Nexus Validation Test Phase 2.5

Introduction

This is an addendum to the NVT phase 2 report. Please reference the NVT phase 2 report for Network Topologies and feature descriptions. The NVT phase 2 report can be found here … http://www.cisco.com/en/US/docs/switches/datacenter/sw/nvt/phase2/NVT-Overview_Phase_2.pdf

Please see below for additional test results for NVT2.5 which includes new software releases.

Network Hardware and Software version Details

DC 1 Image Versions

	Model No.	NVT 2.5
N7K	**N7K-SUP1**	6.1.4a
N5K	**N5K-C5548UP-SUP**	5.2.1.N1.4
N3K	**N3K-C3048TP-1GE-SUP**	5.0.3.U5.1b
ASR9K	**A9K-RSP-4G**	4.2.3
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)*
	UCS-B22-M3	2.1(2a)*
	UCS-2208XP-FEX	2.1(2a)*
	UCS-6296UP-FI	2.1(2a)*

DC 2 Image Versions

	Model No.	NVT 2.5
N7K	**N7K-SUP2E**	6.1.4a
N5K	**N5K-C5548P -SUP**	5.2.1.N1.4
	N5K-C5548UP-SUP	5.2.1.N1.4
N3K	**N3K-C3548P-10G-SUP**	5.0.3.A1.2
ASR9K	**A9K-RSP-4G**	4.2.3
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
UCS	**UCS-5108**	2.1(2a)*
	UCS-B200-M2	2.1(2a)*
	UCS-B22-M3	2.1(2a)*
	UCS-2208XP-FEX	2.1(2a)*
	UCS-6296UP-FI	2.1(2a)*

Caveats for NVT 2.5

CSCuj43329

Symptom: copp core seen when doing ISSD from 6.1.4a to 5.2.9

Conditions: copp cores after switchover while doing ISSD on N7K from 6.1.4a to 5.2.9. Once the standby comes up with 5.2.9 on sup1, ISSD fails with copp core.

Workaround: None

Severity: severe

Platform Seen: N7000

Resolved Releases: 6.2(2)

Applicable Releases: 6.1(4a)

CSCul05316

Symptom: ISSU failed from 6.1.4 cco image to 6.1.4a - with pss error

Conditions: After ISSU from 6.1.4 image to 6.1.4a on sup2e, PSS errors are seen and spm cores as well. This happens once standby upgrade to 6.1.4a

Workaround: None

Severity: severe

Platform Seen: N7000

Resolved Releases: None

Applicable Releases: 6.1(4), 6.1(4a)

CSCuj31644

Symptom: SNMPd Core observed while running an SNMPwalk.

Conditions: When multiple snmp mibwalks are done on N7K, snmpd cores

Workaround: None

Severity: moderate

Platform Seen: N7000

Resolved Releases: 6.1(4a)

Applicable Releases:

CSCuj31512
Symptom: Fabricpath forwarding not working after link disruption & ISSD to 5.2.9
Conditions: On performing ISSD from 6.1.4a to 5.2.9 once the module connecting fabricpath downgrades, there will be packetloss to fabricpath leaves
Workaround: Reload the switch. After reloading, the fabricpath is in good state.
Severity: severe Platform Seen: N7000
Resolved Releases: None
Applicable Releases: 6.1(4a)

CSCuj12664
Symptom: 'vpc' cored on N7k vdc, when loading 6.1.4a
Conditions: On a sup1 vpc setup when you load 6.1.4a, vpc core is seen.
Workaround: None
Severity: severe
Platform Seen: N7000
Resolved Releases: 6.1(4a)
Applicable Releases:

CSCuj08139
Symptom: pixm cored on 6.1.4a vpc peer after ISSD
Conditions: pixm core is seen when you do ISSD from 6.1.4a to 5.2.9. After switchover, once the lincecards and fex's start downgrading pixm cores
Workaround: None
Severity: severe
Platform Seen: N7000
Resolved Releases: 6.1(4a)
Applicable Releases:

CSCuj05629
Symptom: Supervisor failed diag due to: DIAGCLIENT-2-EEM_ACTION_HM_SHUTDOWN
Conditions: After running 6.1.4a for sometime, diag fails with the error: DIAGCLIENT-2-EEM_ACTION_HM_SHUTDOWN

Workaround: None
Severity: severe
Platform Seen: N7000
Resolved Releases: 6.1(4a)
Applicable Releases:

CSCui92399
Symptom: fex traffic duplicating after SSO on n7k running 6.1.4a
Conditions: If fex and its downstream ports are connected in vpc, then a switchover will cause fex traffic duplication
Workaround: Shut and no-shut the fex dowstream ports to stop the flooding
Severity: moderate
Platform Seen: N7000
Resolved Releases: 6.2(2)
Applicable Releases: 6.1(4a)

CSCui92396
Symptom: Console flooded with BFD messages (On a vdc on N7k)
Conditions: After loading 6.1.4a, the uplink L2 links were flooded with BFD messages untill the link was flapped: BFD-5-SESSION_ACTIVE_PARAMS_CHANGE
Workaround: None
Severity: moderate
Platform Seen: N7000
Resolved Releases: None
Applicable Releases: 6.1(4a)

CSCug44374
Symptom: Igmp report sync between 6.2(2) and 6.1(4) VPC does not work
Conditions: With one pair of VPC peers, one peer from has upgraded from 6.1.4 to 6.2.2 and the other is running 6.1.4. igmp report sync from 614 -> 622 is having issue. This happens on both VPC and VPC+.
Workaround: None
Severity: moderate

Platform Seen: N7000
Resolved Releases: 6.2(2)
Applicable Releases: 6.1(4), 6.1(4a)

CSCug41055
Symptom: SYSMGR-2-SERVICE_CRASHED: Service "msdp" (PID 26636) hasn't caught signal
Conditions: This crash was seen while unsuspending a suspended VDC
Workaround: None
Severity: severe
Platform Seen: N7000
Resolved Releases: 6.2(2)
Applicable Releases: 6.1(4), 6.1(4a)

DC1 test results

				NVT 2.5	
Heading		**Test Case**	**Pass/Fail Verification**	**Status**	**Bugs**
1. DC1 Setup		**DC1 Setup**			
1.1. Common Configuration		Common Configuration for all switches	Verify SSH works through the management network on a dedicated vrf Verify Tacacs+ (tacacs.interop.cisco.com) and primary/backup servers Verify NTP 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 on all relevant interfaces. Verify SSO/NSF and GR Verify CoPP function 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	pass	
1.2. Edge/Core to Public Network Setup					
1.2.1. DC1-Core-N7k-1		Setup interfaces from DC1-Core-N7k-1 to Public network [AS1-1,AS1-2]	BGP: Verify Ipv4 eBGP peering between DC1-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC1-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath.	pass	

			PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies. Acl: Verify ACL policies are properly programmed in hardware and are functioning as expected. QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established. For each feature enable label sharing and ensure it is actually deployed by checking the number of used TCAM entries (identify all the features that share labels). Verify bank chaining of the TCAM.		
1.2.2. DC1-Core-N7k-2		Setup interfaces from DC1-Core-N7k-2 to Public network [AS1-1,AS1-2]	BGP: Verify IPv4/IPv6 eBGP peering between DC1-Core-n7k-2 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC1-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies. Acl: Verify ACL policies are properly programmed in hardware and are functioning as expected. QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established. For each feature enable label sharing and ensure it is actually deployed by checking the number of used TCAM entries (identify all the features that share labels). Verify bank chaining of the TCAM.	pass	
1.2.3. DC1-Core-ASR9k-3		Setup interfaces from DC1-Core-ASR9k-3 to Public network [AS1-1,AS1-2]	BGP: Verify IPv4/IPv6 eBGP peering between DC1-Core-ASR9k-3 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC1-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies. Acl: Verify ACL policies are functioning as expected.	not done	

			QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established.		
1.3. Core to Distribution Setup					
1.3.1. DC1-Core-N7k-1		Setup interfaces from DC1-Core-N7k-1 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache	pass	
1.3.2. DC1-Core-N7k-2		Setup interfaces from DC1-Core-N7k-2 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache	pass	
1.3.3. DC1-Core-ASR9k-3		Setup interfaces from DC1-Core-ASR9k-3 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache	not done	
1.4. Distribution to Core Setup					
1.4.1. DC1-Dist-N7k-101		Setup interfaces from Distribution N7k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. OTV: Verify OTV ISIS adjacencies are properly established and OTV routing table. Verify the primary AS is being used. On the primary AS, verify all edge devices show up in the unicast replication list using “show otv adjacency-server replication-list”.	pass	
1.4.2. DC1-Dist-N7k-102		Setup interfaces from Distribution N7k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. OTV: Verify OTV ISIS adjacencies are properly established and OTV routing table. Verify the primary AS is being used. On the primary AS, verify all edge devices show up in the unicast replication list using “show otv adjacency-server replication-list”.	pass	
1.4.3. Distribution Interop					
1.4.3.1. DC1-Dist-C6kE8-103-VSS		Setup interfaces from Distribution C6kE8 VSS to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering.	pass	

			PIM: Verify PIM peering.		
1.4.3.2. DC1-Dist-C6kE8-104		Setup interfaces from Distribution C6kE8 to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.3. DC1-Dist-C6kE7-105-VSS		Setup interfaces from Distribution C6kE7 VSS to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.4. DC1-Dist-C6kE7-106		Setup interfaces from Distribution C6kE7 to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.5. DC1-Dist-C4k-107		Setup interfaces from Distribution C4k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.5. Distribution to ToR Setup					
1.5.1. DC1-Dist-N7k-101		Setup interfaces from Distribution N7k to the ToR	vPC: Verify vPC peer-gateway, vPC peer-switch, vPC Object tracking, vPC auto recovery. Verify vPC peer status, vPC priority and consistency parameters. Check MAC/ARP/igmp snooping synchronization. OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP Ipv4/IPv6 peering between s5 and s6. Verify HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking. DHCP Relay Agent: Verify DHCP relay functionality. BOOTP: Verify BOOTP functionality. OTV: Verify OTV AS adjacencies state and verify VLAN load-balancing for each of theOTV edge devices. Verify remote MAC learning in the OTV MAC table.	pass	
1.5.1.1. ToR FEX vPC		Setup interface from DC1-Dist-N7k-101 to	Verify FEX association with configured port-channels and that the FEX devices are up.	pass	

		ToR FEX vPC			
1.5.1.2. ToR Layer 2 Switch		Setup interface from DC1-Dist-N7k-101 to ToR Layer 2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.1.3. ToR N5k vPC		Setup interface from DC1-Dist-N7k-101 to ToR N5k vPC	Verify vPC status and consistency parameters. Verify spanning tree status on all vlans.	pass	
1.5.1.4. ToR UCS Fabric Interconnect vPC		Setup interface from DC1-Dist-N7k-101 to ToR Fabric Interconnect vPC	Verify vPC status and consistency parameters		
1.5.2. DC1-Dist-N7k-102		Setup interfaces from Distribution N7k to the ToR	FabricPath: Verify FabricPath route and mac-table are built as expected. Verify IS-IS database. Verify multi-destination trees for unknown unicast, broadcast and multicast with root configured on the spine switches. Verify fabricpath load-balance works as expected OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP Ipv4/IPv6 peering between s51 & s52; s53 & s54. Verify HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch with G flag. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking. DHCP Relay Agent: Verify DHCP relay functionality. BOOTP: Verify BOOTP functionality. OTV: Verify OTV AS adjacencies state and verify VLAN load-balancing for each of the OTV edge devices. Verify remote MAC learning in the OTV MAC table.	pass	
1.5.2.1. ToR FEX		Setup interface from distribution DC1-Dist-N7k-102 to ToR FEX	Verify FEX association with configured port-channels and that the FEX devices are up.	pass	
1.5.2.2. ToR Layer 2 Switch		Setup interface from DC1-Dist-N7k-102 to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.2.3. ToR N5k FabricPath		Setup interface from DC1-Dist-N7k-102 to ToR N5k FabricPath	Verify FabricPath route and mac-table are built as expected. Verify the unknown unicast, broadcast and multicast multi-destination trees are built as expected. Verify fabricpath load-balance works as expected	pass	

			Verify IS-IS database, topology and route distribution.		
1.5.2.4. ToR UCS Fabric Interconnect vPC+		Setup interface from DC1-Dist-N7k-102 to ToR Fabric interconnect vPC+	Verify vPC+ status and consistency parameters.		
1.5.2.5. ToR Layer 2 Switch vPC+		Setup interface from DC1-Dist-N7k-102 to ToR L2 Switch vPC+	Verify vPC+ status and consistency parameters.	pass	
1.5.2.6. ToR N3k Layer 3		Setup interface from DC1-Dist-N7k-102 to ToR N3k Layer 3	Verify OSPF/OSPFv3 peering. Verify PIM peering.	pass	
1.5.3. Distribution Interop					
1.5.3.1. DC1-Dist-C6kE8-103-VSS		Setup interfaces from Distribution DC1-Dist-C6kE8-103-VSS to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. VSS: Verify VSS active/standby roles and VSL/MEC status. Verify Fast-redirect optimization IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP configuration. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.5.3.1.1. ToR Layer 2 Switch		Setup interface from DC1-Dist-C6kE8-103-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.1.2. ToR UCS Fabric Interconnect		Setup interface from DC1-Dist-C6kE8-103-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.2. DC1-Dist-C6kE8-104		Setup interfaces from Distribution C6k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP peering between s5 and s6.	pass	

			STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.		
1.5.3.2.1. ToR Layer 2 Switch		Setup interface from DC1-Dist-C6kE8-104 to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.2.2. ToR UCS Fabric Interconnect MEC		Setup interface from DC1-Dist-C6k-006-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.2.3. ToR N5k MEC		Setup interface from DC1-Dist-C6kE8-104 to ToR N5k MEC	Verify spanning tree status on all vlans.	pass	
1.5.3.2.4. ToR N3k Layer 3		Setup interface from DC1-Dist-C6kE8-104 to ToR N3k Layer 3	Verify OSPF/OSPFv3. Verify PIM peering.	pass	
1.5.3.3. DC1-Dist-C6kE7-105-VSS		Setup interfaces from Distribution C6k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. VSS: Verify VSS active/standby roles and VSL/MEC status. Verify Fast-redirect optimization IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP configuration. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.5.3.3.1. ToR Layer 2 Switch		Setup interface from DC1-Dist-C6kE7-105-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.3.2. ToR UCS Fabric Interconnect		Setup interface from DC1-Dist-C6kE7-105-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.4. DC1-Dist-C6kE7-106		Setup interfaces from Distribution C6k to the	OSPF: Verify OSPFv2/OSPFv3 peering.	pass	

		ToR	PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP peering between s5 and s6. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.		
1.5.3.4.1. ToR Layer 2 Switch		Setup interface from DC1-Dist-C6kE8-008-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.4.2. ToR UCS Fabric Interconnect MEC		Setup interface from DC1-Dist-C6kE7-106 to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.4.3. ToR N5k MEC		Setup interface from DC1-Dist-C6kE7-106 to ToR N5k MEC	Verify spanning tree status on all vlans.	pass	
1.5.3.5. DC1-Dist-C4k-107		Setup interfaces from Distribution C4k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP peering between s5 and s6. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking.	pass	

			DHCP Relay Agent: Verify DHCP relay functionality. BOOTP Relay Agent: Verify BOOTP relay functionality.		
1.5.3.5.1. ToR UCS Fabric Interconnect		Setup interface from DC1-Dist-C4k-107 to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.6. ToR to Distribution Setup					
1.6.1. ToR Layer 2 Switch vPC					
1.6.1.1. DC1-Dist-N7k-101		Setup vPC interface from ToR Layer 2 Switch to DC1-Dist-N7k-101	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.	pass	
1.6.2. ToR Layer 2 Switch vPC+					
1.6.2.1. DC1-Dist-N7k-102		Setup interfaces from ToR Layer 2 Switch vPC+ to the DC1-Dist-N7k-102	IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.6.3. ToR N3k Layer 3					
1.6.3.1. DC1-Dist-N7k-102		Setup interface from ToR N3k Layer 3 to DC1-Dist-N7k-102	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. IGMP/MLD Snooping: Verify IGMP/MLD Snooping ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.6.3.2. DC1-Dist-C6kE8-104		Setup interface from ToR N3k Layer 3 to DC1-Dist-C6kE8-104	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. IGMP/MLD Snooping: Verify IGMP/MLD Snooping ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines.	pass	

			ACL: 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.		
1.6.4. ToR N5k vPC					
1.6.4.1. DC1-Dist-N7k-101		Setup interface from ToR N5k vPC Switch to DC1-Dist-N7k-101	vPC: Verify vPC peer status and consistency parameters. Check MAC/ARP/igmp snooping synchronization. IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.6.5. ToR N5k FabricPath					
1.6.5.1. DC1-Dist-N7k-102		Setup interfaces from ToR N5k FabricPath to the DC1-Dist-N7k-102	FabricPath: Verify FabricPath route and mac-table are built as expected. Verify IS-IS database. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected HSRP: Verify HSRP MAC address is programmed in the mac table IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.7. ToR to Hosts Setup					
1.7.1. FEX					
1.7.1.1. End Host		Setup interface from FEX to End Host (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.1.2. End Host vPC		Setup interface from FEX to End Host vPC (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.1.3. UCS Fabric Interconnect		Setup interface from FEX to UCS Fabric Interconnect	Verify spanning tree status (edge) on all vlans for the host ports.		

			Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.1.4. UCS Fabric Interconnect vPC		Setup interface from FEX to UCS Fabric Interconnect vPC	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.1.5. UCS Fabric Interconnect vPC+		Setup interface from FEX to UCS Fabric Interconnect vPC+	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.2. ToR Layer 2 Switch					
1.7.2.1. End Host		Setup interface from ToR Layer 2 Switch to End Host (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.2.2. UCS Fabric Interconnect		Setup interface from ToR Layer 2 Switch to UCS Fabric Interconnect	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.3. ToR N3k Layer 3					
1.7.3.1. End Host		Setup interface from ToR N3k Layer 3 Switch to End Host (traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.4. ToR N5k vPC					
1.7.4.1. FEX vPC		Setup interface from ToR N5k FEX to End Host vPC (traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.4.1. UCS Fabric Interconnect vPC		Setup interface from ToR N5k vPC to UCS Fabric Interconnect vPC	Verify spanning tree status on all vlans. Verify mac table is populated correctly.		

			Verify IGMP/MLD snooping.		
1.7.5. ToR N5k Fabricpath Leaf					
1.7.5.1. UCS Fabric Interconnect vPC+		Setup interface from ToR N5k FP to UCS Fabric Interconnect vPC+	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.2. End Host vPC+		Setup interface from ToR N5k FP to End Host vPC+ (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.3. End Host		Setup interface from ToR N5k FP to End Host (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.4. ToR L2 switch		Setup interface from ToR N5k FP to ToR L2 switch	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.5.5. ToR L2 switch vPC+		Setup interface from ToR N5k FP to ToR L2 switch vPC+	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.5.6. FEX vPC+		Setup interface from N5k FP ToR FEX vPC+ to End Hosts (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.8. UCS Setup					
1.8.1. Fabric Interconnect					
1.8.1.1. DC1-Dist-N7k-101					
1.8.1.1.1. UCS to N7K FEX	1.8.1.1.1.1	Setup for UCS 6296UP FI to FEX	Verify the two FI's are in a cluster. Verify FI end host mode configuration.	pass	

				Verify uplink port-channels towards FEX. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and pinning.		
1.8.1.1.2.	UCS to N7K VPC	1.8.1.1.2.1	Setup for UCS 6296UP FI to FEX	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards ToR FEX. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.1.3. Switch	UCS to Layer 2	1.8.1.1.3.1	Setup for UCS 6296UP FI to Layer 2 Switch	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards layer 2 switch. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and pinning.	pass	
1.8.1.1.4.	UCS to N5k VPC	1.8.1.1.4.1	Setup for UCS 6248UP FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.1.5. VPC	UCS to N7K FEX	1.8.1.1.5.1	Setup for UCS 6248UP FI to N7K FEX VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N7k VPC. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.1.6. VPC	UCS to N5K FEX	1.8.1.1.6.1	Setup for UCS 6296UP FI to N5K FEX VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration.	pass	

			Verify uplink port-channels towards N7k VPC. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.2. DC1-Dist-N7k-102					
1.8.1.2.1. UCS to N7K FabricPath VPC+	1.8.1.2.1.1	Setup for UCS 62xx FI to N7k FabricPath VPC+	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N7k VPC+. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.2.2. UCS to Layer 2 Switch	1.8.1.2.2.1	Setup for UCS 6248UP FI to Layer 2 Switch	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards the layer 2 switch. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.2.3. UCS to N5K FabricPath VPC+	1.8.1.2.3.1	Setup for UCS 6248UP/6296UP FI to N5k VPC+	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC+. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.2.4. UCS to N5K FEX FabricPath VPC+	1.8.1.2.4.1	Setup for UCS 6296UP FI to N5k FEX VPC+	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC+. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.	pass	
1.8.1.3. DC1-Dist-C6kE8-103-VSS					

1.8.1.3.1. UCS to C6kE8 VSS	1.8.1.3.1.1	Setup for UCS 6248UP FI to C6kE8 VSS	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.4. DC1-Dist-C6kE8-104 Standalone					
1.8.1.4.1. UCS to C6kE8 Standalone	1.8.1.4.1.1	Setup for UCS 62xx FI to C6kE8 Standalone	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.4.2. UCS to N5k VPC	1.8.1.4.2.1	Setup for UCS 62xx FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.5. DC1-Dist-C6kE7-105-VSS					
1.8.1.5.1. UCS to C6kE7 VSS	1.8.1.5.1.1	Setup for UCS 62xx FI to C6kE7 VSS	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.6. DC1-Dist-C6kE7-106 Standalone					
1.8.1.6.1. UCS to C6kE7 Standalone	1.8.1.6.1.1	Setup for UCS 6248UP FI to C6kE7 Standalone	Verify the two FI's are in a cluster. Verify FI end host mode configuration.		

			Verify uplink port-channels towards C6k. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.6.2. UCS to N5k VPC	1.8.1.6.2.1	Setup for UCS 62xx FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.7. DC1-Dist-C4k-107					
1.8.1.7.1. UCS to C4k	1.8.1.7.1.1	Setup for UCS 62xx FI to C4k	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C4k. Verify static pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.2. UCS Setup					
1.8.2.1 UCSM Initial Configuration	1.8.2.1.1	Setup network parameters for the FI cluster.	Verify that the primary FI's System Name, Admin Password, Management IP Address, Management IP Netmask, Default Gateway, DNS Server IP, and Domain Name are all properly configured. Verify that the secondary FI is configured to be in a cluster. Verify that the FI cluster is reachable. Verify successful user authentication.		
1.8.2.2. Hypervisor Installation	1.8.2.2.1	Setup ESXi 5.1 for server virtualization	Verify the ESXi 5.1 software installation on the B2xx Mx blade. Verify server's IP address can be pinged. Verify the configured VM's are up and running. Verify the distributed virtual switch is functional. Verify successful installation of operating systems. Verify traffic can be generated by the servers.		

1.8.2.3 VM Provisioning	1.8.2.3.1	Configure 5 virtual machines with 10 virtual network adapters [per each ESXi host].	Verify through the VM's CLI that the virtual network interfaces are up and associated to a vNIC on UCSM. Verify through the VM's CLI and vCenter 5.1, that the proper MAC addresses are associated to each of the VM's virtual network interfaces. Verify through the VM's CLI and vCenter 5.1, that the proper IP addresses are associated to each of the VM's virtual network interfaces via DHCP. Verify that the VMs are able to be accessed through SSH/Telnet. Verify that the VMs are reachable through the management interface. Verify that the VMs in the same subnet are reachable with one another.		
1.8.2.4. VM-FEX Installation	1.8.2.4.1	Setup VM-FEX	Verify through UCSM and vCenter that VM-FEX port profiles are properly mapped to the network adapters in VMDirectPath (High-Performance) mode.		
	1.8.2.4.2	Create datacenter in UCSM under VM tab	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters via DHCP. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	1.8.2.4.3	Create folder under datacenter in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter.		

	1.8.2.4.4	Create distributed virtual switch under folder in UCSM.	Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
1.8.2.4. Nexus 1000V Installation (Pod 106)	1.8.2.4.1	Setup Nexus 1000V	Verify that the Nexus 1000V is installed following the Java Installer procedure. Verify the network configurations for control, packet and management ports are configured with the proper vlans. Verify the configured VEMs and VSMs are up and running. Verify that the VSMs are properly configured in cluster-mode. Verify the n1kv distributed virtual switch is functional. Verify successful installation of operating systems. Verify traffic can be generated by the servers.		
	1.8.2.4.2	Configure uplink port profile on the Nexus 1000V	Verify through UCSM and vCenter that all port profiles are properly mapped to the network adapters in standard performance mode. Verify that vCenter executes the command properly and that it is reflecting the proper operation. Using the NXOS CLI, Verify that the operation is properly updated during the entire process. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify the configured VEMs and VSMs are up and running. Verify that the VSMs are properly configured in cluster-mode. Verify through UCSM and vCenter that all port profiles are properly mapped to the network adapters in standard performance mode.		

	1.8.2.4.3	Configure server-side port profiles on the Nexus 1000V	Fault monitoring verification on vCenter and NXOS CLI. Verify the expected behavior is properly executed following the best practice and user guide. Verify that vCenter executes the command properly and that it is reflecting the proper operation. Using the NXOS CLI, Verify that the operation is properly updated during the entire process. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify the configured VEMs and VSMs are up and running. Verify that the VSMs are properly configured in cluster-mode. Verify through UCSM and vCenter that all port profiles are properly mapped to the network adapters in standard performance mode.		
	1.8.2.4.4	Configure ESXi hosts to use the Cisco Nexus 1000V in vCenter 5.1	Fault monitoring verification on vCenter and NXOS CLI. Verify the expected behavior is properly executed following the best practice and user guide. Verify that vCenter executes the command properly and that it is reflecting the proper operation. Using the NXOS CLI, Verify that the operation is properly updated during the entire process. Verify that the configured VEMs and VSMs are up and running. Verify that the VSMs are properly configured in cluster-mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify through UCSM and vCenter that all port profiles are properly mapped to the network adapters in standard performance mode.		
	1.8.2.4.5	Associate ESXi hosts to use the Cisco Nexus 1000V in vCenter 5.1	Fault monitoring verification on vCenter and NXOS CLI. Verify the expected behavior is properly executed following the best practice and user guide. Verify that vCenter executes the command properly and that it is reflecting the proper operation. Using the NXOS CLI, Verify that the operation is properly updated during the entire process. Verify that the configured VEMs and VSMs are up and running. Verify that the VSMs are properly configured in cluster-mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify through UCSM and vCenter that all port profiles are properly mapped to the network adapters in standard performance mode. Fault monitoring verification on vCenter and NXOS CLI.		

