

## Deploying MC-LAG to VPLS Solution

This solution requires that the two Nexus 7000 aggregation switches be converted to vPC and connected to both the PE routers via mLACP. The PWs on the both the N-PEs are in active state because of the decoupled mode of operation where the state of the attachment circuits controlled by mLACP is independent from the state of the PWs.

This section provides information about the test scope, hardware and software version used during design validation and includes key configuration details to implement this solution.

### Scope

This engagement validated the MC-LAG to VPLS LAN extension solution with Cisco ASR 9000 series routers as PE. The testing was performed for the following features:

**Table 5-1** Test Scope

Features/Tests	Description
L2 Interconnect	Verify L2 Interconnect between data centers and VLAN extension using VPLS
VPLS	Verify VPLS using ASR 9000 (one VFI/VLAN) as N-PE with Nexus 7000 as aggregation devices
mLACP	Verify mLACP on ASR 9000 to achieve N-PE redundancy
RSTP and HSRP	Verify RSTP and HSRP functionality
Storm Control	Verify Storm Control functionality
VLAN extension	Extend 500 and 1200 VLANs between data centers
<b>Negative Tests</b>	
Reload	Reload Nexus 7000 and ASR 9000
Link Failure	Shut/No Shut links between various nodes
SSO	SSO on Nexus 7000
ISSU	ISSU on Nexus 7000

# Hardware and Software

Table 5-2 lists all the hardware and software used in validating MC-LAG based VPLS solution to interconnect data centers.

**Table 5-2 Hardware and Software Information**

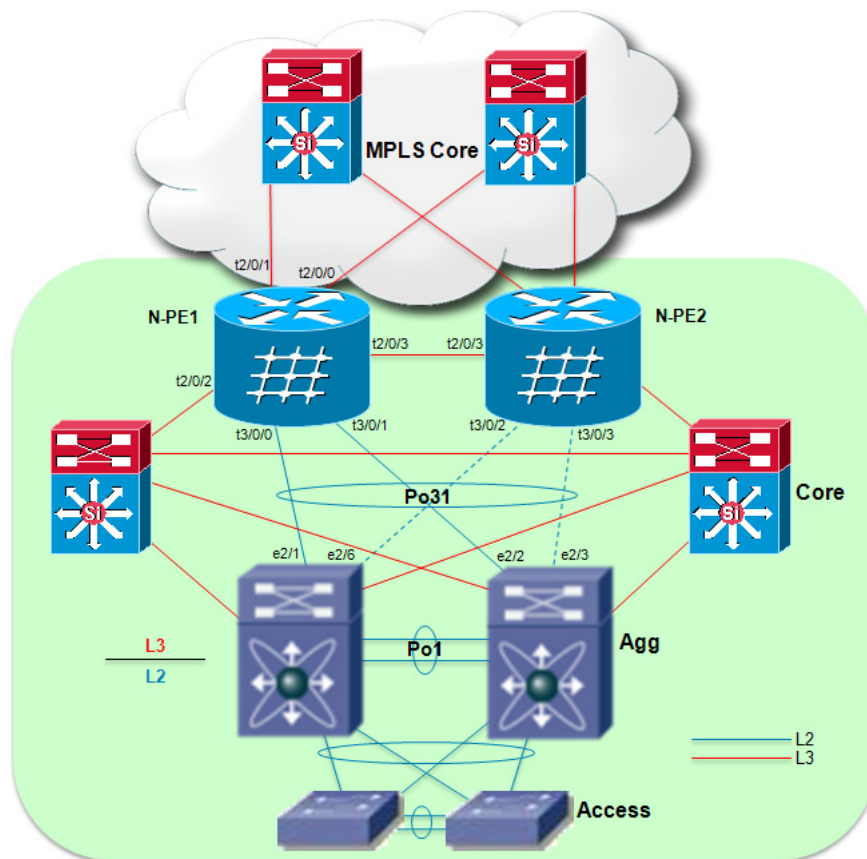
Qty	Hardware Platform	Software Version	Role	Line cards / Interfaces
2	Nexus 7010	Engineering image based on 5.1.2 <sup>1</sup>	DC3 Aggregation	N7K-M108X2-12L N7K-M148GT-11
4	Catalyst 6500	12.2(33)SX13	DC2 and DC3 core	WS-6704-10GE WS-6708-10GE
2	Cisco 7600	12.2(33)SRE2	DC2 PE	7600-ES+4TG3CXL 7600-ES+XT-2TG3CXL 7600-ES+2TG3CXL
2	ASR 9000	4.0.1	DC3 PE	
2	Nexus 5000	4.2(1)N2(1)	DC3 Access	
1	Cisco WS-C4948E	12.2 (54)SG	DC3 Access	
1	Nexus 1000V	4.0(4)SV1(3a)	Virtual Switch	
2	Fabric Interconnect	4.1(3)N2(1.3i)	Fabric Interconnect for UCS	
2	UCS chassis	1.3(1i)	Blade server	
2	Cisco C4900M	12.2(50)SG	DC2 Access	

1. Engineering image based on NX-OS version 5.1(2) was used during testing on Nexus 7000 switches since issues attributing to high convergence times under failure conditions were uncovered. Thus, it is strongly recommended to use this NX-OS software release 5.2 (and higher) to deploy the solution object of this paper.

## Configuration Details

The following list provides snippets of configuration from the PEs and aggregation devices and output from various show commands for verification. All the configuration samples refer to the DC3 site, which is represented again for reference in [Figure 5-1](#).

Figure 5-1 DC3 Site Under Test

**Note**

In the actual testbed, an active/active MC-LAG deployment was validated, leveraging a redundant set of connections between the aggregation devices and the PE routers (as previously shown in Figure 15). For the sake of simplicity, only one set of cables is shown in Figure 5-1. Also, multiple access layer switches were actually connected to the two aggregation devices (to allow carrying up to 1200 VLANs toward the PE routers), even if the network diagram above shows only a pair of Nexus 5000 for simplicity sake.

