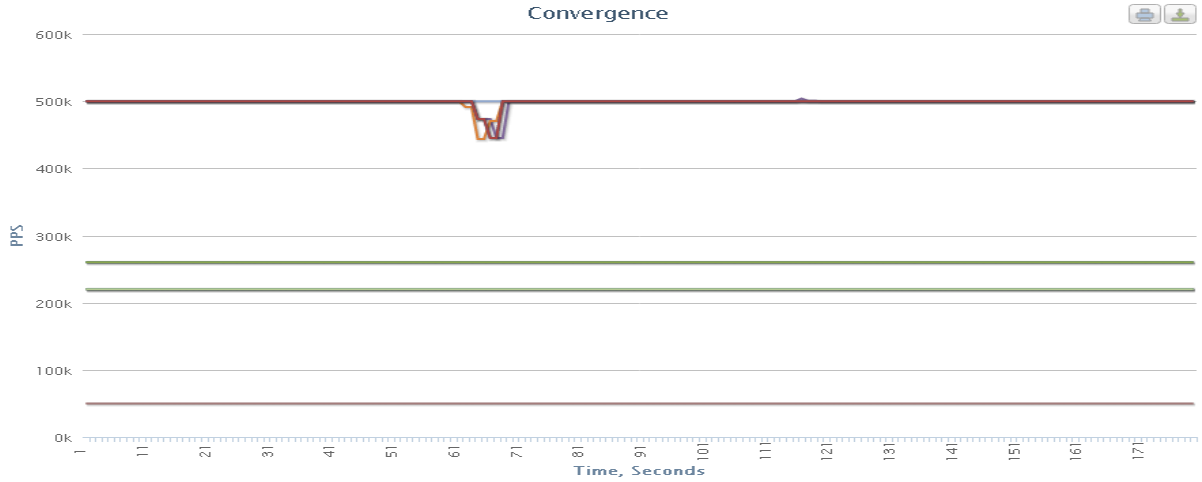

Disadvantages:

When this command is used, the switch learns (S, G) mroutes at a slower rate which will cause slower switchover from shared to source tree (the switch supports source (S, G) route injections at a rate of only 500 routes every two minutes). Furthermore, the multicast routing table must have at least 500 entries free for source (S, G) routes.

2.4 DC35

NVT has added a new topology based on the N3548 in Phase 3.2. The physical and logical designs are similar to DC32, except that it uses OSPF instead of BGP for unicast routing. VPC is a new feature for the N3548 and is covered as part of this topology.