			Verify the expected behavior is properly executed following the best practice and user guide.		
2. Network Disruptions Test Cases		**Network Disruptions Test Cases** Common checks for all network disruptions	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 UCS end host mode on FI and VM-FEX functionality.		
2. Network Disruptions Test Cases		**Network Disruptions Test Cases** Common checks for all network disruptions	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 UCS end host mode on FI and VM-FEX functionality. Verify UCS unicast/multicast traffic convergence		
2.1. L2 Link Failure/Recovery	2.1.1	L2 Port-channel Failure/Recovery between Distribution and ToR devices	Verify STP port states after link disruption are in the expected forwarding mode. Verify that the STP root does not change. Verify HSRP peers status does not change. Verify HSRP MAC in ARP table. Verify HSRP 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. On the aggregation switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. 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 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 link at the access switch 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.	pass	

			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, HSRP 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.		
	2.1.2	L2 port-channel member failure/recovery between Distribution and ToR devices	Verify port-channel load balancing and rbh assignment Verify that IGMP/MLD membership is not affected. 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 will be proportionate to number of members failed Multicast DR should not change. Verify that there is no protocol flapping.	pass	
	2.1.3	vPC leg failure/recovery between Distribution and ToR devices	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.	pass	
	2.1.4	vPC leg member failure/recovery between Distribution and ToR devices	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 there are 2 members on each vPC leg). 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.	pass	
	2.1.5	vPC peer-link failure/recovery between Distribution vPC peer switches	Verify that the operational secondary vPC peer will bring down the vPC member ports.	pass	CSCuj12664

			Verify that secondary peer will suspend the vpc vlan svi's. Verify that on recovery, the original states will be re-established.		
	2.1.6	vPC Peer-keepalive failure/recovery between Distribution vPC peer switches	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.	pass	
	2.1.7	vPC peer-link and keep-alive failure between Distribution vPC peer switches	If the keep-alive fails first followed by vPC peer link, then both vPC peers will become active. Verify dual-active scenario is encountered and with the peer-switch feature enabled, ensure the downstream device does not detect any spanning-tree misconfigurations. If the vPC peer-link fails first followed by the keep-alive link, the secondary should keep it's vPC member ports suspended. With vPC auto-recovery configured if the vPC peer-link fails first followed by the keep-alive link, the secondary will keep it's vPC member ports suspended for the duration of three consecutive keepalive failures. After the timer expires the member ports will be unsuspended and the system will change role to primary causing Dual-active scenario.	pass	
	2.1.8	vPC peer-link and keep-alive recovery from Dual-active between Distribution vPC peer switches	If keep-alive is recovered first, the active/secondary switch is determined by the role priority and the secondary switch will suspend vPC member ports and vpc svi's. If vpc peer link is recovered first followed by keep alive, the active/secondary switch is determined by the role priority and the system resumes.	pass	
	2.1.9	OTV VDC L2 Link Failure/Recovery	Verify traffic will recover after link recovery.		
2.2. L3 Link Failure/Recovery	2.2.1	L3 Port-channel Failure/Recovery between Edge and Public Network[Interop between N7K, C6K]	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 PIM neighbor status. Verify GRE Tunnel re-route due to transport disruption. Verify MTU fragmentation and reassembling at tunnel edge. Verify AutoRP mapping and boundaries.	pass	
	2.2.2	L3 Port-channel Failure/Recovery between Core and Distribution Layers[Interop between	Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link.	pass	

		N7K, ASR9k, C6K, C4k]				
			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.			
	2.2.3	L3 Port-channel Failure/Recovery between Distribution to ToR N3k Layer 3 [Interop between N7K & N3K; C6K &N3k]	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.	pass		

			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.		
	2.2.4	L3 port-channel member failure/recovery	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.	pass	

			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.		
	2.2.5	OTV VDC L3 Link Failure/Recovery	Verify traffic will recover after link recovery.		
2.3. Clear OSPF Neighbors/Process/Routes	2.3	Clear OSPF Neighbors/Process/Routes	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.		
2.4. Clear IPv4/IPv6 Multicast Routes	2.4	Clear IPv4/IPv6 Multicast Routes	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	pass	

			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.		
2.5. Reload and Power Cycle Switch	2.5.1	Reload and Power Cycle Edge/Core Switch	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.	pass	

			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.		
	2.5.2	Reload and Power Cycle Distribution Switch	Verify STP port states during and after reload. Verify HSRP 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 HSRP MAC in ARP table. Verify HSRP 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 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.	pass	

			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.		
	2.5.3	vPC peer switch VDC reload	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) 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)	pass	
2.6. Supervisor and Fabric HA	2.6.1	Supervisor HA on the edge/core layer	Compare startup/running configuration on Active Sup and Standby Sup before and after SSO. 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.	pass	CSCuj05629

			Verify BGP recursive lookup scenario.		
			Verify BGP reconvergence (control-plane & data-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 SSO.		
			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 SSO.		
			Verify BFD peer should not flap during and after SSO.		
			No traffic loss is expected.		
	2.6.2	Supervisor HA on the Distribution layer	Compare startup/running configuration on Active Sup and Standby Sup before and after SSO.	pass	
			Verify STP port states during and after SSO.		
			Verify HSRP peers status during and after SSO.		
			Verify CDP/LLDP status after SSO.		
			Verify ARP tables remain unaffected	fail	CSCui92399
			Verify HSRP MAC in ARP table.		
			Verify OTV ARP optimization/ARP caching works as expected after SSO.		
			Verify head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information.		
			Verify automated mapping of OTV sites multicast groups to transport multicast group.		
			Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding after SSO.		

			Verify IGMP snooping entries remain unaffected. 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 SSO. Verify SNMP traps are sent to SNMP collector. Verify traffic destined for CoPP classes is policed as expected. 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 SSO. 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 SSO. Verify BFD peer should not flap during and after SSO. Verify vPC peer status (role, peer link, keepalive link and consistency parameters) before and after SSO No traffic loss is expected.		
	2.6.3	Fabric Failover on the Edge/Core and Distribution Layers	Verify there is no impact to data plane and control plane on Fabric failover with no oversubscription		
2.7. Line Card OIR and Reset	2.7.1	L3 port-channel member failure/recovery, on OIR/reset line card	Verify hitless operation for non-affected ports Verify traffic load-balancing for distributed port-channels before and after OIR/reset Verify BGP/ IGP/ PIM reconvergence (control-plane & data plane) Verfiy BFD peer detection and client notifications Verify LACP interoperability for distributed port-channels Verify that CDP/LLDP does not lose peer information for non-affected line card. Verify that CDP/LLDP peer is removed for disrupted line cards.	pass	

	2.7.2	L2 port-channel member failure/recovery, on OIR/reset line card	Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. 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. All unicast and multicast traffic should re-converge with minimal packet loss. Verify traffic destined for CoPP classes is policed as expected. Verify port-channel load balancing and rbh assignment Verify LACP interoperability for distributed port-channels Verify STP port states after OIR/reset are in the expected forwarding mode. Verify HSRP peers status after OIR/reset. Verify HSRP MAC in ARP table. Verify IGMP/MLD snooping entries are deleted for the links of affected line card and re-learnt correctly on the alternative link after query from the IGMP snooping router. Verify that IGMP/MLD membership is not affected.	pass		
	2.7.3	vPC leg failure/recovery, on OIR/reset line card	Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. 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. The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping. 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.	pass		

	2.7.4	vPC leg member failure/recovery on OIR/reset line card	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 there are 2 members on each vPC leg). 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.	pass	
	2.7.5	vPC peer-link failure/recovery on OIR/reset line card	Verify that the operational secondary vPC peer will bring down the vPC member ports. Verify that secondary peer will suspend the vpc vlan svi's. Verify that on recovery, the original states will be re-established.	pass	CSCuj08139
	2.7.6	vPC Peer-keepalive failure/recovery on OIR/reset line card	There are 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.	pass	
	2.7.7	vPC peer-link and peer-keepalive failure on OIR/reset line card	If the keep-alive fails first followed by vPC peer link, then both vPC peers will become active. Verify dual-active scenario is encountered and with the peer-switch feature enabled, ensure the downstream device does not detect any spanning-tree misconfigurations. If the vPC peer-link fails first followed by the keep-alive link, the secondary should keep it's vPC member ports suspended. With vPC auto-recovery configured if the vPC peer-link fails first followed by the keep-alive link, the secondary will keep it's vPC member ports suspended for the duration of three consecutive keepalive failures. After the timer expires the member ports will be unsuspended and the system will change role to primary causing Dual-active scenario.	pass	
	2.7.8	vPC peer-link and peer-keepalive recovery on OIR/reset line card	If keep-alive is recovered first, the active/secondary switch is determined by the role priority and the secondary switch will suspend vPC member ports and vpc svi's. If vpc peer link is recovered first followed by keep alive, the active/secondary switch is determined by the role priority and the system resumes.	pass	
2.8. ISSU/ISSD	2.8.1	ISSU/ISSD	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 HSRP peers status during and after ISSU/ISSD. Verify CDP/LLDP status after ISSU/ISSD.	pass	CSCuj43329

			Verify HSRP MAC in ARP table. Verify HSRP 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 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.		
2.9. Configuration Change	2.9.1 2.9.2 2.9.3 2.9.4 2.9.5 2.9.6 2.9.7	Perform VPC Vlan add and delete Perform VPC SVI add and delete Perform Non-VPC Vlan add and delete Perform Non-VPC SVI add and delete Remove VDC and add it back Enable/Disable IGMP snooping Perform HSRP active/standby switchover by changing priority	Verify STP port states after each change are in the expected forwarding mode. Verify HSRP peers status after each change. Verify the L2 forwarding table should be updated correctly after each change. Verify HSRP MAC in ARP table. 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 and re-learnt correctly upon each disruption. DHCP relay configured on the spine switches should remain unaffected after each change. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP after each change. All unicast and multicast traffic should re-converge with expected packet loss. Verify SNMP traps are sent to SNMP collector. Verify that all unicast/multicast traffic convergence.		
2.10.FabricPath – Network disruptions					
2.10.1. FabricPath – Link Failure/Recovery	2.10.1.1	FabricPath - Core Link Failure/Recovery	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify HSRP peers status does not change. Verify HSRP MAC in ARP table. Verify HSRP 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 SNMP traps are sent to SNMP collector. 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.	pass	

			On the aggregation switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. 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 link 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 SPAN is mirroring packets correctly. 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, HSRP 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 traffic destined for CoPP classes is policed as expected.		
	2.10.1.2	Fabricpath - Core Link member failure/recovery	Verify port-channel load balancing and RBH assignment. Verify IS-IS database, topology and route distribution for metric change. Verify that IGMP/MLD membership is not affected. Verify that IGMP snooping entries change based on multi-destination tree topology change. The maximum traffic disruption for unicast/multicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping.	pass	
	2.10.1.3	Fabricpath - vPC+ leg failure/recovery	The maximum traffic disruption for unicast will be half for both upstream and downstream traffic or no loss. 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.	pass	
	2.10.1.4	Fabricpath - vPC+ leg member failure/recovery	The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic.	pass	

			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 there are 2 members on each vPC+ leg). 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.		
	2.10.1.5	Fabricpath - vPC+ peer-link failure/recovery (spine/leaf)	Verify that the operational secondary vPC+ peer will bring down the vPC+ member ports. Verify that secondary peer will not suspend the vPC+ vlan SVI's if "*dual-active exclude vlans*" is configured Verify on recovery that the operational secondary vPC+ peer will bring up the vPC+ member ports after the configured "*delay restore*" timer	pass	
	2.10.1.6	Fabricpath - vPC+ Peer-keepalive failure/recovery	There are no expected effects; both vPC+ peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.	pass	
	2.10.1.7	Fabricpath - vPC+ peer-link and Peer-keepalive failure/recovery	When the keep-alive fails first followed by vPC+ peer link, the peers should continue to see each other through fabricpath network. The effect should be same as just peer-link failure. The recovery should be same as the peer-link recovery.	pass	
2.10.2. FabricPath – Reload	2.10.2.1	FabricPath - Spine Node failure/recovery	Verify Fabricpath multi-destination trees reconverge after root change on node failure. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding. Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on the other spine routers Verify IGMP/MLD snooping entries are deleted for the affected link at the access switch 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 other spine routers. Verify SPAN is mirroring packets correctly. Verify SNMP traps are sent to SNMP collector. DHCP relay configured on the aggregation switches should remain unaffected.	pass	

			Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. All unicast and multicast traffic should re-converge with minimal packet loss. Verify traffic destined for CoPP classes is policed as expected. Verify that the MAC table, FP ISIS route table, ARP table, IP routing table, IGMP membership table, IGMP snooping table, Multicast routing table return to original state on recovery Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on recovery		
	2.10.2.2	FabricPath - Leaf Node failure/recovery	Verify Fabricpath multi-destination trees reconverge after leaf node failure. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify HSRP peers status does not change when CE or leaf switches are reloaded. Verify IGMP/MLD snooping entries are deleted for the affected link at the access switch 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 spine routers. Verify that the MAC table, FP ISIS route table, IGMP snooping table return to original state on recovery Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on recovery	pass	
2.10.3. FabricPath – Supervisor and Fabric HA	2.10.3.1	FabricPath – Supervisor HA on the spine nodes	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Compare startup/running configuration on Active Sup and Standby Sup before and after SSO. Verify STP port states during and after SSO. Verify HSRP peers status during and after SSO. Verify CDP/LLDP status after SSO. Verify HSRP MAC in ARP table. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding after SSO.	pass	