**Step 1** Configure IGP, MPLS and targeted-LDP on PE Routers.

- a. OSPF (IGP) configuration on ASR9000 routers: Cisco recommends enabling BFD to detect failures in the path between adjacent L3 enabled interfaces. Also, it is recommended to tune the OSPF timers (throttle lsa, spf and lsa min-arrival) to ensure to run a faster SPF calculation after notification of a topology change event (arrival of an LSA).

```
router ospf ospf300
  bfd minimum-interval 200
  bfd multiplier 3
  timers throttle lsa all 100 100 5000
  timers throttle spf 100 100 5000
  timers lsa min-arrival 80
  area 0
  interface Loopback100
  !
  interface TenGigE0/2/0/0 << MPLS core facing interface
```

```

bfd fast-detect
network point-to-point
!
interface TenGigE0/2/0/1 << MPLS core facing interface
bfd fast-detect
network point-to-point
!
interface TenGigE0/2/0/2 << Local datacenter core facing interface
bfd fast-detect
network point-to-point
!
interface TenGigE0/2/0/3 << Connected to N-PE2 (Redundant PE in the same datacenter)
bfd fast-detect
network point-to-point
!
!
RP/0/RSP1/CPU0:DC3-ASR9K-NPE1#show ospf neighbor

* Indicates MADJ interface

Neighbors for OSPF ospf300

Neighbor ID      Pri   State           Dead Time   Address      Interface
150.1.1.103      1     FULL/ -         00:00:38   150.3.36.3   TenGigE0/2/0/0
    Neighbor is up for 12:55:08
150.1.1.104      1     FULL/ -         00:00:32   150.3.46.4   TenGigE0/2/0/1
    Neighbor is up for 12:55:07
120.3.1.4        1     FULL/ -         00:00:38   120.3.46.4   TenGigE0/2/0/2
    Neighbor is up for 12:55:08
150.3.3.5        1     FULL/ -         00:00:39   150.3.33.5   TenGigE0/2/0/3
    Neighbor is up for 12:55:07

Total neighbor count: 4

```

#### b. Configure MPLS and LDP.

```

mpls ldp
router-id 150.3.3.5
session protection << MPLS LDP session protection
interface TenGigE0/2/0/0 << MPLS enabled interfaces
!
interface TenGigE0/2/0/1
!
interface TenGigE0/2/0/3
!
!
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#show mpls interfaces
Interface          LDP      Tunnel  Enabled
-----
TenGigE0/2/0/0     Yes      No      Yes
TenGigE0/2/0/1     Yes      No      Yes
TenGigE0/2/0/3     Yes      No      Yes

```

#### Step 2 Configure Inter-chassis Communication Protocol (ICCP) on PEs.

Table 5-3 shows two mLAG groups, group 31 and 32, are configured for load-sharing VLANs between the PEs (as previously discussed in the “Active/Active Deployment with Redundant Physical Connections” section in Chapter 3). Notice that the required configuration on the two PE devices is

pretty much identical, with the exception of few commands. Particularly important is to configure the same mLACP system MAC value, because the key principle of MC-LAG is to make the two PE devices appearing as a single entity to the Nexus 7000 switches deployed at the aggregation layer.

**Table 5-3** LACP groups 31 and 32 configured for load-sharing VLANs between the N-PEs

On N-PE1:	On N-PE2:
<pre> redundancy iccp   group 31     mlacp node 1 &lt;&lt; Unique on each POA     mlacp system mac 0000.0000.0031 &lt;&lt; LACP system ID. Recommended to be same on both N-PEs     mlacp system priority 1 &lt;&lt; Recommended be lower than DHD (Aggregation)     mlacp connect timeout 0   member     neighbor 150.3.3.6 &lt;&lt; N-PE2 loopback address !   backbone     interface TenGigE0/2/0/0 &lt;&lt; Facing MPLS core     interface TenGigE0/2/0/1 &lt;&lt; Facing MPLS core !     isolation recovery-delay 100 &lt;&lt; Used when core isolation condition is cleared !   group 32     mlacp node 2     mlacp system mac 0000.0000.0032     mlacp system priority 2     mlacp connect timeout 0   member     neighbor 150.3.3.6 !   backbone     interface TenGigE0/2/0/0     interface TenGigE0/2/0/1 !     isolation recovery-delay 100 ! ! ! </pre>	<pre> redundancy iccp   group 31     mlacp node 2 &lt;&lt; Unique on each POA     mlacp system mac 0000.0000.0031 &lt;&lt; LACP system ID. Recommended to be same on both N-PEs     mlacp system priority 1 &lt;&lt; Recommended be lower than DHD (Aggregation)     mlacp connect timeout 0   member     neighbor 150.3.3.5 &lt;&lt; N-PE1 loopback address !   backbone     interface TenGigE0/2/0/0 &lt;&lt; Facing MPLS core     interface TenGigE0/2/0/1 &lt;&lt; Facing MPLS core !     isolation recovery-delay 100 &lt;&lt; Used when core isolation condition is cleared !   group 32     mlacp node 1     mlacp system mac 0000.0000.0032     mlacp system priority 1     mlacp connect timeout 0   member     neighbor 150.3.3.5 !   backbone     interface TenGigE0/2/0/0     interface TenGigE0/2/0/1 !     isolation recovery-delay 100 ! ! ! </pre>

#### On N-PE1:

```

RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#show iccp group 31
Redundancy Group 31
  member ip:150.3.3.6 (DC3-ASR9K-NPE2), up (connected)
  monitor: route-watch (up)
  backbone interface Te0/2/0/0: up
  backbone interface Te0/2/0/1: up
  enabled applications: mLACP, IGMP Snooping
  isolation recovery delay timer: 100 s, not running

```

#### On N-PE2:

```

RP/0/RSP1/CPU0:DC3-ASR9K-NPE2#show iccp group 31
Redundancy Group 31

```

```

member ip:150.3.3.5 (DC3-ASR9K-NPE1), up (connected)
monitor: route-watch (up)
backbone interface Te0/2/0/0: up
backbone interface Te0/2/0/1: up
enabled applications: mLACP, IGMP Snooping
isolation recovery delay timer: 100 s, not running

```

**Step 3** Configure Attachment Circuits on each N-PE (interface towards aggregation switch)

- a. Configure interfaces that are members of port-channel connecting to aggregation switches. On each PE the 4 physical interfaces were assigned to the two separate MC-LAG groups previously defined to achieve active/active traffic load-balancing.

```

interface TenGigE0/3/0/0
bundle id 31 mode active << Bundle 31
cdp
lACP period short << lACP fast hellos
lACP period short transmit 100
carrier-delay up 0 down 0
!
interface TenGigE0/3/0/1
bundle id 31 mode active
cdp
lACP period short
lACP period short transmit 100
carrier-delay up 0 down 0
!
interface TenGigE0/3/0/2
bundle id 32 mode active
cdp
lACP period short
lACP period short transmit 100
carrier-delay up 0 down 0
!
interface TenGigE0/3/0/3
bundle id 32 mode active
cdp
lACP period short
lACP period short transmit 100
carrier-delay up 0 down 0

```

- b. Configure Bundle interfaces: these are the logical interfaces equivalent to a port-channel on ASR 9000 devices. Each defined bundle interface is assigned to a specific ICCP group (previously created).

```

interface Bundle-Ether31
lACP switchover suppress-flaps 100
mlACP icCP-group 31 << Attach this bundle to previously defined ICCP group
mlACP switchover type revertive << Automatic switchback to Primary role upon failure
recovery
mlACP switchover recovery-delay 40 << Timer to wait before switching back
bundle wait-while 0 << Recommended to be set to 0 to improve convergence
bundle maximum-active links 2 << Defines maximum active links in port-channel.
!
interface Bundle-Ether32
lACP switchover suppress-flaps 100
mlACP icCP-group 32 << Attach this bundle to previously defined ICCP group
mlACP switchover type revertive << Automatic switchback to Primary role upon failure
recovery
mlACP switchover recovery-delay 40 << Timer to wait before switching back
bundle wait-while 0 << Recommended to be set to 0 to improve convergence
bundle maximum-active links 2 << Defines maximum active links in port-channel.