Figure 6 DC35 Topology

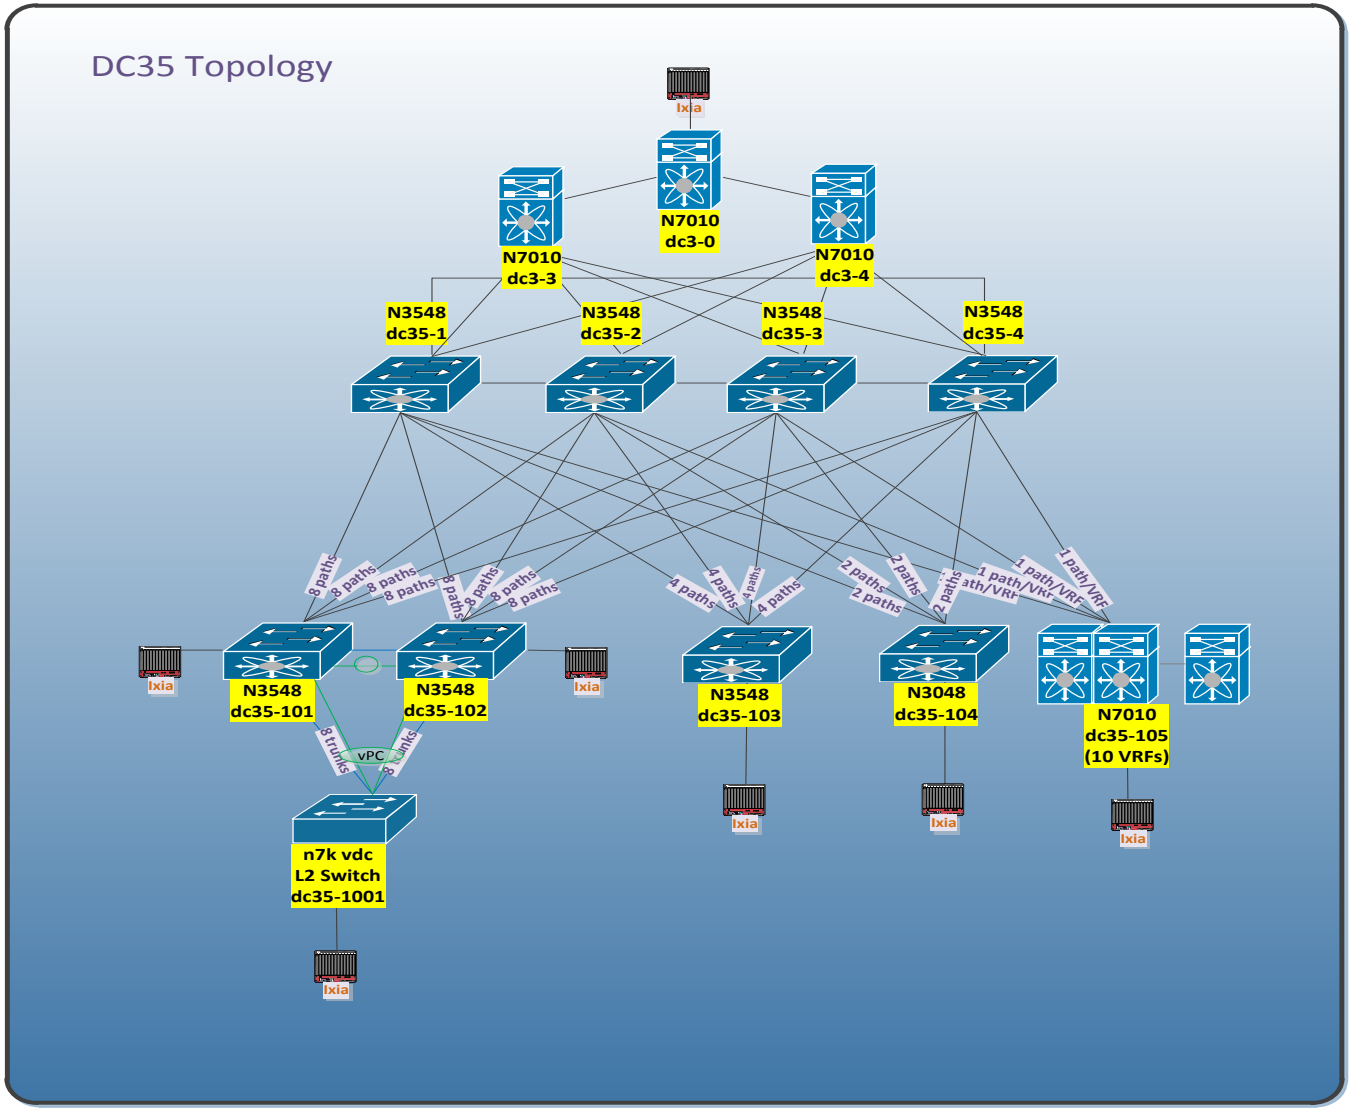

Platform	Model No.	Software Version
N3548	N3K-C3548P-10G-SUP	6.0(2)A3(1)
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N7K	N7K-SUP2E	6.2(8a)
CAT6K	WS-SUP720-3BXL	15.1(1)SY

2.4.1 VPC

vPC has been configured between two N3548 leafs, DC35-101 and DC35-102 with DC35-102 being configured as the primary device. Orphan ports on the two vPC leafs are also connected to Ixia traffic generators, simulating hosts.

2.5 DC36

The DC36 topology for Phase3.0 did not have multicast coverage. NVT has added multicast to DC36 for this phase.

Platform	Model No.	Software Version
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N3064	N3K-C3064PQ-10GE-SU	6.0(2)U3(1)
Cat 6500	WS-C6509-E	15.1(1)SY1
N7010	N7K-SUP2E	6.2(8a)

2.6 DC37

NVT has added a new topology as part of Phase 3.2 to cover the Nexus 3172. The logical topology is similar to the other phase 3 spine/leaf topologies discussed in the Phase 3.0 document. Interoperability with the N7K, N3548, N30xx, and Cat6K series of switches are also covered. Because the topology is logically similar to DC33 from Phase 3.0, the test cases are also similar.