			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 SSO. Verify SNMP traps are sent to SNMP collector. Verify traffic destined for CoPP classes is policed as expected. 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 SSO. 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 SSO. Verify BFD peer should not flap during and after SSO. Verify vPC+ peer status (role, peer link, keepalive link and consistency parameters) before and after SSO No traffic loss is expected.		
	2.10.3.2	FabricPath - Fabric Failover on spine nodes	Verify there is no impact to data plane and control plane on Fabric failover with no oversubscription	pass	
2.10.4. FabricPath – Line card OIR and Reset	2.10.4.1	FabricPath – Line card OIR and Reset on spine nodes	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify hitless operation for non-affected ports Verify traffic load-balancing for distributed port-channels before and after OIR/reset Verfiy BFD peer detection and client notifications	pass	

			Verify LACP interoperability for distributed port-channels Verify STP port states after OIR/reset are in the expected forwarding mode. Verify HSRP peers status after OIR/reset. Verify that CDP/LLDP does not lose peer information for non-affected line card. Verify that CDP/LLDP peer is removed for disrupted line cards. Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. Verify HSRP MAC in ARP table. 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 links of affected line card and re-learnt correctly on the alternative link after query from the IGMP snooping router. Verify SPAN is mirroring packets correctly. Verify SNMP traps are sent to SNMP collector. All unicast and multicast traffic should re-converge with minimal packet loss. Verify traffic destined for CoPP classes is policed as expected.		
	2.10.4.2	FabricPath – FP core port-channel member failure/recovery, on OIR/reset line card	Verify port-channel load balancing and rbh assignment Verify that IGMP/MLD membership is not affected. The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping.	pass	
	2.10.4.3	FabricPath – vPC+ leg failure/recovery on OIR/reset line card	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.	pass	
	2.10.4.4	FabricPath – vPC+ leg member failure/recovery on OIR/reset line card	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 there are 2 members on each vPC+ leg).	pass	

			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.		
	2.10.4.5	FabricPath – vPC+ peer-link failure/recovery on OIR/reset line card	Verify that the operational secondary vPC+ peer will bring down the vPC+ member ports. Verify that secondary peer will not suspend the vPC+ vlan SVI's if "*dual-active exclude vlans*" is configured Verify on recovery that the operational secondary vPC+ peer will bring up the vPC+ member ports after the configured "*delay restore*" timer	pass	
	2.10.4.6	FabricPath – vPC+ Peer-keepalive failure/recovery on OIR/reset line card	There are no expected effects; both vPC+ peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.	pass	
	2.10.4.7	Fabricpath - vPC+ peer-link and Peer-keepalive failure/recovery on OIR/reset line card	When the keep-alive fails first followed by vPC+ peer link, the peers should continue to see each other through fabricpath network. The effect should be same as just peer-link failure. The recovery should be same as the peer-link recovery.	pass	
2.10.5. FabricPath – ISSU/ISSD	2.10.5.1	FabricPath – ISSU/ISSD	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 FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify STP port states during and after ISSU/ISSD. Verify HSRP peers status during and after ISSU/ISSD. Verify CDP/LLDP status after ISSU/ISSD. Verify HSRP MAC in ARP table. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding after ISSU/ISSD.	pass	

			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. All unicast and multicast traffic should re-converge. Verify traffic destined for CoPP classes is policed as expected. 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.	fail	CSCuj31512
2.10.6. FabricPath – Move /Add/Change Hosts	2.10.6.1	FabricPath – MAC move	Verify ARP tables remain unaffected, MAC table shows mac move. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. 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 properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. All unicast and multicast traffic should re-converge with minimal packet loss. Verify SNMP traps are sent to SNMP collector.		
	2.10.6.2	FabricPath – End Hosts Add	Verify ARP and MAC tables add the new hosts. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines. Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP.		
	2.10.6.3	FabricPath – End Hosts Change	Verify ARP and MAC tables change as expected. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines.		

			Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. Monitor all unicast/multicast traffic convergence.		
2.10.7. FabricPath – Configuration Change	2.10.7.1 2.10.7.2 2.10.7.3 2.10.7.4 2.10.7.5 2.10.7.6	Perform FP Vlan add and delete Perform FP SVI add and delete Perform Non-FP Vlan add and delete Perform Non-FP SVI add and delete Perform FP MT root move by changing priority Enable/Disable IGMP snooping	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. 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 after each change. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding after each change. Verify that no flooding happens after traffic convergence after each change. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines after each change. Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches after each change. DHCP relay configured on the spine switches should remain unaffected after each change. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP after each change. 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 on all the affected FP switches. All unicast and multicast traffic should re-converge with minimal packet loss. Verify SNMP traps are sent to SNMP collector. Monitor all unicast/multicast traffic convergence.	pass	
2.11.OTV – Network Disruptions					
2.11.1. OTV – Reload	2.11.1.1	OTV – Reload	Verify HSRP isolation across OTV sites works as expected after reload/recovery. Verify OTV ARP optimization/ARP caching works as expected after reload/recovery. Verify unknown unicast traffic doesn't flood. Verify STP is blocked across OTV sites.		

			Verify the Secondary Adj. Server will take over after primary Adj. Server failover. 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 head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information. Verify automated mapping of OTV sites multicast groups to transport multicast group. Verify IGMP snooping entries are properly relearned on the affected OTV switches. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.11.2. OTV – Move/Add/Change Hosts	2.11.2.1	OTV – MAC move/Add/Change Hosts	Verify HSRP isolation across OTV sites works as expected. Verify OTV ARP optimization/ARP caching works as expected. Verify unknown unicast traffic doesn't flood. Verify the new hosts's macs are learnt across OTV sites. Verify STP is blocked across OTV sites. 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 head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information. Verify automated mapping of OTV sites multicast groups to transport multicast group. Verify IGMP snooping entries are properly relearned on the affected OTV switches. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.11.3. OTV – Configuration Change	2.11.3.1	Add and delete OTV VLAN	Verify HSRP isolation across OTV sites works as expected		
	2.11.3.2	Add and delete OVT SVI	Verify OTV ARP optimization/ARP caching/ARP suppression works as expected.		
	2.11.3.3	Enable and disable proxy ARP	Verify unknown unicast traffic doesn't flood.		
	2.11.3.4	Enable and disable suppression ARP	Verify STP is blocked across OTV sites.		
	2.11.3.5	Enable and disable igmp snooping	Verify new Adj. Server works as expected.		
	2.11.3.6	Add and delete overlay interface	Verify the new hosts's macs are learnt across OTV sites.		
	2.11.3.7	Dynamically changing Adj Server	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.		

	2.11.3.8	Add/remove/flush MAC entries	Verify head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information.		
	2.11.3.9	Add/remove/flush ARP entries	Verify automated mapping of OTV sites multicast groups to transport multicast group.		
	2.11.3.10	Add/remove/flush multicast group entries	Verify IGMP snooping entries are properly relearned on the affected OTV switches.		
	2.11.3.11	Add/remove/flush active multicast source entries	Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.12.UCS – Disruptions					
2.12.1. UCS – Link Failure/Recovery	2.12.1.1	UCS - Link Failure/Recovery Between FI and N7K: VPC	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify VM does not lose network connectivity. Measure traffic convergence for each disruption		
	2.12.1.2	FI Uplink port-channel member failure/recovery: 101-01 n7k vpc	Verify traffic recovery within the expected time frame. Verify that rehashing is performed according to the port-channel protocol (LACP) deployed. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify there is no mac address learning on FI uplink. Verify MAC learning on FI server links is not impacted.		
	2.12.1.3	FI Uplink port-channel failure/recovery: 101-01 n7k vpc	Verify traffic should switch to other FI and re-converge with expected packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on other FI server links.		
	2.12.1.4	FI to IOM port-channel member failure/recovery: 101-01 n7k vpc	Verify traffic recovery within the expected time frame. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify there is no mac address learning on FI uplink.		

	2.12.1.5	FI to IOM port-channel failure/recovery:	Verify mac learning on FI server links is not impacted. Verify traffic recovery within the expected time frame. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on other FI server links.		
	2.12.1.6	FI cluster link member failure/recovery: 101-01 n7k vpc	Verify traffic should have no impact. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links is not impacted.		
	2.12.1.7	FI to FI isolation/recovery: 101-01 n7k vpc	Verify traffic should re-converge after FI cluster link recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after FI cluster link recovery. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink after FI cluster link recovery. Verify mac learning on other FI server links after FI cluster link recovery.		
2.12.2. UCS – Fabric Interconnect Reload and Power Cycle	2.12.2.1	UCS – Fabric Interconnect Reload and Power Cycle: 101-01 n7k vpc	Verify traffic recovery within the expected time frame. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify that traffic flows accordingly through the uplink switches following the VPC model. Verify there is no mac address learning on other FI uplink. Verify mac learning on other FI server links. Verify FI uplink static pinning works as expected.		

			Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify VM does not lose network connectivity. Measure traffic convergence for each disruption		
2.12.3. UCS – IOM OIR	2.12.3.1	UCS – IOM OIR	Verify traffic recovery within the expected time frame. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify there is no mac address learning on other FI uplink. Verify mac learning on other FI server links. Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify VM does not lose network connectivity.		
2.12.4. UCS – Blade OIR	2.12.4.1	UCS – Blade OIR	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after blade recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify mac learning on FI server links after blade recovery. Verify that no flooding happens after traffic convergence after blade recovery.		

	2.12.4.2	Perform live blade OIR (same slot, same chassis)	Verify that IGMP snooping is working as expected after blade recovery. Verify when blade is re-inserted that hypervisor and VMs are restored. Remove live blade and re-insert into the same slot within the same chassis. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.4.3	Perform live blade OIR (different slot, same chassis)	Remove live blade and decomission from slot. Then re-insert the blade into a different slot within the same chassis, and associate the service profile to the blade. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink.		

	2.12.4.4	Perform maintenance blade oir (different slot, different chassis)	Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Gracefully shutdown VMs and blade. Dissassociate service profile from blade. Remove the blade and accept notifications. Insert the blade into a different slot in a different chassis, and associate the service profile to the blade. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink.		
	2.12.4.5	Perform a blade swap (B200 with B22) for a blade upgrade	Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the same HDDs are retained throughout the process. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved.		

	2.12.4.6	In a B-Series chassis perform a blade upgrade/downgrade (B22/B200)	Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the same HDDs are retained throughout the process. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved.		
	2.12.4.7	In a B-Series chassis perform a complete blade upgrade/downgrade (B22/B200)	Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved.		

			Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify that the HDD OIR in a RAID 1 Mirrored system does not impact the VMs.		
2.12.5. UCS – Chassis Reload and Power Cycle	2.12.5.1	UCS – Chassis Reload and Power Cycle	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify traffic should re-converge after chassis IOM and blade recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after chassis IOM and blade recovery. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after chassis IOM and blade recovery. Verify that no flooding happens after traffic convergence after chassis IOM and blade recovery. Verify that IGMP snooping is working as expected after chassis IOM and blade recovery. Verify VM network connectivity is restored.		
2.12.6. UCS – FI image and IOM Firmware Upgrade	2.12.6.1	UCS – FI image and IOM Firmware Upgrade	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify traffic should re-converge after IOM firmware upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after IOM firmware upgraded. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after IOM firmware upgraded.		

			Verify that no flooding happens after traffic convergence after IOM firmware upgraded. Verify that IGMP snooping is working as expected after IOM firmware upgraded. Verify VM network connectivity is restored.		
2.12.7. UCS – Blade adapter Firmware upgrade	2.12.7.1	UCS – Blade adapter Firmware upgrade	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify traffic should re-converge after blade adapter firmware upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade adapter firmware upgraded. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade adapter firmware upgraded. Verify that no flooding happens after traffic convergence after blade adapter firmware upgraded. Verify that IGMP snooping is working as expected after blade adapter firmware upgraded. Verify VM network connectivity is restored.		
2.12.8. UCS – Blade BIOS upgrade	2.12.8.1	UCS – Blade BIOS upgrade	Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify traffic should re-converge after blade BIOS upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade BIOS upgraded. Verify proper MAC address learning on both fabric interconnects and Nexus 7000 switches. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade BIOS upgraded. Verify that no flooding happens after traffic convergence after blade BIOS upgraded. Verify that IGMP snooping is working as expected after blade BIOS upgraded. Verify VM network connectivity is restored.		
2.12.9. UCS – VMotion for Blade Maintenance	2.12.9.1	Migrate live VM across different blades, same chassis, same FI pair (VM-FEX)	Verify that the VM's vNICs and port profiles are still associated and configured properly before and after the migration through monitoring the CLI. Verify that the MAC address of the migrated VM is learned on the destined Fabric Interconnect and the corresponding upstream switch throughout the migration.		