```

- c. Configure sub-interfaces for all the VLANs to be extended on both N-PEs: a separate sub-interface is created for each VLAN that needs to be extended across the MPLS core.

```
interface Bundle-Ether31.100 l2transport
 encapsulation dot1q 100
 rewrite ingress tag pop 1 symmetric
!
interface Bundle-Ether31.101 l2transport
 encapsulation dot1q 101
 rewrite ingress tag pop 1 symmetric
!
interface Bundle-Ether31.102 l2transport
 encapsulation dot1q 102
 rewrite ingress tag pop 1 symmetric
!
interface Bundle-Ether31.103 l2transport
 encapsulation dot1q 103
 rewrite ingress tag pop 1 symmetric
!
...
!
```

**Step 4** Pseudowire configuration on PEs (one VFI per VLAN).

The VPLS configuration that was validated leverages a separate VFI and bridge domain for each VLAN that needed to be extended via VPLS. In IOS-XR this can be achieved by defining a “bridge group” container, under which all the different bridge domains and VFIs are created. ASR9000 supports the use of BGP for auto-discovery of the neighbor PE devices, but in this specific validation effort leveraged static neighbor configuration.

```
l2vpn << l2vpn configuration mode
pw-status
 logging
  pseudowire << Enable Pseudowire status logging
!
pw-class vpls-pw-class << PW class to enable mpls encapsulation
 encapsulation mpls
!
!
bridge group group1
  bridge-domain vlan100 << Define bridge domain
  interface Bundle-Ether31.100 << Aggregation facing sub-interface for VFI
  !
  vfi vfi100 << Define VFI
  neighbor 150.2.2.5 pw-id 100 << Pseudowire peer with VC identifier
  pw-class vpls-pw-class
  !
  neighbor 150.2.2.6 pw-id 100 << Pseudowire peer
  pw-class vpls-pw-class
  !
  neighbor 150.3.3.6 pw-id 100 << Pseudowire peer
  pw-class vpls-pw-class
  !
  neighbor 150.11.11.5 pw-id 100 << Pseudowire peer
  pw-class vpls-pw-class
  !
  neighbor 150.11.11.6 pw-id 100 << Pseudowire peer
  pw-class vpls-pw-class
  !
!
!
bridge-domain vlan101
 interface Bundle-Ether31.101
 !
 vfi vfi101
```





**On Agg2:**

```

interface Ethernet2/1
  lacp rate fast
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1200-1449
  channel-group 32 mode active
  no shutdown

interface Ethernet2/2
  lacp rate fast
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,76-80,100-349
  channel-group 31 mode active
  no shutdown

interface Ethernet2/3
  lacp rate fast
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1,76-80,100-349
  channel-group 31 mode active
  no shutdown

interface Ethernet2/6
  lacp rate fast
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1200-1449
  channel-group 32 mode active
  no shutdown

```

- b. Configure port-channel interface on both Nexus 7000 aggregation switches

**On Agg1:**

```

interface port-channel31
  switchport
  switchport mode trunk
  ip port access-group HSRP_Deny in << Filter HSRP hellos from remote datacenters
  switchport trunk allowed vlan 1,76-80,100-349
  spanning-tree port type edge trunk << Define edge port to improve convergence
  spanning-tree bpdudfilter enable << Filter BPDU for STP isolation
  lacp max-bundle 1 << Maximum links to be active. Configure 1 on both N7k
  vpc 31 << vPC domain for virtual port-channel

interface port-channel32
  switchport
  switchport mode trunk
  ip port access-group HSRP_Deny in
  switchport trunk allowed vlan 1200-1449
  spanning-tree port type edge trunk
  spanning-tree bpdudfilter enable
  lacp max-bundle 1
  vpc 32

```

**On Agg2:**

```

interface port-channel31
  switchport
  switchport mode trunk
  ip port access-group HSRP_Deny in

```

```

switchport trunk allowed vlan 1,76-80,100-349
spanning-tree port type edge trunk
spanning-tree bpdufilter enable
lacp max-bundle 1
vpc 31

```

```

interface port-channel32
  switchport
  switchport mode trunk
  ip port access-group HSRP_Deny in
  switchport trunk allowed vlan 1200-1449
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  lacp max-bundle 1
  vpc 32

```

**c. Access list to filter HSRP hellos configured on both aggregation switches**

```

ip access-list HSRP_Deny
  statistics per-entry
  10 deny udp any 224.0.0.102/32 eq 1985 << Filter specific to HSRP v2 hellos (for
  HSRP v1 that are the default version used on Nexus 7000the address to use would be
  224.0.0.2/32)
  20 permit ip any any

```

**d. Verify vPC and port-channel status**

```
DC3-N7K-AGG1# show vpc brief
```

Legend:

(\*) - local vPC is down, forwarding via vPC peer-link

```

vPC domain id          : 1
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Type-2 consistency status : success
vPC role               : primary
Number of vPCs configured : 5
Peer Gateway           : Disabled
Dual-active excluded VLANs : -

```

vPC Peer-link status

```

-----
id   Port   Status Active vlans
--   -
1    Po1    up     1,76-80,100-999,1100-2199,3051-3100
-----

```

vPC status

```

-----
id   Port   Status Consistency Reason           Active vlans
--   -
..
31   Po31   up     success    success           1,76-80,100-599
..

```

**Step 6** Verify that MC-LAG operation on ASR9000 and Nexus 7000 Configure Nexus 7000 aggregation switches for virtual port-channel (vPC).

**Step 7** The final results of the LACP negotiation between the PE routers and the aggregation layer switches is to activate, on a per redundancy group basis, all the links connected to the Primary PE device. The connections to the Secondary PE must end up in Hot-Standby state, and this needs to happen both on the PE and aggregation devices, to avoid traffic black holing. The following CLI commands allow verifying this behavior.