Figure 7 DC37 Topology

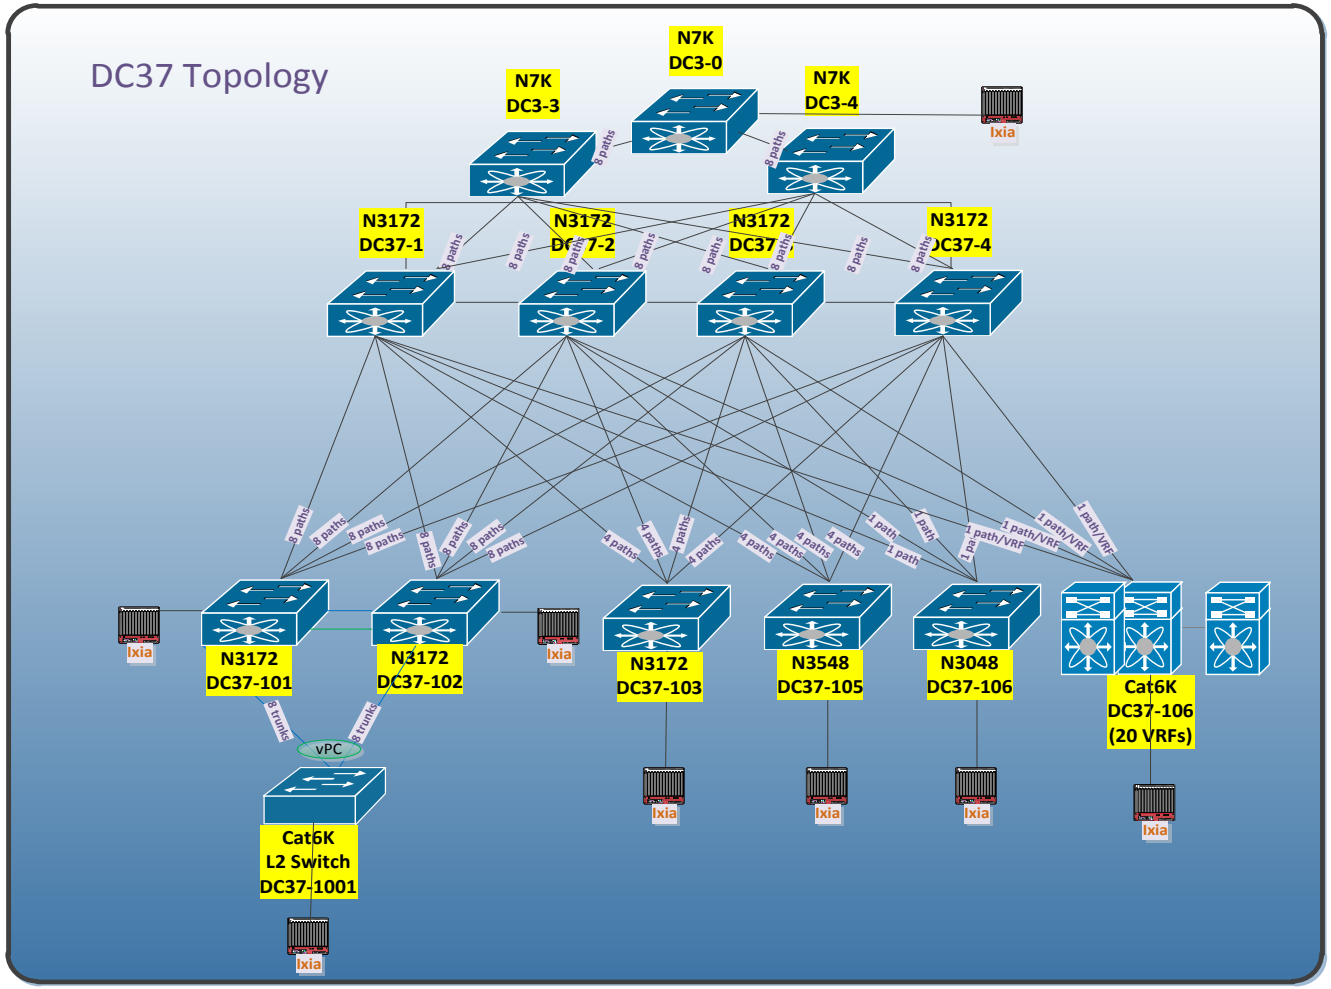

Platform	Model No.	Software Version
N3172	N3K-C3172PQ-10GE-SU	6.0(2)U3(1)
N3548	N3K-C3548P-10G-SUP	6.0(2)A3(1)
N3048	N3K-C3048TP-1GE-SUP	6.0(2)U3(1)
N7010	N7K-SUP2E	6.2(8a)
CAT6K	VS-SUP2T-10G	15.0(1)SY6

3. NVT Findings/Conclusion

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*

3.1 Caveats for DC31/DC34 (Nexus 6000)

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". There is no functional impact; this is just a transient state until the registration process is complete on the DR.
Workaround: None, there is no functional impact.
Severity: Moderate
Status: Closed
Platform Seen: N6000
Resolved Releases:
Applicable Releases: 6.0(2)N2(2)

CSCun64020
Symptom: 40G to 10G breakout port configuration lost.
Conditions: Upon upgrading from 6.0(2)N2(3) to 7.0(0)N1(1) using boot variable change and reload may cause loss of breakout configuration (eg. interface breakout slot 1 port 4-12 map 10g-4x).
Workaround: Save config prior to upgrade and reapply breakout port configuration.
Severity: Severe
Status: New
Platform Seen: Nexus 6000
Resolved Releases:
Applicable Releases: 7.0(0)N1(1) 7.0(1) N1(1) 7.0(2)N1(1)

CSCun84500

Symptom: Some source (S,G)s are not registered between Anycast RPs
Conditions: When Anycast RP peering is configured, each RP should update the other peers with any active sources using PIM NULL registers. With this bug, some of these NULL registers are missing, resulting in the peers not being aware of the affected sources. Any PIM joins to the affected peers will not receive multicast traffic for the missing sources.
Workaround: There's no workaround for this issue; network requiring this feature must upgrade to 7.0(2)N1(1).
Severity: Severe
Status: Resolved
Platform Seen: Nexus 6000
Resolved Releases: 7.0(2)N1(1)
Applicable Releases: 7.0(1)N1(1) 7.0(0)N1(1)

CSCuo46566
Symptom: Multicast traffic is not received on the orphan port attached to the PIM NDR.
Conditions: When two Nexus 6000s are configured in vPC, the multicast traffic is not received on the orphan port attached to the PIM NDR.
Workaround: There's no workaround for this issue; network requiring this feature must upgrade to 7.0(2)N1(1).
Severity: Severe
Status: Resolved
Platform Seen: Nexus 6000
Resolved Releases: 7.0(2)N1(1)
Applicable Releases: 7.0(0)N1(1) 7.0(1)N1(1)

CSCuo49778
Symptom: Not all of the source (S,G)s are properly stored in the MSDP peer SA caches.
Conditions: When Anycast RP with MSDP is configured, the RP should send source active messages to its peers. With this bug, some source active (SA) updates are not sent to other RPs. Consequently, any PIM joins toward the RP with the missing MSDP SA will not receive traffic from those sources.
Workaround: There's no workaround for this issue; network requiring this feature must upgrade to 7.0(2)N1(1).
Severity: Severe
Status: Resolved
Platform Seen: Nexus 6000
Resolved Releases: 7.0(2)N1(1)
Applicable Releases: 7.0(0)N1(1) 7.0(1)N1(1)

3.2 Caveats for DC32 (Nexus 3548)

CSCuo00926
Symptom: SNMP walk operations causing /var/tmp/vdc.log to grow without any upper bound in file size.
Conditions: Polling objects in the entityMIB will result in log entries added to the /var/tmp/vdc.log file in bootflash without any upper bound in file size. All the space on the bootflash may be consumed. This may affect other applications that require space from the bootflash.
Workaround: Use the CLI "filesys delete /var/tmp/vdc.log" to periodically delete the log file.
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3548
Resolved Releases: 6.0(2)A3(1)
Applicable Releases: 6.0(2)A1(1c) 6.0(2)A2(1)

CSCuo14680
Symptom: AFM may core during NAT configuration
Conditions: Adding static nat translations to switch configuration causes AFM to core and switch to reload.
Workaround: None
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3548
Resolved Releases: 6.0(2)A3(1)
Applicable Releases: None

CSCuo30063
Symptom: Downgrade may cause switch to lose config and login info.
Conditions: Performing install all downgrade from 6.0(2)A3(1) to 6.0(2)A1(1e) may cause loss of config and login info.
Workaround: Save config prior to upgrade and reapply lost configuration. If unable to login, perform password recovery procedure.
Severity: Severe
Status: Unreproducible
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A3(1)

CSCup25035
Symptom: PIM join from one vPC peer to the other peer over the peerlink does not work

Conditions: A vPC peer receiving PIM join from the other vPC peer through the peerlink does not add forwarding entries for the join. Consequentlly, no multicast traffic is routed over the peerlink.
Workaround: Use a dedicated L3 link between the vPC peers for multicast routing.
Severity: Severe
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A3(1)

CSCuo49467
Symptom: On recovery of L3 ECMP interface between spine and leaf, multicast traffic is lost.
Conditions: When a L3 ECMP interface is recovered from failed state, the leaf switch may send a join to the spine over the new link. The spine switch adds a new OIF to the (*,g) route, but does not add it to some of the (s,g) routes. As a result, traffic for the affected (s,g) routes will not be forwarded to the leaf.
Workaround: Shut and no-shut the affected interface.
Severity: Moderate
Status: Assigned
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A3(1)

CSCup38101
Symptom: Traffic destined for TCP ports 80 or 443 get punted to CPU.
Conditions: With WARP mode and VRFs configured, TCP traffic with destination port 80 and 443 forwarded in VRFs will be punted to CPU. The CoPP will rate limit the packets and may potentially cause drops.
Workaround: Revert to normal mode operation.
Severity: Severe
Status: Resolved
Platform Seen: Nexus 3548
Resolved Releases: 6.0(2)A4(1)
Applicable Releases: 6.0(2)A3(1)

The following bugs are carried over from 6.0(2)A1(1c) release that are still observed in 6.0(2)A3(1) software release.