	2.12.9.2	Migrate live VM across different blades, different chassis, same FI pair (VM-FEX)	Verify that the VMs within the testbed remain pingable between one another during and after the migration. Verify that the VM's network interfaces remain pingable from our management network before and after the migration. Verify that the VM is still reachable through an SSH, or Telnet session. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that no faults are raised on either UCSM or vCenter during the operation. Verify that the VM migration is properly executed while following the best practices and user guide. Verify that the VM's vNICs and port profiles are still associated and configured properly before and after the migration through monitoring the CLI. Verify that the MAC address of the migrated VM is learned on the destined Fabric Interconnect and the corresponding upstream switch throughout the migration.		
	2.12.9.3	Migrate live VM across different blades, same chassis, same FI pair (VMWare vDS)	Verify that the VMs within the testbed remain pingable between one another during and after the migration. Verify that the VM's network interfaces remain pingable from our management network before and after the migration. Verify that the VM is still reachable through an SSH, or Telnet session. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that no faults are raised on either UCSM or vCenter during the operation. Verify that the VM migration is properly executed while following the best practices and user guide. Verify that the VM's vNICs and port profiles are still associated and configured properly before and after the migration through monitoring the CLI. Verify that the MAC address of the migrated VM is learned on the destined Fabric Interconnect and the corresponding upstream switch throughout the migration.		
	2.12.9.4	Migrate live VM across different blades,	Verify that the VMs within the testbed remain pingable between one another during and after the migration. Verify that the VM's network interfaces remain pingable from our management network before and after the migration. Verify that the VM is still reachable through an SSH, or Telnet session. Verify that no faults are raised on either UCSM or vCenter during the operation. Verify that the VM migration is properly executed while following the best practices and user guide. Verify that the VM's vNICs and port profiles are still associated and configured properly before		

		different chassis, same FI pair (VMWare vDS)	and after the migration through monitoring the CLI. Verify that the MAC address of the migrated VM is learned on the destined Fabric Interconnect and the corresponding upstream switch throughout the migration. Verify that the VMs within the testbed remain pingable between one another during and after the migration. Verify that the VM's network interfaces remain pingable from our management network before and after the migration. Verify that the VM is still reachable through an SSH, or Telnet session. Verify that no faults are raised on either UCSM or vCenter during the operation. Verify that the VM migration is properly executed while following the best practices and user guide.		
	2.12.9.5	Migrate live VM across different blades, different chassis, different FI pair (VMWare vDS)	Verify that the VM's vNICs and port profiles are still associated and configured properly before and after the migration through monitoring the CLI. Verify that the MAC address of the migrated VM is learned on the destined Fabric Interconnect and the corresponding upstream switch throughout the migration. Verify that the VMs within the testbed remain pingable between one another during and after the migration. Verify that the VM's network interfaces remain pingable from our management network before and after the migration. Verify that the VM is still reachable through an SSH, or Telnet session. Verify that no faults are raised on either UCSM or vCenter during the operation. Verify that the VM migration is properly executed while following the best practices and user guide.		
2.12.10. UCS – NIC Bonding	2.12.10.1	Configure Active / Standby nic bonding	Modify ifcfg-eth8 configuration file Modify ifcfg-eth9 configuration file Create ifcfg-bond0 configuration file Create Modprobe.conf file for mode1 active/standby nics Verify that the bonding is successful Perform an ifdown on eth8 which is the active nic Verify standby nic eth9 becomes active after failover. Perform an ifup on eth8 and verify it becomes standby Verify ping and ssh sessions are all active Verify FI uplink static pinning works as expected.		

			Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify there is no mac address learning on FI uplink.			
	2.12.10.2	Configure Adaptive Load Balancing nic bonding	Modify ifcfg-eth8 configuration file Modify ifcfg-eth9 configuration file Create ifcfg-bond0 configuration file Create Modprobe.conf file for mode6 (ALB) nics Verify that the bonding is successful Perform an ifdown on eth8 Verify traffic continues without loss as secondary nic continues to forward traffic. Perform ifup on eth8 and verify traffic continues to load balance between links. Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify there is no mac address learning on FI uplink.			
	2.12.10.3	Perform FI Failover from Fi-A to Fi-B	Verify ping and ssh sessions are all active login to FI CLI and enter local-mgmt and preform reload on FI-A verify that the FI recovers and there are no critical error messages verify that the vifs failover to FI-B and traffic resumes verify that the vifs resume on FI-A and traffic resumes Verify ping and ssh sessions are all active Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify there is no mac address learning on FI uplink.			
	2.12.10.4	Perform Network uplink failover	Shut the network uplink portchannel on the FI Verify that enm pinning fails verify that the vifs failover to FI-B and traffic resumes			

			No-Shut the network uplink portchannel on the FI verify that the vifs resume on FI-A and traffic resumes Verify ping and ssh sessions are all active Verify FI uplink static pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify there is no mac address learning on FI uplink.		
2.12.11. UCS – Port Profile Tests	2.12.11.1	Remove a port profile in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.11.2	Toggle port profile's I/O Performance mode	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.11.3	Create a profile client in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation.		

			Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.			
	2.12.11.4	Associate a port profile to a VM	Verify vCenter executes the command properly and that UCSM is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.			
	2.12.11.5	Remove associated port profile and profile client in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process.			

	2.12.11.6	Unassociate port profile from a VM	Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify vCenter executes the command properly and that UCSM is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process.		
	2.12.11.7	Remove unassociated port profile and profile client in UCSM	Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process.		
	2.12.11.8	Modify port profile and LAN pin group in UCSM	Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved.		

	2.12.11.9	Create duplicate port profile in UCSM	Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
2.12.12. UCS – VM-FEX Tests	2.12.12.1	Create duplicate associated distributed virtual switch (VM-FEX) from the same FI cluster in UCSM	Verify UCSM detects and reflects the proper duplication error. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the data plane interfaces are configured in VMDirectPath mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		

	2.12.12.2	Associate/Sync distributed virtual switch to ESXi hosts in vCenter	Verify vCenter executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.12.3	Remove associated distributed virtual switch in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.12.4	Create duplicate associated distributed virtual switch from a different FI cluster in UCSM	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode.		

	2.12.12.5	Remove duplicate associated distributed virtual switch from different FI-pair in UCSM	Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI and vCenter GUI. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
2.12.13. UCS – Server Clustering Tests	2.12.13.1	Convert pod to cluster setting in vCenter 5.1	Verify vSphere GUI executes the command properly and that it is reflecting the proper operation. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that vCenter 5.1 acknowledges the creation of the cluster and its components. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.13.2	Configure and associate a shared datastore for cluster High Availability in vCenter 5.1	Verify vSphere GUI executes the command properly and that it is reflecting the proper operation. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that vCenter 5.1 acknowledges the creation of the cluster and its components.		

	2.12.13.3	Enable VM Monitoring within the High Availability cluster	Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify vSphere GUI executes the command properly and that it is reflecting the proper operation. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify that vCenter 5.1 acknowledges the creation of the cluster and its components. Verify through UCSM and vCenter that VM-FEX port profiles for all necessary data plane traffic are properly mapped to the network adapters in VMDirectPath (High-Performance) mode. Verify through UCSM and vCenter that VM-FEX port profiles for management plane traffic are properly mapped to the network adapters in standard performance mode. Fault monitoring verification on vCenter. Verify the expected behavior is properly following the best practice and user guide.		
2.12.14. UCS – Service Profile Testing	2.12.14.1	From UCSM GUI perform server shutdown for a scheduled maintenance.	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.14.2	From UCSM GUI perform boot server to recover after a schedule maintenance	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters.		

	2.12.14.3	From UCSM GUI perform a blade reset to simulate a blade failure	Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters.		
	2.12.14.4	From UCSM GUI perform a server profile (SP) rename for management purposes	Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters.		
	2.12.14.5	From UCSM GUI perform a server profile (SP) clone for management purposes	Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		

	2.12.14.6	From UCSM GUI perform a server profile (SP) template creation for portability and usability purposes	Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.14.7	From UCSM GUI perform service profile (SP) dis-association for a blade maintanance	Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters. Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		
	2.12.14.8	From UCSM GUI perform a bind to a template for the reprovisioning of a newly inserted blade	Verify when blade is re-inserted that hypervisor and vm are properly restored. Verify UCSM executes the command properly and that vCenter is reflecting the operation. Verify syncing between UCSM GUI, vCenter GUI and KVM consoles. Using the CLI, verify that the vNICs, MAC, and IP addresses are properly associated on all of the VMs' network adapters.		

			Verify that the reachability of all the affected interfaces is properly restored after each disruption and the network convergence is achieved. Verify that the reachability of all the interfaces of non-affected VMs is preserved throughout the entire process. Verify FI uplink static pinning works as expected. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Fault monitoring verification on both UCSM and vCenter. Verify the expected behavior is properly following the best practice and user guide.		

DC2 test results

				NVT 2.5	
Heading		**Test Case**	**Pass/Fail Verification**	**Status**	**Bugs**
1. DC2 Setup		**DC2 Setup**			
1.1. Common Configuration		Common Configuration for all switches	Verify SSH works through the management network on a dedicated vrf 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 on all relevant interfaces. Verify SSO/NSF and GR Verify CoPP function 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	pass	CSCuj31644
1.2. Edge/Core to Public Network Setup					
1.2.1. DC2-Core-N7k-1		Setup interfaces from DC2-Core-N7k-1 to Public network [AS1-1,AS1-2]	BGP: Verify Ipv4 eBGP peering between DC2-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC2-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath.	pass	

			PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies. Acl: Verify ACL policies are properly programmed in hardware and are functioning as expected. QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established. For each feature enable label sharing and ensure it is actually deployed by checking the number of used TCAM entries (identify all the features that share labels). Verify bank chaining of the TCAM.		
1.2.2. DC2-Core-N7k-2		Setup interfaces from DC2-Core-N7k-2 to Public network [AS1-1,AS1-2]	BGP: Verify IPv4/IPv6 eBGP peering between DC2-Core-n7k-2 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC2-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies. Acl: Verify ACL policies are properly programmed in hardware and are functioning as expected. QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established. For each feature enable label sharing and ensure it is actually deployed by checking the number of used TCAM entries (identify all the features that share labels). Verify bank chaining of the TCAM.	pass	
1.2.3. DC2-Core-ASR9k-3		Setup interfaces from DC2-Core-ASR9k-3 to Public network [AS1-1,AS1-2]	BGP: Verify IPv4/IPv6 eBGP peering between DC2-Core-ASR9k-3 and AS1-1,AS1-2. Verify eBGP multipath. BGP: Verify Ipv6 eBGP peering between DC2-Core-n7k-1 and AS1-1,AS1-2. Verify eBGP multipath. PIM: Verify PIM peering. Redistribute: Verify routes are redistributed according to configured policies.		

			Acl: Verify ACL policies are functioning as expected. QoS: Verify QoS marking and policing. NAT: Verify NAT translation is properly handled at uplink interfaces including the GRE tunnel EP. NDE: Verify Netflow enabled interfaces monitor and export flow entries to external flow collector. GRE: Ensure GRE tunnels are up and all configured protocol peerings are fully established.		
1.3. Core to Distribution Setup					
1.3.1. DC2-Core-N7k-1		Setup interfaces from DC2-Core-N7k-1 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache	pass	
1.3.2. DC2-Core-N7k-2		Setup interfaces from DC2-Core-N7k-2 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache	pass	
1.3.3. DC2-Core-ASR9k-3		Setup interfaces from DC2-Core-ASR9k-3 to Distribution blocks	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache		
1.4. Distribution to Core Setup					
1.4.1. DC2-Dist-N7k-101		Setup interfaces from Distribution N7k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. OTV: Verify OTV ISIS adjacencies are properly established and OTV routing table. Verify the primary AS is being used. On the primary AS, verify all edge devices show up in the unicast replication list using “show otv adjacency-server replication-list”.	pass	
1.4.2. DC2-Dist-N7k-102		Setup interfaces from Distribution N7k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. OTV: Verify OTV ISIS adjacencies are properly established and OTV routing table. Verify the primary AS is being used. On the primary AS, verify all edge devices show up in the unicast replication list using “show otv adjacency-server replication-list”.	pass	
1.4.3. Distribution Interop					