**On N-PE1:**

```
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#show bundle bundle-ether 31

Bundle-Ether31
  Status: Up
  Local links <active/standby/configured>: 2 / 0 / 2
  Local bandwidth <effective/available>: 20000000 (20000000) kbps
  MAC address (source): 18ef.63e4.249b (Chassis pool)
  Minimum active links / bandwidth: 1 / 1 kbps
  Maximum active links: 2
  Wait while timer: Off
  LACP: Operational
    Flap suppression timer: 100 ms
  mLACP: Operational
    ICCP Group: 31
    Role: Active
    Foreign links <active/configured>: 0 / 2
    Switchover type: Revertive
    Recovery delay: 40 s
    Maximize threshold: 1 link
  IPv4 BFD: Not configured

Port          Device          State          Port ID          B/W, kbps
-----
Te0/3/0/0     Local           Active          0x82d9, 0x9001  10000000
  Link is Active
Te0/3/0/1     Local           Active          0x82d9, 0x9002  10000000
  Link is Active
Te0/3/0/2     150.3.3.6      Standby        0x82da, 0xa003  10000000
  Link is marked as Standby by mLACP peer
Te0/3/0/3     150.3.3.6      Standby        0x82da, 0xa004  10000000
  Link is marked as Standby by mLACP peer
```

**On N-PE2:**

```
RP/0/RSP1/CPU0:DC3-ASR9K-NPE2#show bundle bundle-ether 31

Bundle-Ether31
  Status: mLACP hot standby
  Local links <active/standby/configured>: 0 / 2 / 2
  Local bandwidth <effective/available>: 0 (0) kbps
  MAC address (source): 18ef.63e4.249b (Peer)
  Minimum active links / bandwidth: 1 / 1 kbps
  Maximum active links: 2
  Wait while timer: Off
  LACP: Operational
    Flap suppression timer: 100 ms
  mLACP: Operational
    ICCP Group: 31
    Role: Standby
    Foreign links <active/configured>: 2 / 2
    Switchover type: Revertive
    Recovery delay: 40 s
    Maximize threshold: 1 link
  IPv4 BFD: Not configured

Port          Device          State          Port ID          B/W, kbps
-----
Te0/3/0/2     Local           Standby        0x82da, 0xa003  10000000
  mLACP peer is active
Te0/3/0/3     Local           Standby        0x82da, 0xa004  10000000
  mLACP peer is active
Te0/3/0/0     150.3.3.5      Active          0x82d9, 0x9001  10000000
```

```

Link is Active
Te0/3/0/1      150.3.3.5      Active      0x82d9, 0x9002  10000000
Link is Active

```

**On Agg1:**

```

DC3-N7K-AGG1#show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        S - Switched      R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports      Channel
-----
1      Po1(SU)      Eth       LACP           Eth1/2(P)         Eth1/3(P)
..
31     Po31(SU)     Eth       LACP           Eth2/1(P)         Eth2/6(H)
..
DC3-N7K-AGG1#..

```

**On Agg2:**

```

DC3-N7K-AGG2#show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        S - Switched      R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports      Channel
-----
1      Po1(SU)      Eth       LACP           Eth1/2(P)         Eth1/3(P)
..
31     Po31(SU)     Eth       LACP           Eth2/2(P)         Eth2/3(H)
.....

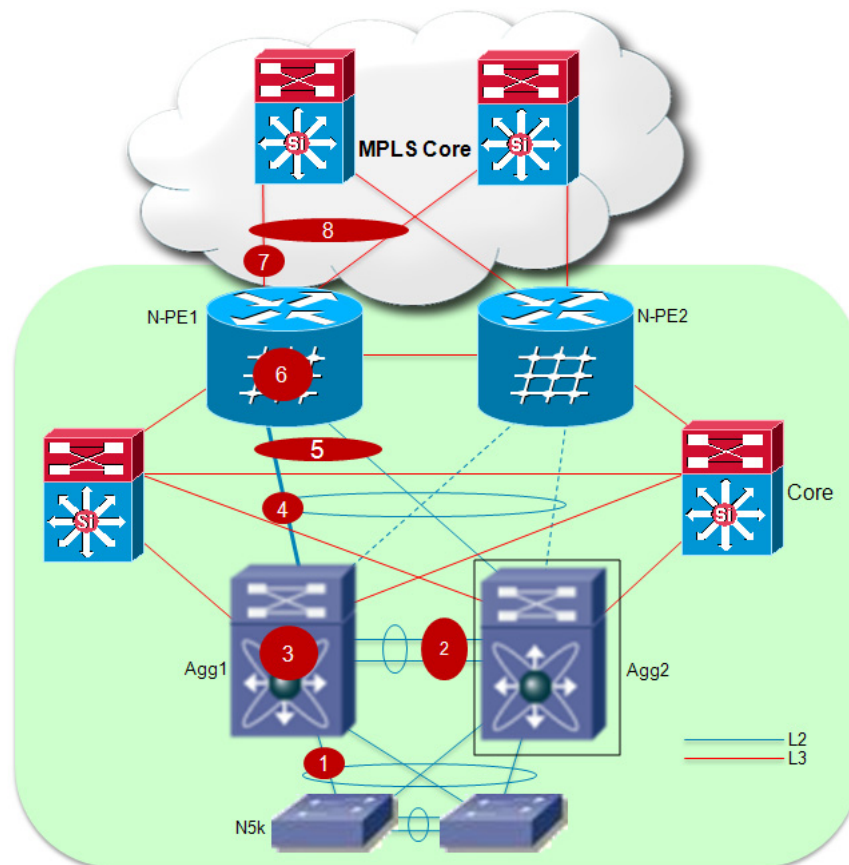
```

## Convergence Tests

Convergence testing was performed to measure convergence times for unicast and multicast traffic during various link and node failures. Convergence was measured from the data source to the receiver (end-to-end network convergence) by determining packet loss for each flow. For example, a packet rate of 1000 packets per second (pps) corresponds to 1-millisecond (ms) convergence time for each packet dropped.

Figure 5-2 provides a view of the different failure and recovery scenarios that were validated.

Figure 5-2 Failure/Recovery Test Cases



Each failure and recovery scenario will be analyzed, discussing the mechanisms leveraged to recover traffic flows and presenting the specific test results achieved during the validation effort. Test results will be presented in the two specific scenarios where 500 and 1200 VLANs were extended by leveraging two separate MC-LAG groups, as shown in Figure 2-1. For the 500 VLANs extension case, each MC-LAG group was carrying 250 VLANs; for the 1200 VLANs extension case, one MC-LAG group (the one to which failure/recovery scenarios were applied) was carrying 1000 VLANs, whereas the other one was carrying 200 VLANs.

To better understanding the convergence results achieved, it is important to describe the various traffic flows that were used during testing, distinguishing the cases where 500 VLANs or 1200 VLANs were extended.