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) 6.0(2)A3(1)

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) 6.0(2)A3(1)

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 causing traffic disruption for the affected flows.
Workaround: None
Severity: Severe
Status: New
Platform Seen: N3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c) 6.0(2)A3(1)

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. No functional impact for the affected multicast entries.
Workaround: There is no functional impact however the use of routed port-channels instead of individual routed links is recommended to avoid this bug.
Severity: Moderate
Status: New
Platform Seen: Nexus 3548
Resolved Releases:
Applicable Releases: 6.0(2)A1(1c) 6.0(2)A3(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 - security".
Workaround: None, there is no functional impact.
Severity: Minor
Status: New
Platform Seen: N3000
Resolved Releases:
Applicable Releases: 6.0(2)U2(1) 6.0(2)A3(1)

3.3 Caveats for DC33 (Nexus 3048)

Following bugs are carried over from the 6.0(2)U1(3) and are still being observed in the 6.0(2)U3(1) release.

CSCuj64147
Symptom: Nexus 3000 configuration loss upon image upgrade
Conditions: Upgrading the image from 6.0(2)U1(3) to 6.0(2)U2(1) after changing the boot variable and using the reload command 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: Save config prior to upgrade and reapply lost port configuration.
Severity: Severe

Status: Assigned
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U2(1) 6.0(2)U3(1)

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) 6.0(2)U3(1)

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: Use “hardware profile multicast prefer-source-tree eternity” to always prefer the source tree multicast entry which will eliminate the duplication in this scenario.
Severity: Severe
Status: Closed
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U3(1)

CSCul08871
Symptom: Nexus 3000 temporary packet duplication upon spine router failure
Conditions: Spine router reload might cause temporary packet duplication.
Workaround: None
Severity: Severe
Status: New
Platform Seen: Nexus 3000

Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U3(1)

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: There is no functional impact however the use of routed port-channels instead of individual routed links is recommended to avoid this bug.
Severity: Severe
Status: Closed
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U3(1)

CSCul39829
Symptom: Nexus 3000 PIM register is polarized to one RP when deploying anycast RP.
Conditions: The PIM register message are polarized to the same RP even when deployed with Multicast with Multipath.
Workaround: None
Severity: Severe
Status: Assinged
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U3(1)

CSCul46458
Symptom: Nexus 3000 “show port-channel load-balance” CLI and running-config may return wrong information
Conditions: After changing the load-balance algorithm value, the CLI and the running-config will still show the port-channel load-balance default value.
Workaround: None
Severity: Minor
Status: New
Platform Seen: Nexus 3000
Resolved Releases:
Applicable Releases: 6.0(2)U1(3) 6.0(2)U3(1)

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) 6.0(2)U3(1)

3.4 Caveats for DC36 (Nexus 3064)

CSCuo81954
Symptom: Show feature does not work for LLDP
Conditions: Even though LLDP feature is enabled globally, the “show feature” CLI does not show that it’s enabled.
Workaround: “show lldp neighbor” will provide lldp neighbor information when lldp is enabled.
Severity: Moderate
Status: Verified
Platform Seen: Nexus 3XXX
Resolved Releases: 6.0(2)U3(1)
Applicable Releases: None

3.5 Caveats for DC37 (Nexus 3172)

CSCup30207
Symptom: Duplicated multicast packets sent to receivers connected via vPC.
Conditions: Multicast data traffic rerouted from vPC leg back on same VLAN when source and receivers are on the same VLAN
Workaround: None
Severity: Severe
Status: Assigned
Platform Seen: Nexus 3172
Resolved Releases:
Applicable Releases: 6.0(2)U3(1)

4. NVT Test Results

The following section contains test case and results for:

- DC32
- DC33
- DC35
- DC36
- DC37

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

4.1 DC32 NAT Test Cases

NAT Test Cases	Pass/Fail	Defect
shut dc32-1002 e1/1 -> 101 (ip nat inside uplink)	pass	
no-shut dc32-1002 e1/1 -> 101 (ip nat inside uplink)	pass	
shut dc32-1002 e1/2-> 101 (ip nat inside uplink)	pass	
no-shut dc32-1002 e1/2 -> 101 (ip nat inside uplink)	pass	
shut dc32-1002 e1/1-2 -> 101/102 (both ip nat inside uplinks 1002)	pass	
no-shut dc32-1002 e1/1-2 -> 101/102 (both ip nat inside uplinks 1002)	pass	
reload dc32-101	pass	
shut dc32-1002 vpc peerlink po 2003	pass	
no-shut dc32-1002 vpc peerlink po 2003	pass	
shut dc32-1002 vpc keepalive link e1/6	pass	