1.4.3.1. DC2-Dist-C6kE8-103-VSS		Setup interfaces from Distribution C6kE8 VSS to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.2. DC2-Dist-C6kE8-104		Setup interfaces from Distribution C6kE8 to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.3. DC2-Dist-C6kE7-105-VSS		Setup interfaces from Distribution C6kE7 VSS to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.4. DC2-Dist-C6kE7-106		Setup interfaces from Distribution C6kE7 to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.4.3.5. DC2-Dist-C4k-107		Setup interfaces from Distribution C4k to the core switches	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering.	pass	
1.5. Distribution to ToR Setup					
1.5.1. DC2-Dist-N7k-101		Setup interfaces from Distribution N7k to the ToR	vPC: Verify vPC peer-gateway, vPC peer-switch, vPC Object tracking, vPC auto recovery. Verify vPC peer status, vPC priority and consistency parameters. Check MAC/ARP/igmp snooping synchronization. OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP Ipv4/IPv6 peering between s5 and s6. Verify HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking. DHCP Relay Agent: Verify DHCP relay functionality. BOOTP: Verify BOOTP functionality. OTV: Verify OTV AS adjacencies state and verify VLAN load-balancing for each of theOTV	pass	

			edge devices. Verify remote MAC learning in the OTV MAC table.		
1.5.1.1. ToR FEX vPC		Setup interface from DC2-Dist-N7k-101 to ToR FEX vPC	Verify FEX association with configured port-channels and that the FEX devices are up.	pass	
1.5.1.2. ToR Layer 2 Switch		Setup interface from DC2-Dist-N7k-101 to ToR Layer 2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.1.3. ToR N5k vPC		Setup interface from DC2-Dist-N7k-101 to ToR N5k vPC	Verify vPC status and consistency parameters. Verify spanning tree status on all vlans.	pass	
1.5.1.4. ToR UCS Fabric Interconnect vPC		Setup interface from DC2-Dist-N7k-101 to ToR Fabric Interconnect vPC	Verify vPC status and consistency parameters		
1.5.2. DC2-Dist-N7k-102		Setup interfaces from Distribution N7k to the ToR	FabricPath: Verify FabricPath route and mac-table are built as expected. Verify IS-IS database. Verify multi-destination trees for unknown unicast, broadcast and multicast with root configured on the spine switches. Verify fabricpath load-balance works as expected OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP Ipv4/IPv6 peering between s51 & s52; s53 & s54. Verify HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch with G flag. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking. DHCP Relay Agent: Verify DHCP relay functionality. BOOTP: Verify BOOTP functionality. OTV: Verify OTV AS adjacencies state and verify VLAN load-balancing for each of the OTV edge devices. Verify remote MAC learning in the OTV MAC table.	pass	
1.5.2.1. ToR FEX		Setup interface from distribution DC2-Dist-N7k-102 to ToR FEX	Verify FEX association with configured port-channels and that the FEX devices are up.	pass	
1.5.2.2. ToR Layer 2 Switch		Setup interface from DC2-Dist-N7k-102 to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.2.3. ToR N5k FabricPath		Setup interface from DC2-Dist-N7k-102 to ToR N5k FabricPath	Verify FabricPath route and mac-table are built as expected. Verify the unknown unicast, broadcast and multicast multi-destination trees are built as expected.	pass	

			Verify fabricpath load-balance works as expected Verify IS-IS database, topology and route distribution.		
1.5.2.4. ToR UCS Fabric Interconnect vPC+		Setup interface from DC2-Dist-N7k-102 to ToR Fabric interconnect vPC+	Verify vPC+ status and consistency parameters.	pass	
1.5.2.5. ToR Layer 2 Switch vPC+		Setup interface from DC2-Dist-N7k-102 to ToR L2 Switch vPC+	Verify vPC+ status and consistency parameters.	pass	
1.5.2.6. ToR N3k Layer 3		Setup interface from DC2-Dist-N7k-102 to ToR N3k Layer 3	Verify OSPF/OSPFv3 peering. Verify PIM peering.		
1.5.3. Distribution Interop					
1.5.3.1. DC2-Dist-C6kE8-103-VSS		Setup interfaces from Distribution DC2-Dist-C6kE8-103-VSS to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. VSS: Verify VSS active/standby roles and VSL/MEC status. Verify Fast-redirect optimization IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP configuration. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.5.3.1.1. ToR Layer 2 Switch		Setup interface from DC2-Dist-C6kE8-103-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.1.2. ToR UCS Fabric Interconnect		Setup interface from DC2-Dist-C6kE8-103-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.2. DC2-Dist-C6kE8-104		Setup interfaces from Distribution C6k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping	pass	

			HSRP: Verify HSRP peering between s5 and s6. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.		
1.5.3.2.1. ToR Layer 2 Switch		Setup interface from DC2-Dist-C6kE8-104 to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.2.2. ToR UCS Fabric Interconnect MEC		Setup interface from DC2-Dist-C6k-006-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.2.3. ToR N5k MEC		Setup interface from DC2-Dist-C6kE8-104 to ToR N5k MEC	Verify spanning tree status on all vlans.	pass	
1.5.3.2.4. ToR N3k Layer 3		Setup interface from DC2-Dist-C6kE8-104 to ToR N3k Layer 3	Verify OSPF/OSPFv3. Verify PIM peering.	pass	
1.5.3.3. DC2-Dist-C6kE7-105-VSS		Setup interfaces from Distribution C6k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. VSS: Verify VSS active/standby roles and VSL/MEC status. Verify Fast-redirect optimization IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP configuration. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.5.3.3.1. ToR Layer 2 Switch		Setup interface from DC2-Dist-C6kE7-105-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.3.2. ToR UCS Fabric Interconnect		Setup interface from DC2-Dist-C6kE7-105-VSS to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		

1.5.3.4. DC2-Dist-C6kE7-106		Setup interfaces from Distribution C6k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP peering between s5 and s6. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.5.3.4.1. ToR Layer 2 Switch		Setup interface from DC2-Dist-C6kE8-008-VSS to ToR L2 Switch	Verify spanning tree status on all vlans.	pass	
1.5.3.4.2. ToR UCS Fabric Interconnect MEC		Setup interface from DC2-Dist-C6kE7-106 to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.5.3.4.3. ToR N5k MEC		Setup interface from DC2-Dist-C6kE7-106 to ToR N5k MEC	Verify spanning tree status on all vlans.	pass	
1.5.3.5. DC2-Dist-C4k-107		Setup interfaces from Distribution C4k to the ToR	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. MSDP: Verify MSDP peering and SA-cache PIM Snooping: Verify PIM snooping. IGMP/MLD Snooping: Verify IGMP/MLD Snooping HSRP: Verify HSRP peering between s5 and s6. STP: Verify RSTP parameters and port status. ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: Verify that all the policies are properly programmed in hardware. QoS: Verify QoS marking.	pass	

			DHCP Relay Agent: Verify DHCP relay functionality. BOOTP Relay Agent: Verify BOOTP relay functionality.		
1.5.3.5.1. ToR UCS Fabric Interconnect		Setup interface from DC2-Dist-C4k-107 to ToR Fabric Interconnect	Verify spanning tree status on all vlans.		
1.6. ToR to Distribution Setup					
1.6.1. ToR Layer 2 Switch vPC					
1.6.1.1. DC2-Dist-N7k-101		Setup vPC interface from ToR Layer 2 Switch to DC2-Dist-N7k-101	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.	pass	
1.6.2. ToR Layer 2 Switch vPC+					
1.6.2.1. DC2-Dist-N7k-102		Setup interfaces from ToR Layer 2 Switch vPC+ to the DC2-Dist-N7k-102	IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.6.3. ToR N3k Layer 3					
1.6.3.1. DC2-Dist-N7k-102		Setup interface from ToR N3k Layer 3 to DC2-Dist-N7k-102	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. IGMP/MLD Snooping: Verify IGMP/MLD Snooping ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines. ACL: 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.	pass	
1.6.3.2. DC2-Dist-C6kE8-104		Setup interface from ToR N3k Layer 3 to DC2-Dist-C6kE8-104	OSPF: Verify OSPFv2/OSPFv3 peering. PIM: Verify PIM peering. IGMP/MLD Snooping: Verify IGMP/MLD Snooping ARP & MAC : Verify ARP and MAC addresses are properly learnt across all the forwarding engines.	pass	

			ACL: 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.		
1.6.4. ToR N5k vPC					
1.6.4.1. DC2-Dist-N7k-101		Setup interface from ToR N5k vPC Switch to DC2-Dist-N7k-101	vPC: Verify vPC peer status and consistency parameters. Check MAC/ARP/igmp snooping synchronization. IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.6.5. ToR N5k FabricPath					
1.6.5.1. DC2-Dist-N7k-102		Setup interfaces from ToR N5k FabricPath to the DC2-Dist-N7k-102	FabricPath: Verify FabricPath route and mac-table are built as expected. Verify IS-IS database. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected HSRP: Verify HSRP MAC address is programmed in the mac table IGMP/MLD Snooping: Verify IGMP/MLD Snooping STP: Verify RSTP parameters and port status. VACL, PACL: Verify that all the policies are properly programmed in hardware.	pass	
1.7. ToR to Hosts Setup					
1.7.1. FEX					
1.7.1.1. End Host		Setup interface from FEX to End Host (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.1.2. End Host vPC		Setup interface from FEX to End Host vPC (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.1.3. UCS Fabric Interconnect		Setup interface from FEX to UCS Fabric Interconnect	Verify spanning tree status (edge) on all vlans for the host ports.		

			Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.1.4. UCS Fabric Interconnect vPC		Setup interface from FEX to UCS Fabric Interconnect vPC	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.1.5. UCS Fabric Interconnect vPC+		Setup interface from FEX to UCS Fabric Interconnect vPC+	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.2. ToR Layer 2 Switch					
1.7.2.1. End Host		Setup interface from ToR Layer 2 Switch to End Host (traffic generator)	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.2.2. UCS Fabric Interconnect		Setup interface from ToR Layer 2 Switch to UCS Fabric Interconnect	Verify spanning tree status (edge) on all vlans for the host ports. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.3. ToR N3k Layer 3					
1.7.3.1. End Host		Setup interface from ToR N3k Layer 3 Switch to End Host (traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.4. ToR N5k vPC					
1.7.4.1. FEX vPC		Setup interface from ToR N5k FEX to End Host vPC (traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.4.1. UCS Fabric Interconnect vPC		Setup interface from ToR N5k vPC to UCS Fabric Interconnect vPC	Verify spanning tree status on all vlans. Verify mac table is populated correctly.		

			Verify IGMP/MLD snooping.		
1.7.5. ToR N5k Fabricpath Leaf					
1.7.5.1. UCS Fabric Interconnect vPC+		Setup interface from ToR N5k FP to UCS Fabric Interconnect vPC+	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.2. End Host vPC+		Setup interface from ToR N5k FP to End Host vPC+ (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.3. End Host		Setup interface from ToR N5k FP to End Host (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.7.5.4. ToR L2 switch		Setup interface from ToR N5k FP to ToR L2 switch	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.5.5. ToR L2 switch vPC+		Setup interface from ToR N5k FP to ToR L2 switch vPC+	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.	pass	
1.7.5.6. FEX vPC+		Setup interface from N5k FP ToR FEX vPC+ to End Hosts (Traffic generator)	Verify spanning tree status on all vlans. Verify mac table is populated correctly. Verify IGMP/MLD snooping.		
1.8. UCS Setup					
1.8.1. Fabric Interconnect					
1.8.1.1. DC2-Dist-N7k-101					
1.8.1.1.1. UCS to N7K FEX		Setup for UCS 62xx FI to FEX	Verify the two FI's are in a cluster. Verify FI end host mode configuration.		

			Verify uplink port-channels towards FEX. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and pinning.		
1.8.1.1.2. UCS to N7K VPC		Setup for UCS 62xx FI to FEX	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards ToR FEX. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.1.3. UCS to Layer 8 Switch		Setup for UCS 62xx FI to Layer 2 Switch	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards layer 2 switch. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and pinning.		
1.8.1.1.4. UCS to N5k VPC		Setup for UCS 62xx FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.1.5. UCS to N7K VPC		Setup for UCS 62xx FI to N7K VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N7k VPC. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.2. DC2-Dist-N7k-102					
1.8.1.2.1. UCS to N7K FabricPath VPC+		Setup for UCS 62xx FI to N7k FabricPath VPC+	Verify the two FI's are in a cluster.		

			Verify FI end host mode configuration. Verify uplink port-channels towards N7k VPC+. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.2.2. UCS to Layer 2 Switch		Setup for UCS 62xx FI to Layer 2 Switch	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards the layer 2 switch. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.2.3. UCS to N5K FabricPath VPC+		Setup for UCS 62xx FI to N5k VPC+	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC+. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.3. DC2-Dist-C6kE8-103-VSS					
1.8.1.3.1. UCS to C6kE8 Standalone		Setup for UCS 62xx FI to C6kE8 Standalone	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.4. DC2-Dist-C6kE8-104					
1.8.1.4.1. UCS to C6kE8 VSS		Setup for UCS 62xx FI to C6kE8 VSS	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify dynamic pinning on the FI uplinks.		

			Verify IOM to FI connectivity and port-channel mode.		
1.8.1.4.2. UCS to N5k VPC		Setup for UCS 62xx FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.5. DC2-Dist-C6kE7-105-VSS					
1.8.1.5.1. UCS to C6kE7 Standalone		Setup for UCS 62xx FI to C6kE7 Standalone	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.6. DC2-Dist-C6kE7-106					
1.8.1.6.1. UCS to C6kE7 VSS		Setup for UCS 62xx FI to C6kE7 VSS	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards C6k. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.6.2. UCS to N5k VPC		Setup for UCS 62xx FI to N5k VPC	Verify the two FI's are in a cluster. Verify FI end host mode configuration. Verify uplink port-channels towards N5k VPC. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.1.7. DC2-Dist-C4k-107					
1.8.1.7.1. UCS to C4k		Setup for UCS 62xx FI to C4k	Verify the two FI's are in a cluster.		