#### 500 VLANs

- **L2: Intra-VLAN-100-349:** these are pure L2 traffic flows (250 VLANs) extended between DC2 and DC3 sites leveraging the first MC-LAG group. Most of the failure/recovery scenarios directly affected links and devices relative to this MC-LAG group, so it is expected to notice traffic outage.
- **L2: Intra-VLAN-1200-1449:** these are pure L2 traffic flows (250 VLANs) extended between DC2 and DC3 sites leveraging the second MC-LAG group. In most of the validated failure scenarios, traffic flowing across this connection should remain unaffected. This is because most of the link/node failures are performed to only affect the connections forming the first MC-LAG group.
- **L3L2: Inter-VLAN:** these are traffic flows that are first routed at the aggregation layer between a source and a destination VLAN. The destination VLAN is then extended to the remote site, so the behavior is expected to be very similar to the pure L2 traffic flows mentioned above.

- **L3:Inter-VLAN:** these are pure routed flows that leverage the dedicated L3 connections established between the aggregation and core layre devices. Most of the failire scenarios affecting the MC-LAG connection should not affect these flows.
- **Multicast L2: Intra-VLAN:** these are 100 multicast flows (one in each unique VLAN) characterized by having 1 source and 1 receiver in the same VLAN in separate data center site.

**Note**

L3 multicast was not the main focus so convergence numbers are not specifically referenced. However, it was validated that L3 multicast stream always recovered in all failure scenarios.

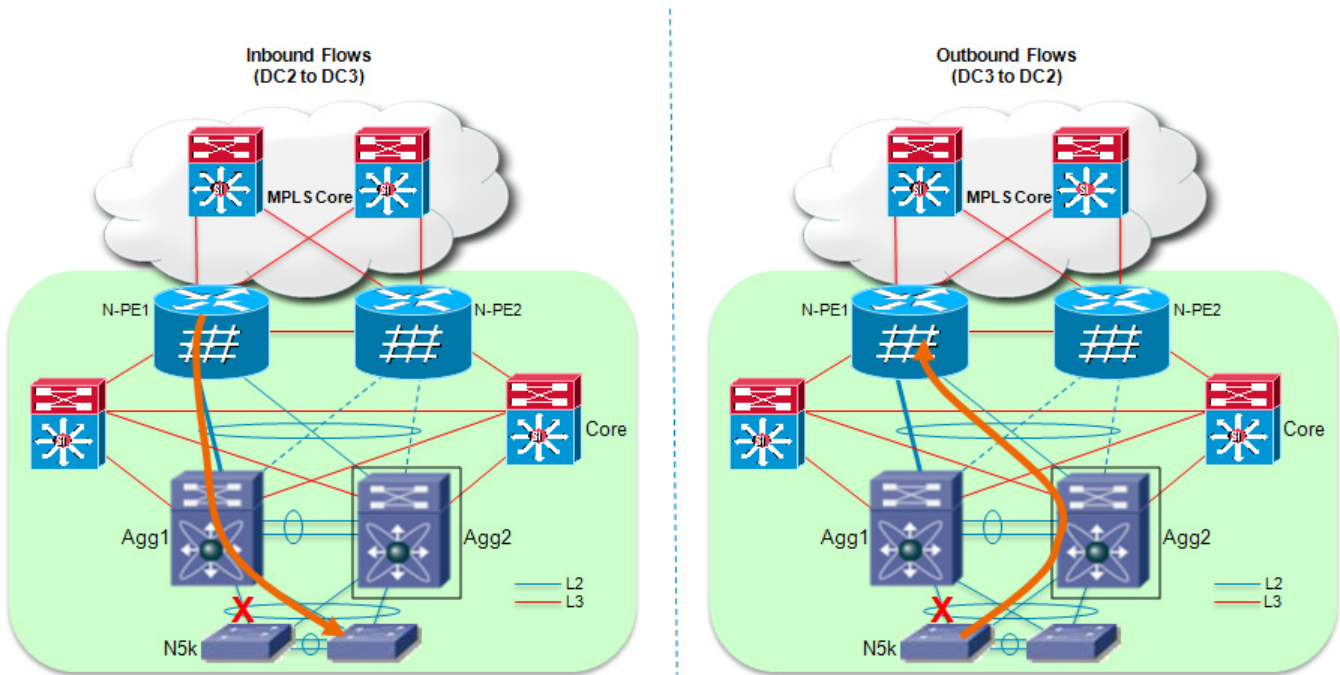
**1200 VLANs**

- **L2:Intra-VLAN-100-349:** these are pure L2 traffic flows (250 VLANs) carried between a specific access layer device and the Nexus 7000 in aggregation. These flows are then combined with the ones described below on the MC-LAG connection between Nexus 7000 switches and PE routers (in order to achieve a total of 1000 VLANs on that connection).
- **L2:Intra-VLAN-350-999-1100-1199:** these are pure L2 traffic flows (750 VLANs) carried between another access layer device and the aggregation switches. As mentioned above, these flows are then carried on the main MC-Lag connection under test toward the PE routers.
- **L2:Intra-VLAN-1200-1399:** these are pure L2 traffic flows (200 VLANs) extended between DC2 and DC3 sites leveraging the second MC-LAG group. In most of the validated failure scenarios, traffic flowing across this connection should remain unaffected.
- **L3L2: Inter-VLAN:** these are traffic flows that are first routed at the aggregation layer between a source and a destination VLAN. The destination VLAN is then extended to the remote site, so the behavior is expected to be very similar to the pure L2 traffic flows mentioned above.
- **L3:Inter-VLAN:** these are pure routed flows that leverage the dedicated L3 connections established between the aggregation and core layre devices. Most of the failire scenarios affecting the MC-LAG connection should not affect these flows.

## Test 1: Access to Aggregation Uplink Failure and Recovery

When the uplink between access and aggregation layer devices fails, the traffic recovering mechanism is the same independently from the direction of the traffic ([Figure 5-3](#)).

Figure 5-3 Failure/Recovery Test Cases



Inbound flows (originated in the remote site DC2 and destined to local DC3) received by the left aggregation device are switched via the remaining link connecting to the access layer.

Outbound flows (originated in the local DC3 site and destined to remote DC2) that were originally sent via the failed uplink need also to be shifted to the remaining available uplink connecting to the right aggregation switch. The access layer device is responsible for detecting the uplink failure event and performing etherchannel re-hashing. Traffic flows that were originally hashed to the right uplinks are unaffected by the failure.



**Note** The opposite behavior is required when the uplink is recovered.

The convergence results achieved when extending 500 VLANs are shown in Table 5-4.