no-shut dc32-1002 vpc keepalive link e1/6	pass	
shut dc32-1002 vpc keepalive & peer link	pass	
no-shut dc32-1002 vpc keepalive & peer link	pass	
shut dc32-1002 vpc keepalive & peer link /w autorecovery	pass	
no-shut dc32-1002 vpc keepalive & peer link /w autorecovery	pass	
shut dc32-1002 po1004 (ip nat outside link)	pass	
noshut dc32-1002 po1004 (ip nat outside link)	pass	
shut dc32-1002 po2004 (vpc PO)	pass	
no shut dc32-1002 po 2004 (vpc PO)	pass	
reload dc32-1002	pass	
shut dc32-1003 e1/1 -> 101 (ip nat inside uplink)	pass	
no-shut dc32-1003 e1/1 -> 101 (ip nat inside uplink)	pass	
shut dc32-1003 e1/2 -> 102 (ip nat inside uplink)	pass	
no-shut dc32-1003 e1/2 -> 102 (ip nat inside uplink)	pass	
shut dc32-1003 e1/1-2 -> 101/102(both ip nat inside uplinks 1003)	pass	
no-shut dc32-1003 e1/1-2 -> 101/102 (both ip nat inside uplinks 1003)	pass	
reload dc32-102	pass	
reload dc32-1003	pass	
shut dc32-1003 po 1004 L2 PO	pass	
no-shut dc32-1003 po 1004 L2 PO	pass	
reload dc32-1004	pass	
shut dc32-1003 po2004 (vpc L2 PO)	pass	
no-shut dc32-1003 po2004 (vpc L2 PO)	pass	
dc32-1002 e1/5 (L3 ip int)	fail	CSCup25035
dc32-1003 e1/5 (L3 ip int)	fail	CSCup25035
Remove Nat configs	pass	
Re-add Nat configs	pass	

4.2 Frontend Report Summary

	Folders	Total # of Iteration	Total # of Test Cases	Total # of Pass	Total # of Pass w/Exception	Total # of Fail	Defect(s)
1	NVT 3.1	17723	3272	3247	13	12	
1.1	DC32	6921	1292	1276	8	8	
1.1.1	Configuration	336	32	32	0	0	
1.1.1.1	Common Configuration	42	4	4	0	0	
1.1.1.2	Ixia Setup/Configuration	42	4	4	0	0	
1.1.1.3	Interface and LACP Configs	42	4	4	0	0	
1.1.1.4	SVI and HSRP Configs	42	4	4	0	0	
1.1.1.5	SPT Configs (MST)	42	4	4	0	0	
1.1.1.6	OSPF Configs	42	4	4	0	0	
1.1.1.7	BGP Configs	42	4	4	0	0	
1.1.1.8	Mcast Configs	42	4	4	0	0	
1.1.2	Spine to Core Setup	42	4	4	0	0	
1.1.2.1	Spine to Core Setup	42	4	4	0	0	
1.1.3	Spine to Leaf Setup	42	4	4	0	0	
1.1.3.1	Spine to Leaf N3500 Setup	42	4	4	0	0	
1.1.4	Leaf to Spine Setup	23	4	4	0	0	
1.1.4.1	Leaf N3500 to Spine Setup	20	3	3	0	0	
1.1.4.2	Leaf N3000 to Spine Setup	3	1	1	0	0	

1.1.5	Leaf to Hosts Ixia Setup	22	4	4	0	0	
1.1.5.1	Leaf to Hosts Ixia Setup	22	4	4	0	0	
1.1.6	Unicast ECMP	3107	615	611	0	4	
1.1.6.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	67	16	16	0	0	
1.1.6.2	L3 Port-channel Failure/Recovery between Spines	68	16	16	0	0	
1.1.6.3	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf	1150	232	229	0	3	CSCuo49467,CSCum63413,CSCul27903
1.1.6.4	L3 Routed Port Failure/Recovery	1185	238	237	0	1	CSCum63413
1.1.6.5	L3 Port-channel Failure/Recovery between Spine and Leaf	524	71	71	0	0	
1.1.6.6	L3 port-channel member Failure/Recovery between Spine and Leaf	87	30	30	0	0	
1.1.6.7	Clear Neighbors	8	4	4	0	0	
1.1.6.8	Clear Ipv4/IPv6 Unicast Routes	8	4	4	0	0	
1.1.6.9	Restart process	10	4	4	0	0	
1.1.7	L2 Link Failure/Recovery	167	21	19	0	2	
1.1.7.1	L2 Port-channel Failure/Recovery between Leaf and ToR devices	123	20	19	0	1	CSCuo81954
1.1.7.2	L2 Port-channel Failure/Recovery between Leaf devices	44	1	0	0	1	CSCun31859
1.1.8	Multicast with Multipath	3114	602	595	7	0	
1.1.8.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	27	4	3	1	0	CSCuo49467
1.1.8.2	L3 Port-channel Failure/Recovery between Spines	68	16	16	0	0	
1.1.8.3	L3 Progressive Routed Port Failure then Recovery between Spine and Leaf	1154	231	229	2	0	CSCul27903, CSCul27880
1.1.8.4	L3 Routed Port Failure/Recovery	1189	240	239	1	0	CSCum63413
1.1.8.5	L3 Port-channel Failure/Recovery between Spine and Leaf	510	70	67	3	0	CSCuo49467, CSCul27903, CSCum63413
1.1.8.6	L3 port-channel member Failure/Recovery between Spine and Leaf	87	29	29	0	0	

1.1.8.7	RP,DR Failure	69	7	7	0	0	
1.1.8.8	Clear PIM Routes	10	5	5	0	0	
1.1.9	Reload and Power Cycle Switch	68	6	3	1	2	
1.1.9.1	Reload Spine	36	4	2	0	2	CSCuo00926,CSCuo30063
1.1.9.2	Reload Leaf	32	2	1	1	0	CSCun31859
1.2	DC33	3879	966	959	5	2	
1.2.1	Configuration	720	24	20	4	0	CSCuj67358,CSCul46510,CSCul39829,CSCul46458
1.2.1.1	Common Configuration	72	2	2	0	0	
1.2.1.2	Ixia Setup/Configuration	72	2	2	0	0	
1.2.1.3	Interface and LACP Configs	72	2	2	0	0	
1.2.1.4	SVI and HSRP Configs	72	2	2	0	0	
1.2.1.5	SPT Configs (MST)	72	2	2	0	0	
1.2.1.6	OSPF Configs	72	2	2	0	0	
1.2.1.7	BGP Configs	72	2	2	0	0	
1.2.1.8	Mcast Configs	72	2	2	0	0	
1.2.1.9	NAT	72	2	2	0	0	
1.2.1.10	BFD	72	2	2	0	0	
1.2.2	Unicast ECMP	1789	476	476	0	0	
1.2.2.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	24	4	4	0	0	
1.2.2.2	L3 Port-channel Failure/Recovery between Spines	118	31	31	0	0	
1.2.2.3	L3 Port-channel member Failure/Recovery between Spines	60	16	16	0	0	
1.2.2.4	L3 Port-channel member Failure/Recovery between Spine and Leaf	190	112	111	1	0	CSCul28087