			Verify FI end host mode configuration. Verify uplink port-channels towards C4k. Verify dynamic pinning on the FI uplinks. Verify IOM to FI connectivity and port-channel mode.		
1.8.2. UCS Blade					
1.8.2.1. Hypervisor Installation		Setup hypervisor for server virtualization	Verify the hypervisor software installation on the B2xx Mx blade. Verify server's IP address can be pinged. Verify the configured VM's are up and running. Verify the distributed virtual switch is functional. Verify successful installation of operating systems. Verify traffic can be generated by the servers.		
1.8.2.2. Nexus 1000V Installation		Setup Nexus 1000V	Verify that the Nexus 1000V is installed. Verify the network configurations for control, packet and management vlans. Verify the configured VEM's are up and running. Verify the distributed virtual switch is functional. Verify successful installation of operating systems. Verify traffic can be generated by the servers.		
1.8.2.3. VM FEX Installation		Setup VM FEX	Verify that policies are applied to the VM servers.		
2. Network Disruptions Test Cases		**Network Disruptions Test Cases** Common checks for all network disruptions	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 UCS end host mode on FI and VM-FEX functionality. Verify UCS unicast/multicast traffic convergence		
2.1. L2 Link Failure/Recovery	2.1.1	L2 Port-channel Failure/Recovery between Distribution and ToR devices	Verify STP port states after link disruption are in the expected forwarding mode. Verify that the STP root does not change. Verify HSRP peers status does not change. Verify HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on	pass	

			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. On the aggregation switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. 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 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 link at the access switch 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, HSRP 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.		
	2.1.2	L2 port-channel member failure/recovery between Distribution and ToR devices	Verify port-channel load balancing and rbh assignment Verify that IGMP/MLD membership is not affected. 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 will be proportionate to number of members failed	pass		

			Multicast DR should not change. Verify that there is no protocol flapping.	
	2.1.3	vPC leg failure/recovery between Distribution and ToR devices	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.	pass
	2.1.4	vPC leg member failure/recovery between Distribution and ToR devices	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 there are 2 members on each vPC leg). 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.	
	2.1.5	vPC peer-link failure/recovery between Distribution vPC peer switches	Verify that the operational secondary vPC peer will bring down the vPC member ports. Verify that secondary peer will suspend the vpc vlan svi's. Verify that on recovery, the original states will be re-established.	pass
	2.1.6	vPC Peer-keepalive failure/recovery between Distribution vPC peer switches	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.	pass
	2.1.7	vPC peer-link and keep-alive failure between Distribution vPC peer switches	If the keep-alive fails first followed by vPC peer link, then both vPC peers will become active. Verify dual-active scenario is encountered and with the peer-switch feature enabled, ensure the downstream device does not detect any spanning-tree misconfigurations. If the vPC peer-link fails first followed by the keep-alive link, the secondary should keep it's vPC member ports suspended. With vPC auto-recovery configured if the vPC peer-link fails first followed by the keep-alive link, the secondary will keep it's vPC member ports suspended for the duration of three consecutive keepalive failures. After the timer expires the member ports will be unsuspended and the system will change role to primary causing Dual-active scenario.	pass
	2.1.8	vPC peer-link and keep-alive recovery from Dual-active between Distribution vPC peer switches	If keep-alive is recovered first, the active/secondary switch is determined by the role priority and the secondary switch will suspend vPC member ports and vpc svi's. If vpc peer link is recovered first followed by keep alive, the active/secondary switch is determined by the role priority and the system resumes.	pass

	2.1.9	OTV VDC L2 Link Failure/Recovery	Verify traffic will recover after link recovery.		
2.2. L3 Link Failure/Recovery	2.2.1	L3 Port-channel Failure/Recovery between Edge and Public Network[Interop between N7K, C6K]	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 PIM neighbor status. Verify GRE Tunnel re-route due to transport disruption. Verify MTU fragmentation and reassembling at tunnel edge. Verify AutoRP mapping and boundaries.	pass	
	2.2.2	L3 Port-channel Failure/Recovery between Core and Distribution Layers[Interop between N7K, ASR9k, C6K, C4k]	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..	pass	

			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.			
	2.2.3	L3 Port-channel Failure/Recovery between Distribution to ToR N3k Layer 3 [Interop between N7K & N3K; C6K &N3k]	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.	pass		

			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.		
	2.2.4	L3 port-channel member failure/recovery	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.	pass	
	2.2.5	OTV VDC L3 Link Failure/Recovery	Verify traffic will recover after link recovery.		
2.3. Clear OSPF Neighbors/Process/Routes	2.3	Clear OSPF Neighbors/Process/Routes	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.		
2.4. Clear IPv4/IPv6 Multicast Routes	2.4	Clear IPv4/IPv6 Multicast Routes	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.	pass	

2.5. Reload and Power Cycle Switch	2.5.1	Reload and Power Cycle Edge/Core Switch	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.	pass	
	2.5.2	Reload and Power Cycle Distribution Switch	Verify STP port states during and after reload. Verify HSRP 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 HSRP MAC in ARP table. Verify HSRP MAC address is programmed as a router/static MAC on the active switch and a dynamic entry on the standby switch.	pass	

			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 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)		CSCui92396
	2.5.3	vPC peer switch VDC reload	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)	pass	
2.6. Supervisor and Fabric HA	2.6.1	Supervisor HA on the edge/core layer	Compare startup/running configuration on Active Sup and Standby Sup before and after SSO. 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 (control-plane & data-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 SSO. 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 SSO. Verify BFD peer should not flap during and after SSO.	pass	

			No traffic loss is expected.			
	2.6.2	Supervisor HA on the Distribution layer	Compare startup/running configuration on Active Sup and Standby Sup before and after SSO. Verify STP port states during and after SSO. Verify HSRP peers status during and after SSO. Verify CDP/LLDP status after SSO. Verify ARP tables remain unaffected Verify HSRP MAC in ARP table. Verify OTV ARP optimization/ARP caching works as expected after SSO. Verify head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information. Verify automated mapping of OTV sites multicast groups to transport multicast group. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding after SSO. Verify IGMP snooping entries remain unaffected. 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 SSO. Verify SNMP traps are sent to SNMP collector. Verify traffic destined for CoPP classes is policed as expected. 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 SSO. Verify PIM neighbor status.	pass		

			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 SSO. Verify BFD peer should not flap during and after SSO. Verify vPC peer status (role, peer link, keepalive link and consistency parameters) before and after SSO No traffic loss is expected.		
	2.6.3	Fabric Failover on the Edge/Core and Distribution Layers	Verify there is no impact to data plane and control plane on Fabric failover with no oversubscription	pass	
2.7. Line Card OIR and Reset	2.7.1	L3 port-channel member failure/recovery, on OIR/reset line card	Verify hitless operation for non-affected ports Verify traffic load-balancing for distributed port-channels before and after OIR/reset Verify BGP/ IGP/ PIM reconvergence (control-plane & data plane) Verfiy BFD peer detection and client notifications Verify LACP interoperability for distributed port-channels Verify that CDP/LLDP does not lose peer information for non-affected line card. Verify that CDP/LLDP peer is removed for disrupted line cards. Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. 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. All unicast and multicast traffic should re-converge with minimal packet loss. Verify traffic destined for CoPP classes is policed as expected.	pass	
	2.7.2	L2 port-channel member failure/recovery, on OIR/reset line card	Verify port-channel load balancing and rbh assignment Verify LACP interoperability for distributed port-channels Verify STP port states after OIR/reset are in the expected forwarding mode. Verify HSRP peers status after OIR/reset. Verify HSRP MAC in ARP table.	pass	

			Verify IGMP/MLD snooping entries are deleted for the links of affected line card and re-learnt correctly on the alternative link after query from the IGMP snooping router. Verify that IGMP/MLD membership is not affected. Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. 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. The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping.		
	2.7.3	vPC leg failure/recovery, on OIR/reset line card	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.	pass	
	2.7.4	vPC leg member failure/recovery on OIR/reset line card	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 there are 2 members on each vPC leg). 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.	pass	
	2.7.5	vPC peer-link failure/recovery on OIR/reset line card	Verify that the operational secondary vPC peer will bring down the vPC member ports. Verify that secondary peer will suspend the vpc vlan svi's. Verify that on recovery, the original states will be re-established.	pass	
	2.7.6	vPC Peer-keepalive failure/recovery on OIR/reset line card	There are 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.	pass	

	2.7.7	vPC peer-link and peer-keepalive failure on OIR/reset line card	If the keep-alive fails first followed by vPC peer link, then both vPC peers will become active. Verify dual-active scenario is encountered and with the peer-switch feature enabled, ensure the downstream device does not detect any spanning-tree misconfigurations. If the vPC peer-link fails first followed by the keep-alive link, the secondary should keep it's vPC member ports suspended. With vPC auto-recovery configured if the vPC peer-link fails first followed by the keep-alive link, the secondary will keep it's vPC member ports suspended for the duration of three consecutive keepalive failures. After the timer expires the member ports will be unsuspended and the system will change role to primary causing Dual-active scenario.	pass	
	2.7.8	vPC peer-link and peer-keepalive recovery on OIR/reset line card	If keep-alive is recovered first, the active/secondary switch is determined by the role priority and the secondary switch will suspend vPC member ports and vpc svi's. If vpc peer link is recovered first followed by keep alive, the active/secondary switch is determined by the role priority and the system resumes.	pass	
2.8. ISSU/ISSD	2.8.1	ISSU/ISSD	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 HSRP peers status during and after ISSU/ISSD. Verify CDP/LLDP status after ISSU/ISSD. Verify HSRP MAC in ARP table. Verify HSRP 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 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.	pass	

			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.		
2.9. Configuration Change	2.9.1	Perform VPC Vlan add and delete	Verify STP port states after each change are in the expected forwarding mode.		
	2.9.2	Perform VPC SVI add and delete	Verify HSRP peers status after each change.		
	2.9.3	Perform Non-VPC Vlan add and delete	Verify the L2 forwarding table should be updated correctly after each change.		
	2.9.4	Perform Non-VPC SVI add and delete	Verify HSRP MAC in ARP table.		
	2.9.5	Remove VDC and add it back	Verify that no flooding happens after traffic convergence.		
	2.9.6	Enable/Disable IGMP snooping	Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines.		
	2.9.7	Perform HSRP active/standby switchover by changing priority	Verify IGMP/MLD snooping entries are deleted and re-learnt correctly upon each disruption. DHCP relay configured on the spine switches should remain unaffected after each change. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP after each change. All unicast and multicast traffic should re-converge with expected packet loss.		

			Verify SNMP traps are sent to SNMP collector. Verify that all unicast/multicast traffic convergence.		
2.10.FabricPath – Network disruptions					
2.10.1. FabricPath – Link Failure/Recovery	2.10.1.1	FabricPath - Core Link Failure/Recovery	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify HSRP peers status does not change. Verify HSRP MAC in ARP table. Verify HSRP 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 SNMP traps are sent to SNMP collector. 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 are programmed as adjacencies for L3 next hop forwarding. 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 link 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 SPAN is mirroring packets correctly. 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, HSRP 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.	pass	

			Verify traffic destined for CoPP classes is policed as expected.		
	2.10.1.2	Fabricpath - Core Link member failure/recovery	Verify port-channel load balancing and RBH assignment. Verify IS-IS database, topology and route distribution for metric change. Verify that IGMP/MLD membership is not affected. Verify that IGMP snooping entries change based on multi-destination tree topology change. The maximum traffic disruption for unicast/multicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping.	pass	
	2.10.1.3	Fabricpath - vPC+ leg failure/recovery	The maximum traffic disruption for unicast will be half for both upstream and downstream traffic or no loss. 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.	pass	
	2.10.1.4	Fabricpath - vPC+ leg member failure/recovery	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 there are 2 members on each vPC+ leg). 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.	pass	
	2.10.1.5	Fabricpath - vPC+ peer-link failure/recovery (spine/leaf)	Verify that the operational secondary vPC+ peer will bring down the vPC+ member ports. Verify that secondary peer will not suspend the vPC+ vlan SVI's if "*dual-active exclude vlans*" is configured Verify on recovery that the operational secondary vPC+ peer will bring up the vPC+ member ports after the configured "*delay restore*" timer	pass	
	2.10.1.6	Fabricpath - vPC+ Peer-keepalive failure/recovery	There are no expected effects; both vPC+ peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.	pass	
	2.10.1.7	Fabricpath - vPC+ peer-link and Peer-keepalive failure/recovery	When the keep-alive fails first followed by vPC+ peer link, the peers should continue to see each other through fabricpath network. The effect should be same as just peer-link failure. The recovery should be same as the peer-link recovery.	pass	

2.10.2. FabricPath – Reload	2.10.2.1	FabricPath - Spine Node failure/recovery	Verify Fabricpath multi-destination trees reconverge after root change on node failure. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding. Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on the other spine routers Verify IGMP/MLD snooping entries are deleted for the affected link at the access switch 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 other spine routers. Verify SPAN is mirroring packets correctly. Verify SNMP traps are sent to SNMP collector. 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, HSRP services, ARP, proxy arp and IGMP. All unicast and multicast traffic should re-converge with minimal packet loss. Verify traffic destined for CoPP classes is policed as expected. Verify that the MAC table, FP ISIS route table, ARP table, IP routing table, IGMP membership table, IGMP snooping table, Multicast routing table return to original state on recovery Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on recovery	pass	
	2.10.2.2	FabricPath - Leaf Node failure/recovery	Verify Fabricpath multi-destination trees reconverge after leaf node failure. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify HSRP peers status does not change when CE or leaf switches are reloaded. Verify IGMP/MLD snooping entries are deleted for the affected link at the access switch 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 spine routers. Verify that the MAC table, FP ISIS route table, IGMP snooping table return to original state on recovery	pass	

			Verify the L2/L3 forwarding entries are synchronized among the hardware forwarding engines on recovery		
2.10.3. FabricPath – Supervisor and Fabric HA	2.10.3.1	FabricPath – Supervisor HA on the spine nodes	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Compare startup/running configuration on Active Sup and Standby Sup before and after SSO. Verify STP port states during and after SSO. Verify HSRP peers status during and after SSO. Verify CDP/LLDP status after SSO. Verify HSRP MAC in ARP table. Verify HSRP 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 are programmed as adjacencies for L3 next hop forwarding after SSO. 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 SSO. Verify SNMP traps are sent to SNMP collector. Verify traffic destined for CoPP classes is policed as expected. 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 SSO. Verify PIM neighbor status. Verify static RP mapping as the backup of auto RP.	pass	