Table 5-4 Test 1 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Access to aggregation uplink failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.099	0.124
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.099	0.124
		L3:Inter-Vlan	0	0
	no shut	L2:Intra-Vlan-100-349	1.453	1.381
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	1.453	1.381
		L3:Inter-Vlan	0	0

Table 5-4 Test 1 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Access to aggregation uplink failure & recovery <b>Multicast Traffic</b>	shut	Multicast L2: Intra-Vlan	0.099	0.188
	no shut	Multicast L2: Intra-Vlan	1.453	1.423

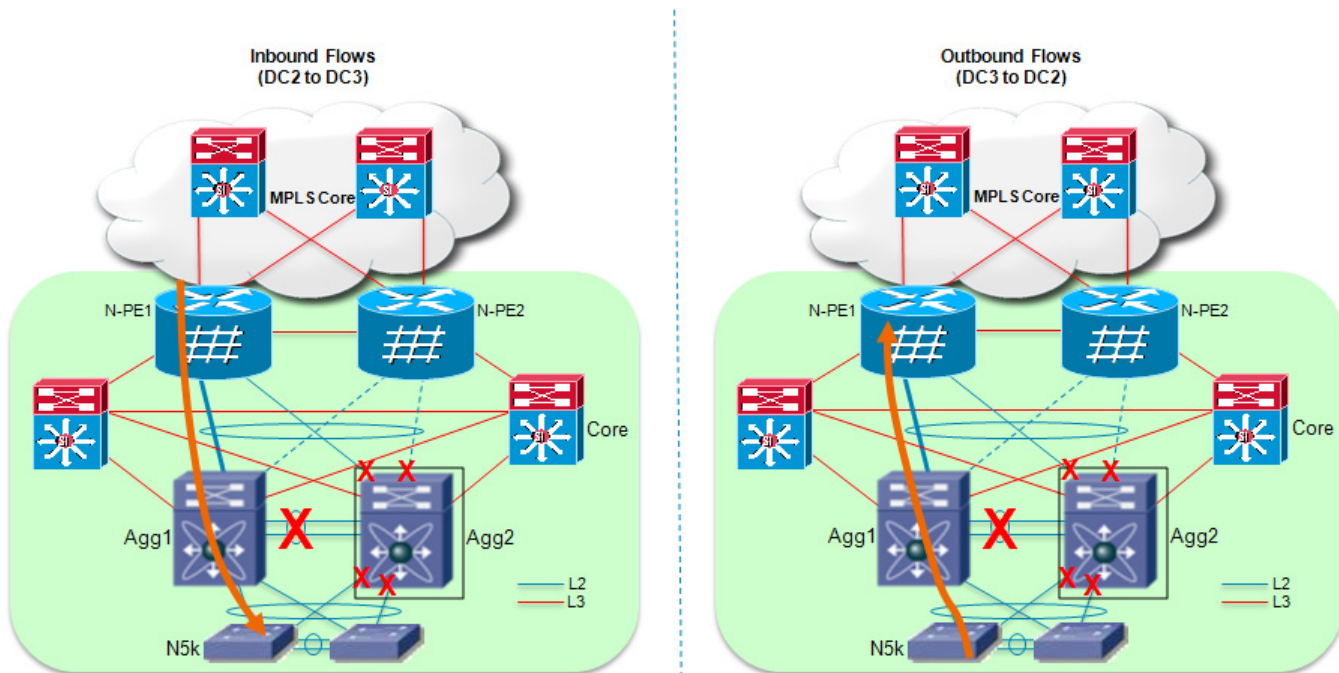
**Note**

Convergence results for 1200 VLANs are not available for this specific failure scenario because Nexus 5000 can support up to 512 VLANs in the software release that was validated as part of this solution.

## Test 2: Complete vPC Peer-Link Failure and Recovery

As a consequence of the vPC peer-link failure, assuming that the two aggregation devices can still communicate via the peer-keepalive link, the device operating in “vPC secondary” role (the right aggregation switch in Figure 5-4) would bring down all the physical interface part of configured vPCs.

Figure 5-4 Complete vPC Peer-Link Failure Scenario



Inbound flows will be re-hashed by the ASR 9000 PE1 device on the remaining links connecting to the left Nexus 7000. Outbound traffic flows will be re-hashed by the access layer devices on the remaining uplinks (similarly to what discussed in Test 1).

When the peer-link connection is re-established, the vPC secondary device will re-enable the vPC physical links and traffic will start flow again on these connections.

The results shown in Table 5-5 and Table 5-6 highlight how also pure L3 traffic is impacted in this scenario. This is expected given that no physical path is available after the peer-link failure and re-routing of inbound L3 traffic is required at the DC Core layer.



**Table 5-5** Test 2 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Complete vPC peer link failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.011	0.126
		L2:Intra-Vlan-1200-1449	0	0.271
		L3L2: Inter-Vlan	0.011	0.271
		L3:Inter-Vlan	1.767	0.215
	no shut	L2:Intra-Vlan-100-349	1.162	0.529
		L2:Intra-Vlan-1200-1449	0	0.534
		L3L2: Inter-Vlan	1.162	0.534
		L3:Inter-Vlan	2.065	0.535
Complete vPC peer link failure and recovery <b>Multicast Traffic</b>	shut	Multicast L2:Intra-Vlan	0.011	0.271
	no shut	Multicast L2:Intra-Vlan	1.160	0.534