1.2.2.5	L3 Routed Port Failure/Recovery	1112	292	292	0	0	
1.2.2.6	L3 Port-channel Failure/Recovery between Spine and Leaf	233	8	8	0	0	
1.2.2.7	Clear Neighbors	27	7	7	0	0	
1.2.2.8	Clear Ipv4/IPv6 Unicast Routes	16	4	4	0	0	
1.2.2.9	Restart process	9	3	3	0	0	
1.2.3	L2 Link Failure/Recovery	110	19	19	0	0	
1.2.3.1	vPC leg failure/recovery between Leaf and ToR	86	15	15	0	0	
1.2.3.2	vPC peer-link failure/recovery between Leaf vPC peer switches	24	4	4	0	0	
1.2.4	Multicast with Multipath	1242	439	439	0	0	
1.2.4.1	L3 Port-channel member Failure/Recovery between Spines	200	46	46	0	0	
1.2.4.2	L3 Routed Port Failure/Recovery	834	276	276	0	0	
1.2.4.3	L3 port-channel member Failure/Recovery between Spine and Leaf	190	112	112	0	0	
1.2.4.4	RP,DR Failure	18	5	5	0	0	
1.2.5	Reload and Power Cycle Switch	18	7	5	0	2	CSCul08871,CSCuj64147
1.2.5.1	Reload Spine	6	2	2	0	0	
1.2.5.2	Reload Leaf	12	3	3	0	0	
1.3	DC35	2861	297	297	0	0	
1.3.1	Configuration	190	20	20	0	0	
1.3.1.1	Common Configuration	19	2	2	0	0	
1.3.1.2	Ixia Setup/Configuration	19	2	2	0	0	
1.3.1.3	Interface and LACP Configs	19	2	2	0	0	
1.3.1.4	SVI and HSRP Configs	19	2	2	0	0	

1.3.1.5	SPT Configs (MST)	19	2	2	0	0	
1.3.1.6	OSPF Configs	19	2	2	0	0	
1.3.1.7	BGP Configs	19	2	2	0	0	
1.3.1.8	Mcast Configs	19	2	2	0	0	
1.3.1.9	NAT	19	2	2	0	0	
1.3.1.10	BFD	19	2	2	0	0	
1.3.2	Leaf to Hosts Setup	105	7	7	0	0	
1.3.2.1	Leaf to N7K Switch Setup	35	2	2	0	0	
1.3.2.2	Leaf to Hosts Ixia Setup	70	5	5	0	0	
1.3.3	Unicast ECMP	1395	147	147	0	0	
1.3.3.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	36	4	4	0	0	
1.3.3.2	L3 Port-channel Failure/Recovery between Spines	2	2	2	0	0	
1.3.3.3	L3 Routed Port Failure/Recovery	527	35	35	0	0	
1.3.3.4	L3 Port-channel Failure/Recovery between Spine and Leaf	447	68	68	0	0	
1.3.3.5	L3 port-channel member Failure/Recovery between Spine and Leaf	383	38	38	0	0	
1.3.4	L2 Link Failure/Recovery	213	12	12	0	0	
1.3.4.1	vPC leg failure/recovery between Leaf and ToR	131	8	8	0	0	
1.3.4.2	vPC peer-link failure/recovery between Leaf vPC peer switches	82	4	4	0	0	
1.3.5	Multicast with Multipath	958	95	95	0	0	
1.3.5.1	RPF Failure/Recovery between first leaf and elected RP	18	2	2	0	0	
1.3.5.2	L3 Port-channel Failure/Recovery between Core and Distribution Layers	36	4	4	0	0	
1.3.5.3	L3 Port-channel Failure/Recovery between Spines	2	2	2	0	0	

1.3.5.4	L3 Routed Port Failure/Recovery	450	35	35	0	0	
1.3.5.5	L3 Port-channel Failure/Recovery between Spine and Leaf	212	28	28	0	0	
1.3.5.6	L3 port-channel member Failure/Recovery between Spine and Leaf	183	22	22	0	0	
1.3.5.7	RP,DR Failure	57	2	2	0	0	
1.3.6	Reload and Power Cycle Switch	57	7	7	0	0	
1.3.6.1	Reload Spine	25	2	2	0	0	
1.3.6.2	Reload Leaf	32	5	5	0	0	
1.3.7	Leaf to Spine Setup	60	5	5	0	0	
1.3.7.1	Leaf N3500 to N3500 Spine	54	3	3	0	0	
1.3.7.2	Leaf N3000 to N3500 Spine	3	1	1	0	0	
1.3.7.3	Leaf C6K to N3500 Spine	3	1	1	0	0	
1.3.8	L3 Link Failure/Recovery	36	4	4	0	0	
1.3.8.1	L3 Port-channel Failure/Recovery between Core and N3500-BGP	36	4	4	0	0	
1.4	DC36	2393	501	500	0	1	
1.4.1	Configuration	633	81	81	0	0	
1.4.1.1	Common Configuration	94	11	11	0	0	
1.4.1.2	Ixia Setup/Configuration	97	12	12	0	0	
1.4.1.3	Interface and LACP Configs	97	12	12	0	0	
1.4.1.4	SVI and HSRP Configs	39	7	7	0	0	
1.4.1.5	SPT Configs (MST)	39	7	7	0	0	
1.4.1.6	OSPF Configs	63	6	6	0	0	
1.4.1.7	BGP Configs	97	12	12	0	0	