			Verify MSDP neighbors and SA cache consistency. Verify multicast HW and SW entries are properly programmed and synchronized after SSO. Verify BFD peer should not flap during and after SSO. Verify vPC+ peer status (role, peer link, keepalive link and consistency parameters) before and after SSO No traffic loss is expected.		
	2.10.3.2	FabricPath - Fabric Failover on spine nodes	Verify there is no impact to data plane and control plane on Fabric failover with no oversubscription	pass	
2.10.4. FabricPath – Line card OIR and Reset	2.10.4.1	FabricPath – Line card OIR and Reset on spine nodes	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify hitless operation for non-affected ports Verify traffic load-balancing for distributed port-channels before and after OIR/reset Verfiy BFD peer detection and client notifications Verify LACP interoperability for distributed port-channels Verify STP port states after OIR/reset are in the expected forwarding mode. Verify HSRP peers status after OIR/reset. Verify that CDP/LLDP does not lose peer information for non-affected line card. Verify that CDP/LLDP peer is removed for disrupted line cards. Verify the L2 forwarding table should be re-learnt correctly after OIR/reset. Verify HSRP MAC in ARP table. 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 links of affected line card and re-learnt correctly on the alternative link after query from the IGMP snooping router. Verify SPAN is mirroring packets correctly. Verify SNMP traps are sent to SNMP collector. All unicast and multicast traffic should re-converge with minimal packet loss.	pass	

			Verify traffic destined for CoPP classes is policed as expected.		
	2.10.4.2	FabricPath – FP core port-channel member failure/recovery, on OIR/reset line card	Verify port-channel load balancing and rbh assignment Verify that IGMP/MLD membership is not affected. The maximum traffic disruption for unicast should be in sub-second range for both upstream and downstream traffic. Multicast DR should not change. Verify that there is no protocol flapping.	pass	
	2.10.4.3	FabricPath – vPC+ leg failure/recovery on OIR/reset line card	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.	pass	
	2.10.4.4	FabricPath – vPC+ leg member failure/recovery on OIR/reset line card	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 there are 2 members on each vPC+ leg). 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.	pass	
	2.10.4.5	FabricPath – vPC+ peer-link failure/recovery on OIR/reset line card	Verify that the operational secondary vPC+ peer will bring down the vPC+ member ports. Verify that secondary peer will not suspend the vPC+ vlan SVI's if "*dual-active exclude vlans*" is configured Verify on recovery that the operational secondary vPC+ peer will bring up the vPC+ member ports after the configured "*delay restore*" timer	pass	
	2.10.4.6	FabricPath – vPC+ Peer-keepalive failure/recovery on OIR/reset line card	There are no expected effects; both vPC+ peers continue to synchronize MAC address tables, IGMP entries, no traffic disruptions.	pass	
	2.10.4.7	Fabricpath - vPC+ peer-link and Peer-keepalive failure/recovery on OIR/reset line card	When the keep-alive fails first followed by vPC+ peer link, the peers should continue to see each other through fabricpath network. The effect should be same as just peer-link failure. The recovery should be same as the peer-link recovery.	pass	
2.10.5. FabricPath – ISSU/ISSD	2.10.5.1	FabricPath – ISSU/ISSD	Verify if ISSU image compatibility for non-disruptive upgrade/downgrade	pass	

			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 FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. Verify STP port states during and after ISSU/ISSD. Verify HSRP peers status during and after ISSU/ISSD. Verify CDP/LLDP status after ISSU/ISSD. Verify HSRP MAC in ARP table. Verify HSRP 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 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. All unicast and multicast traffic should re-converge. Verify traffic destined for CoPP classes is policed as expected. 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.		
2.10.6. FabricPath – Move /Add/Change Hosts	2.10.6.1	FabricPath – MAC move	Verify ARP tables remain unaffected, MAC table shows mac move. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. 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 properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. All unicast and multicast traffic should re-converge with minimal packet loss. Verify SNMP traps are sent to SNMP collector.		
	2.10.6.2	FabricPath – End Hosts Add	Verify ARP and MAC tables add the new hosts. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding.		

	2.10.6.3	FabricPath – End Hosts Change	Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines. Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. Verify ARP and MAC tables change as expected. Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding. Verify that no flooding happens after traffic convergence. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines. Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches. DHCP relay configured on the spine switches should remain unaffected. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. Monitor all unicast/multicast traffic convergence.		
2.10.7. FabricPath – Configuration Change	2.10.7.1 2.10.7.2 2.10.7.3 2.10.7.4 2.10.7.5 2.10.7.6	Perform FP Vlan add and delete Perform FP SVI add and delete Perform Non-FP Vlan add and delete Perform Non-FP SVI add and delete Perform FP MT root move by changing priority Enable/Disable IGMP snooping	Verify FabricPath route and mac-table are built as expected. Verify IS-IS database, topology and route distribution. Verify multi-destination trees for unknown unicast, broadcast and multicast. Verify fabricpath load-balance works as expected. 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 after each change. On the spine switches, verify that the ARP are programmed as adjacencies for L3 next hop forwarding after each change. Verify that no flooding happens after traffic convergence after each change. Verify the L2/L3 forwarding entries are synchronized among the hardwre forwarding engines after each change.	pass	

			Verify IGMP/MLD snooping entries are properly relearned on the affected FP switches after each change. DHCP relay configured on the spine switches should remain unaffected after each change. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP after each change. 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 on all the affected FP switches. All unicast and multicast traffic should re-converge with minimal packet loss. Verify SNMP traps are sent to SNMP collector. Monitor all unicast/multicast traffic convergence.		
2.11.OTV – Network Disruptions					
2.11.1. OTV – Reload	2.11.1.1	OTV – Reload	Verify HSRP isolation across OTV sites works as expected after reload/recovery. Verify OTV ARP optimization/ARP caching works as expected after reload/recovery. Verify unknown unicast traffic doesn't flood. Verify STP is blocked across OTV sites. Verify the Secondary Adj. Server will take over after primary Adj. Server failover. 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 head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information. Verify automated mapping of OTV sites multicast groups to transport multicast group. Verify IGMP snooping entries are properly relearned on the affected OTV switches. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.11.2. OTV – Move/Add/Change Hosts	2.11.2.1	OTV – MAC move/Add/Change Hosts	Verify HSRP isolation across OTV sites works as expected. Verify OTV ARP optimization/ARP caching works as expected. Verify unknown unicast traffic doesn't flood. Verify the new hosts's macs are learnt across OTV sites. Verify STP is blocked across OTV sites.		

			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 head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information. Verify automated mapping of OTV sites multicast groups to transport multicast group. Verify IGMP snooping entries are properly relearned on the affected OTV switches. Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy arp and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.11.3. OTV – Configuration Change	2.11.3.1	Add and delete OTV VLAN	Verify HSRP isolation across OTV sites works as expected		
	2.11.3.2	Add and delete OVT SVI	Verify OTV ARP optimization/ARP caching/ARP suppression works as expected.		
	2.11.3.3	Enable and disable proxy ARP	Verify unknown unicast traffic doesn't flood.		
	2.11.3.4	Enable and disable suppression ARP	Verify STP is blocked across OTV sites.		
	2.11.3.5	Enable and disable igmp snooping	Verify new Adj. Server works as expected.		
	2.11.3.6	Add and delete overlay interface	Verify the new hosts's macs are learnt across OTV sites.		
	2.11.3.7	Dynamically changing Adj Server	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.		
	2.11.3.8	Add/remove/flush MAC entries	Verify head-end replication for multicast traffic on unicast-only transport works as expected, check the data-group mapping table for receiver information.		
	2.11.3.9	Add/remove/flush ARP entries	Verify automated mapping of OTV sites multicast groups to transport multicast group.		
	2.11.3.10	Add/remove/flush multicast group entries	Verify IGMP snooping entries are properly relearned on the affected OTV switches.		
	2.11.3.11	Add/remove/flush active multicast source entries	Verify that secondary addresses provide the same capability and services to nodes through DHCP relay, HSRP services, ARP, proxy ARP and IGMP. Verify SNMP traps are sent to SNMP collector.		
2.12.UCS – Disruptions					
2.12.1. UCS – Link Failure/Recovery		UCS - Link Failure/Recovery	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify DHCP/BOOTP functionalities.		
		FI Uplink port-channel member failure/recovery:	Verify unicast and multicast traffic should re-converge with minimal packet loss.		

			Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links is not impact. Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is not affected.			
		FI Uplink port-channel failure/recovery:	Verify unicast and multicast traffic should switch to other FI and re-converge with expected packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify there is no mac address learning on FI uplink. Verify mac learning on other FI server links. Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is working as expected.			
		FI to IOM port-channel member failure/recovery:	Verify unicast and multicast traffic should re-converge with minimal packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links is not impact. Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is not affected.			
			FI to IOM port-channel failure/recovery: Verify unicast and multicast traffic should switch to other FI and re-converge with expected packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify there is no mac address learning on FI uplink.			

			Verify mac learning on other FI server links. Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is working as expected.		
		FI cluster link failure/recovery:	Verify unicast and multicast traffic should have no impact. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works with no impact. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links is not impact. Verify that IGMP snooping is not affected.		
		FI to FI isolation/recovery:	Verify unicast and multicast traffic should re-converge after FI cluster link recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after FI cluster link recovery. Verify there is no mac address learning on FI uplink after FI cluster link recovery. Verify mac learning on other FI server links after FI cluster link recovery. Verify that no flooding happens after traffic convergence after FI cluster link recovery. Verify that IGMP snooping is working as expected after FI cluster link recovery.		
		MEM and CPU:	Monitor MEM and CPU Usage on FI.		
		Convergence:	Measure unicast/multicast traffic convergence for each disruption		
2.12.2. UCS – Fabric Interconnect Reload and Power Cycle		UCS – Fabric Interconnect Reload and Power Cycle	Verify unicast and multicast traffic should switch to other FI and re-converge with expected packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify there is no mac address learning on other FI uplink. Verify mac learning on other FI server links.		

			Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is working as expected. Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify DHCP/BOOTP functionalities.		
		MEM and CPU:	Monitor MEM and CPU Usage on FI.		
		Convergence:	Measure unicast/multicast traffic convergence for each disruption		
2.12.3. UCS – IOM OIR		UCS – IOM OIR	Verify unicast and multicast traffic should switch to other FI and re-converge with expected packet loss. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected. Verify GARP is sent by other FI after fabric switchover. Verify there is no mac address learning on other FI uplink. Verify mac learning on other FI server links. Verify that no flooding happens after traffic convergence. Verify that IGMP snooping is working as expected. Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption		

2.12.4. UCS – Blade OIR		UCS – Blade OIR	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after blade recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade recovery. Verify that no flooding happens after traffic convergence after blade recovery. Verify that IGMP snooping is working as expected after blade recovery. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption		
2.12.5. UCS – Chassis Reload and Power Cylce		UCS – Chassis Reload and Power Cycle	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after chassis IOM and blade recovery. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after chassis IOM and blade recovery. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after chassis IOM and blade recovery. Verify that no flooding happens after traffic convergence after chassis IOM and blade recovery. Verify that IGMP snooping is working as expected after chassis IOM and blade recovery. Verify DHCP/BOOTP functionalities. MEM and CPU:		

			Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption		
2.12.6. UCS – FI image and IOM Firmware Upgrade		UCS – FI image and IOM Firmware Upgrade	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after IOM firmware upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after IOM firmware upgraded. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after IOM firmware upgraded. Verify that no flooding happens after traffic convergence after IOM firmware upgraded. Verify that IGMP snooping is working as expected after IOM firmware upgraded. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption		
2.12.7. UCS – Blade adapter Firmware upgrade		UCS – Blade adapter Firmware upgrade	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after blade adapter firmware upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade adapter firmware upgraded. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade adapter firmware upgraded. Verify that no flooding happens after traffic convergence after blade adapter firmware upgraded.		

			Verify that IGMP snooping is working as expected after blade adapter firmware upgraded. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption.		
2.12.8. UCS – Blade BIOS upgrade		UCS – Blade BIOS upgrade	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after blade BIOS upgraded. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after blade BIOS upgraded. Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after blade BIOS upgraded. Verify that no flooding happens after traffic convergence after blade BIOS upgraded. Verify that IGMP snooping is working as expected after blade BIOS upgraded. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption.		
2.12.9. UCS – VMotion		UCS – VMotion	Verify FI uplink dynamic pinning works as expected. Verify that CDP/LLDP does not lose peer information for non-affected links. Verify that CDP/LLDP peer is removed for disrupted link. Verify SNMP traps are sent from FI to SNMP collector. Verify unicast and multicast traffic should re-converge after VMotion. Verify RPF check/ Déjà vu check/ Broadcast traffic pinning works as expected after VMotion.		

			Verify there is no mac address learning on FI uplink. Verify mac learning on FI server links after VMotion. Verify that no flooding happens after traffic convergence after VMotion. Verify that IGMP snooping is working as expected after VMotion. Verify DHCP/BOOTP functionalities. MEM and CPU: Monitor MEM and CPU Usage on FI. Convergence: Measure unicast/multicast traffic convergence for each disruption.		