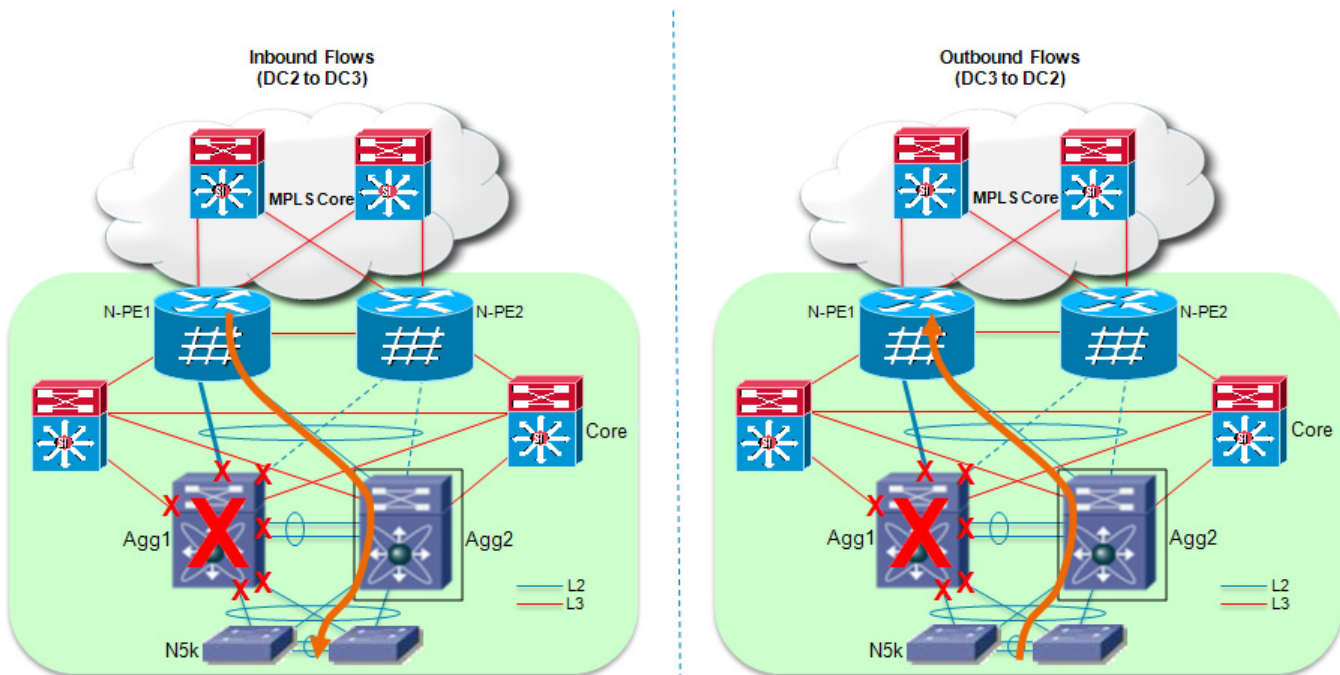
**Table 5-6** Test 2 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Complete vPC peer link failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0	0.141
		L2:Intra-Vlan-350-999-1100-1199	0	0.239
		L2:Intra-Vlan-1200-1399	0.010	0.262
		L3L2:Inter-Vlan	0.010	0.262
		L3:Inter-Vlan	5.636	0.228
	no shut	L2:Intra-Vlan-100-349	0	2.335
		L2:Intra-Vlan-350-999-1100-1199	0	2.063
		L2:Intra-Vlan-1200-1399	3.164	2.065
		L3L2:Inter-Vlan	3.163	2.335
		L3:Inter-Vlan	2.797	2.066

## Test 3: Aggregation Device Failure and Recovery

When one of the two aggregation Nexus 7000 device fails, all the physical connections to and from that device are obviously torn down as well. Inbound and outbound traffic paths after the failure are shown in [Figure 5-5](#).

Figure 5-5 Aggregation Device Failure Scenario



Inbound flows are re-hashed by the ASR 9000 PE device on the remaining active links connected to the right aggregation layer device (this recovery mechanism is similar to the one discussed for the vPC peer-link failure scenario). Outbound flows are instead re-hashed by the access layer switches, similarly to how noticed for the access to aggregation uplink failure test case.

The physical failure of the aggregation device would obviously impact also the pure L3 routed flows, since the dedicated L3 connections will also fail.

Table 5-7 Test 3 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Aggregation device failure and recovery <b>Unicast Traffic</b>	reload	L2:Intra-Vlan-100-349	0.018	0.123
		L2:Intra-Vlan-1200-1449	0.013	0.228
		L3L2: Inter-Vlan	0.018	0.228
		L3:Inter-Vlan	0.048	0.228
	restore	L2:Intra-Vlan-100-349	0.029	1.837
		L2:Intra-Vlan-1200-1449	1.894	0.751
		L3L2: Inter-Vlan	1.893	1.837
		L3:Inter-Vlan	1.284	0.749
Aggregation device failure and recovery <b>Multicast Traffic</b>	reload	Multicast L2: Intra-Vlan	0.018	0.228
	restore	Multicast L2:Intra-Vlan	1.891	1.881

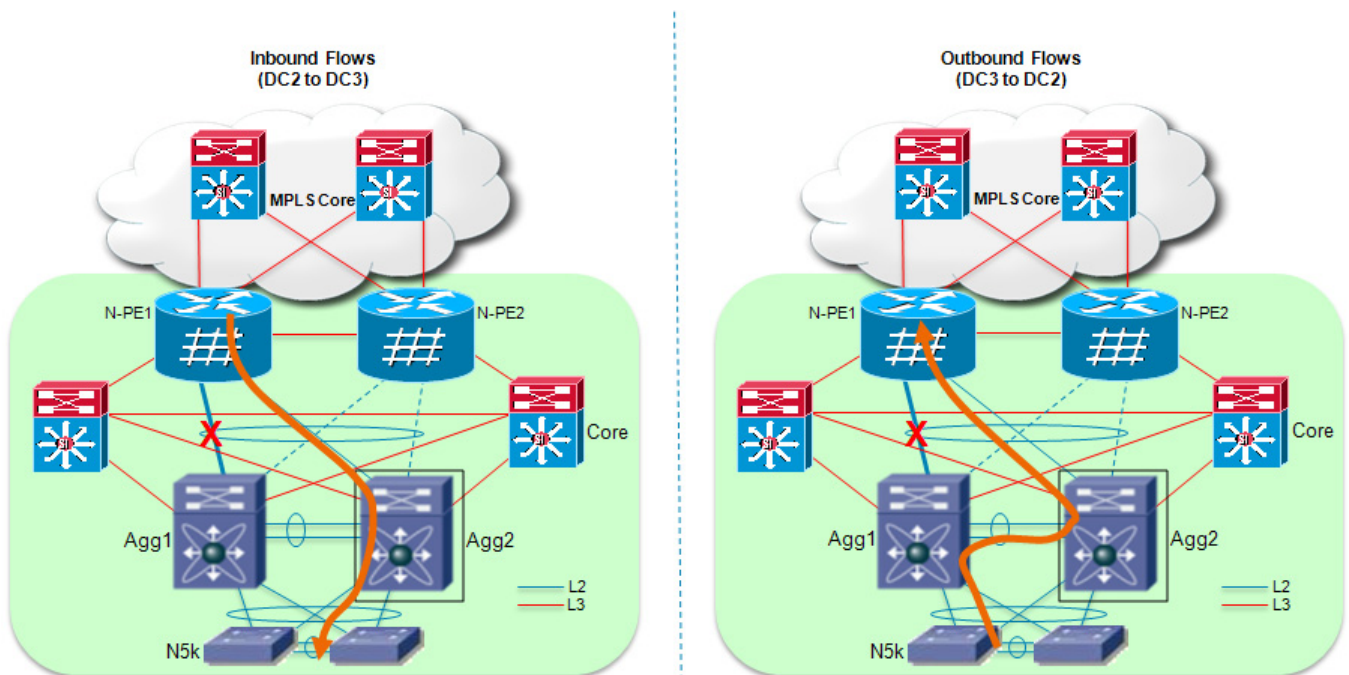
**Table 5-8** Test 3 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Aggregation device failure and recovery <b>Unicast Traffic</b>	reload	L2:Intra-Vlan-100-349	0.017	0.125
		L2:Intra-Vlan-350-999-1100-1199	0	0.250
		L2:Intra-Vlan-1200-1399	0	0.250
		L3L2:Inter-Vlan	0.017	0.250
		L3:Inter-Vlan	0.059	0.250
	restore	L2:Intra-Vlan-100-349	0.023	3.986
		L2:Intra-Vlan-350-999-1100-1199	0	0
		L2:Intra-Vlan-1200-1399	3.245	3.247
		L3L2:Inter-Vlan	0.023	3.985
		L3:Inter-Vlan	4.625	4.606

### Test 4: Aggregation to PE Active Link Failure and Recovery

The failure of an active L2 link between aggregation and PE device would cause the same recovery mechanism for inbound traffic flows already discussed in the previous failure scenario, where the ASR 9000 PE router is responsible to re-hash the flows on the remaining active link connecting to the second aggregation switch (Figure 5-6).