1.4.1.8	Mcast Configs	97	12	12	0	0	
1.4.1.9	NAT	5	1	1	0	0	
1.4.1.10	BFD	5	1	1	0	0	
1.4.2	Spine to Core Setup	23	6	6	0	0	
1.4.2.1	Spine to Core Setup	23	6	6	0	0	
1.4.3	Spine to Leaf Setup	23	6	6	0	0	
1.4.3.1	Spine to Leaf Setup	23	6	6	0	0	
1.4.4	Leaf to Spine Setup	19	7	6	0	1	CSCuo81954
1.4.4.1	Leaf N3048 to N3k Spine	10	3	3	0	0	
1.4.4.2	Leaf Cat6k to N3k Spine	3	1	1	0	0	
1.4.4.3	Leaf N3064 to N3k Spine	6	2	2	0	0	
1.4.5	Leaf to Hosts Ixia Setup	19	6	6	0	0	
1.4.5.1	Leaf to Hosts Ixia Setup	19	6	6	0	0	
1.4.6	Reload and Power Cycle Switch	18	6	6	0	0	
1.4.6.1	Reload Leaf	11	6	6	0	0	
1.4.7	Unicast ECMP	1002	205	205	0	0	
1.4.7.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	48	16	16	0	0	
1.4.7.2	L3 Routed Port Failure/Recovery	603	150	150	0	0	
1.4.7.3	L3 Port-channel Failure/Recovery between Spine and Leaf	247	35	35	0	0	
1.4.7.4	L3 port-channel member Failure/Recovery between Spine and Leaf	48	4	4	0	0	
1.4.8	Multicast with Multipath	669	184	184	0	0	
1.4.8.1	L3 Port-channel Failure/Recovery between Core and Distribution Layers	48	16	16	0	0	

1.4.8.2	L3 Routed Port Failure/Recovery	532	135	135	0	0	
1.4.8.3	L3 Port-channel Failure/Recovery between Spine and Leaf	78	25	25	0	0	
1.4.8.4	RP,DR Failure	7	4	4	0	0	
1.4.8.5	Restart process	4	4	4	0	0	
1.5	DC37	1669	216	215	0	1	CSCup30207
1.5.1	Configuration	430	40	40	0	0	
1.5.1.1	Common Configuration	43	4	4	0	0	
1.5.1.2	Ixia Setup/Configuration	43	4	4	0	0	
1.5.1.3	Interface and LACP Configs	43	4	4	0	0	
1.5.1.4	SVI and HSRP Configs	43	4	4	0	0	
1.5.1.5	SPT Configs (MST)	43	4	4	0	0	
1.5.1.6	OSPF Configs	43	4	4	0	0	
1.5.1.7	BGP Configs	43	4	4	0	0	
1.5.1.8	Mcast Configs	43	4	4	0	0	
1.5.1.9	NAT	43	4	4	0	0	
1.5.1.10	BFD	43	4	4	0	0	
1.5.2	Leaf to Spine Setup	9	5	5	0	0	
1.5.2.1	Leaf N3000 to N3K Spine	9	5	5	0	0	
1.5.3	Leaf to Hosts Setup	4	4	4	0	0	
1.5.3.1	Leaf to N7K Switch Setup	1	1	1	0	0	
1.5.3.2	Leaf to Hosts Ixia Setup	3	3	3	0	0	
1.5.4	Unicast ECMP	472	68	68	0	0	
1.5.4.1	L3 Routed Port Failure/Recovery	214	20	20	0	0	
1.5.4.2	L3 Port-channel Failure/Recovery between Spine and Leaf	258	48	48	0	0	
1.5.5	L2 Link Failure/Recovery	70	10	10	0	0	
1.5.5.1	vPC leg failure/recovery between Leaf and ToR	54	6	6	0	0	
1.5.5.2	vPC peer-link failure/recovery between Leaf vPC peer switches	16	4	4	0	0	

1.5.6	Multicast with Multipath	684	79	79	0	0	
1.5.6.1	RPF Failure/Recovery between first leaf and elected RP	26	2	2	0	0	
1.5.6.2	L3 Routed Port Failure/Recovery	238	20	20	0	0	
1.5.6.3	L3 Port-channel Failure/Recovery between Spine and Leaf	372	48	48	0	0	
1.5.6.4	RP,DR Failure	48	9	9	0	0	
1.5.7	Reload and Power Cycle Switch	48	9	9	0	0	
1.5.7.1	Reload Spine	29	4	4	0	0	
1.5.7.2	Reload Leaf	19	5	5	0	0	

4.3 Excerpt from Frontend Report for DC37

	Folders	Verification	Total # of Iteration	Total # of test cases	Total # of Pass	Total # of Pass w/Exception	Total # of Fail	Defect(s)
1.5	DC37		1669	216	215	0	1	CSCup30207
1.5.1	Configuration		430	40	40	0	0	
1.5.1.1	Common Configuration		43	4	4	0	0	
1.5.1.1.1	DC37-1-none-none		39					
		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 SSH works through the management network on a dedicated vrf						
		Verify RSA key does not change on device						
		Verify error vlans						
		Verify frames delta does not increase.						
1.5.4	Unicast ECMP		472	68	68	0	0	
1.5.4.1	L3 Routed Port Failure/Recovery		214	20	20	0	0	
1.5.4.1.1	DC37-1-LinkShut-e1/1		20					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		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 that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		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.						
		Verify frames delta does not increase before link shut						
		Verify frames delta does not increase after link no shut						
1.5.4.1.2	DC37-1-LinkNoShut-e1/1		19					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		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 that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						

		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.						
		Verify frames delta does not increase before link shut						
		Verify frames delta does not increase after link no shut						
1.5.5	L2 Link Failure/Recovery		70	10	10	0	0	
1.5.5.1	vPC leg failure/recovery between Leaf and ToR		54	6	6	0	0	
1.5.5.1.1	DC37-101-failureL2Shut-Po11		13					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		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 that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		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.						
		Verify frames delta does not increase before link shut						
		Verify error vlans						
		Verify mac sync for VPC setup (Compare mac entries are same in both VPC peers before shut)						
		Verify mac addresses are not learned via vPC Peer-Link before primary link shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,Before Shut.						

		Verify CDP is enabled globally						
		Verify LLDP is enabled globally.						
		Verify cdp status are appropriate before failure						
		verify lldp status are appropriate before Failure						
		Verify IGMP is enabled globally						
		Verify that the STP state of all Vlans are in same state that of corresponding STP interface (RSTP) before shut						
		Verify if the STP interfaces are in FWD state (MST) before shut						
		Verify traffic drop by checking Rx rate in all ixia ports after Shut						
		Verify flooding by checking Rx rate in all ixia ports after Shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,After Shut.						
		Verify mac addresses are removed from the link after link shut						
		Verify mac addresses are moved from primary link to vPc peer-link after primary link is shut						
		Verify cdp peer entries are lost for affected links						
		Verify cdp entries does not lose peer information for unaffected links						
		Verify lldp peer entries are lost for affected links						
		Verify lldp entries does not lose peer information for unaffected links						
		Verify all IGMP snooping entries are same after link shut (vpc link)						
		verify ARP entries after link shut are same as before link shut						
		verify IGMP group membership after link shut is same as before link shut						
		Verify all DHCP Relay entries are same after link shut						
		Verify traffic drop by checking Rx rate in all ixia ports after No Shut						
		Verify flooding by checking Rx rate in all ixia ports after No Shut						
		Verify frames delta does not increase after link no shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,After No Shut.						
		Verify there are no missing MAC addresses after no shut						
		Verify vPc Peer-Link no longer has MAC addresses from initial capture of the primary link						
		Verify that cdp entries after No Shut are same as entries taken before						
		Verify that lldp entries after No Shut are same as entries taken before						
		Verify all IGMP snooping entries after link no shut are same as before the link shut						
		Verify that the STP state of all Vlans are in same state that of corresponding STP interface (RSTP) after no shut						
		Verify if the STP interfaces are in FWD state (MST) after no shut						
		verify ARP entries after link no shut are same as before link shut						
		verify IGMP group membership after link no shut is same as before link shut						

		Verify all DHCP Relay entries are same after link no shut						
		Verify VPC information after link no shut is same as before link shut						
1.5.5.1.2	DC37-101-failureL2NoShut-Po11		13					
		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		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 that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		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.						
		Verify frames delta does not increase before link shut						
		Verify error vlans						
		Verify mac sync for VPC setup (Compare mac entries are same in both VPC peers before shut)						
		Verify mac addresses are not learned via vPC Peer-Link before primary link shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,Before Shut.						
		Verify CDP is enabled globally						
		Verify LLDP is enabled globally.						
		Verify cdp status are appropriate before failure						
		verify lldp status are appropriate before Failure						
		Verify IGMP is enabled globally						

		Verify that the STP state of all Vlans are in same state that of corresponding STP interface (RSTP) before shut						
		Verify if the STP interfaces are in FWD state (MST) before shut						
		Verify traffic drop by checking Rx rate in all ixia ports after Shut						
		Verify flooding by checking Rx rate in all ixia ports after Shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,After Shut.						
		Verify mac addresses are removed from the link after link shut						
		Verify mac addresses are moved from primary link to vPc peer-link after primary link is shut						
		Verify cdp peer entries are lost for affected links						
		Verify cdp entries does not lose peer information for unaffected links						
		Verify lldp peer entries are lost for affected links						
		Verify lldp entries does not lose peer information for unaffected links						
		Verify all IGMP snooping entries are same after link shut (vpc link)						
		verify ARP entries after link shut are same as before link shut						
		verify IGMP group membership after link shut is same as before link shut						
		Verify all DHCP Relay entries are same after link shut						
		Verify traffic drop by checking Rx rate in all ixia ports after No Shut						
		Verify flooding by checking Rx rate in all ixia ports after No Shut						
		Verify frames delta does not increase after link no shut						
		Verify 30 sec Load Interval Input and output rate of the interfaces to be greater than user Specified rate,After No Shut.						
		Verify there are no missing MAC addresses after no shut						
		Verify vPc Peer-Link no longer has MAC addresses from initial capture of the primary link						
		Verify that cdp entries after No Shut are same as entries taken before						
		Verify that lldp entries after No Shut are same as entries taken before						
		Verify all IGMP snooping entries after link no shut are same as before the link shut						
		Verify that the STP state of all Vlans are in same state that of corresponding STP interface (RSTP) after no shut						
		Verify if the STP interfaces are in FWD state (MST) after no shut						
		verify ARP entries after link no shut are same as before link shut						
		verify IGMP group membership after link no shut is same as before link shut						
		Verify all DHCP Relay entries are same after link no shut						
		Verify VPC information after link no shut is same as before link shut						
1.5.6	Multicast with Multipath		684	79	79	0	0	
1.5.6.1	RPF Failure/Recovery between first leaf and elected RP		26	2	2	0	0	

1.5.6.1.1	DC37-103-LinkShut-Po12		13					
		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)						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify frames delta does not increase before link shut						
		Verify frames delta does not increase after link no shut						
1.5.6.1.2	DC37-103-LinkNoShut-Po12		13					
		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)						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP neighbors						
		Verify PIM neighbor status.						
		Verify frames delta does not increase before link shut						
		Verify frames delta does not increase after link no shut						
1.5.7	Reload and Power Cycle Switch		48	9	9	0	0	
1.5.7.1	Reload Spine		29	4	4	0	0	
1.5.7.1.1	DC37-1-Reload-DC37-1		14					

		Verify that all unicast/multicast traffic convergence is comparable to previous releases.						
		Verify STP port states during and after reload.						
		Verify FHRP peers status during and after reload.						
		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 that MEM and CPU Usage for Supervisors and line cards are comparable to previous releases.						
		Verify that there are no dead flows						
		Verify TB, error, crash						
		Verify interfaces in error						
		Verify any core dumps						
		Verify CDP/LLDP status during reload on the peers and after reload on the peers and DUT.						
		Verify frames delta does not increase.						
		Verify rx rate for all ixia ports are as expected (compared to baseline).						
		Verify traffic drop by checking Rx rate in all ixia ports after Shut						