**Figure 5-6** Aggregation Device Failure Scenario



Outbound flows that were originally sent on the failed link need instead to be switched via the peer-link connecting the aggregation devices. Once again, this should be taken into consideration when designing the bandwidth to be dedicated to this connection.

**Table 5-9 Test 4 results with 500 VLANs (Unicast and Multicast)**

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Aggregation to PE active link failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0	0.658
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0	0.658
		L3:Inter-Vlan	0	0
	no shut	L2:Intra-Vlan-100-349	0	0.553
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0	0.553
		L3:Inter-Vlan	0	0
Aggregation to PE active link failure and recovery <b>Multicast Traffic</b>	shut	Multicast L2:Intra-Vlan	0.011	0.657
	no shut	Multicast L2:Intra-Vlan	0.167	0.553

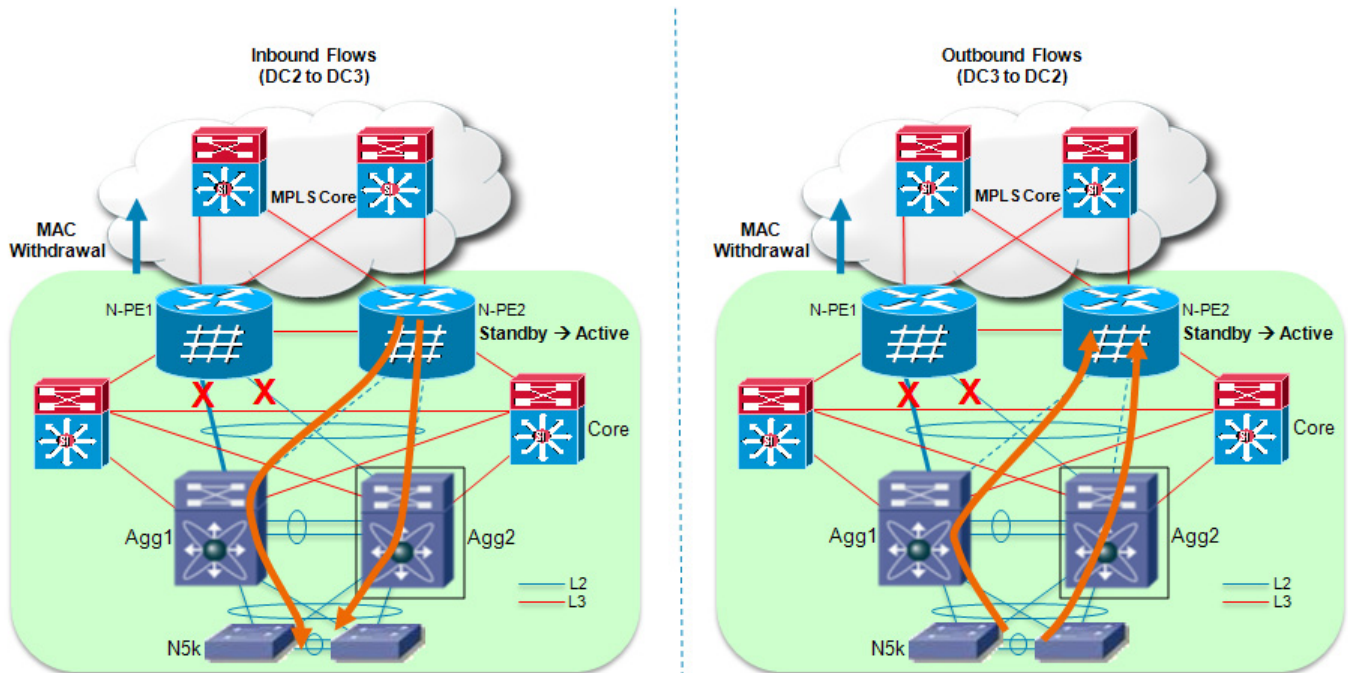
**Table 5-10 Test 4 results with 1200 VLANs (Unicast Traffic)**

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Aggregation to PE active link failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.011	0.849
		L2:Intra-Vlan-350-999-1100-1199	0.011	0.849
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	0.010	0.849
		L3:Inter-Vlan	0	0
	no shut	L2:Intra-Vlan-100-349	1.223	2.086
		L2:Intra-Vlan-350-999-1100-1199	1.231	2.094
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	1.223	2.086
		L3:Inter-Vlan	0	0

## Test 5: PE Dual Active Links to Aggregation Failure and Recovery

The contemporary failure of both links connecting the active PE to the aggregation layer devices is the first scenario where the standby PE needs to become active to allow for L2 traffic flows recovery. In order for that to happen, the initially active PE communicate via ICCP to its peer that it loses both active links, forcing the standby router to become active (Figure 5-7).

Figure 5-7 PE Dual Active Links to Aggregation Failure Scenario



Inbound traffic recovery is mainly dictated by two mechanisms:

- Time required for the standby PE to transition to an active role. This would cause (through LACP negotiation) moving the interfaces connecting the aggregation devices to this PE from an Hot-Standby to an Active state.
- Time required for the remote PE devices to flush their MAC address tables (as a result of a MAC notification originated by the local PE) to ensure that traffic can now be flooded and can reach the newly activated PE router. Flooding will stop once bidirectional communication is established between sites, and the remote PE routers correctly populate the information in their MAC address tables.

Outbound traffic outage is mostly dictated by the first item discussed above (i.e. the time required to transition hot-standby interfaces to an active role).



**Note**

To minimize the occurrence of this failure scenario, Cisco recommends spreading the links connecting each PE router to the aggregation switched on different linecards.

**Table 5-11** Test 5 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
PE dual active links to aggregation failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.378	0.374
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.378	0.374
		L3:Inter-Vlan	0	0
	no shut	L2:Intra-Vlan-100-349	0.592	0.322
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.591	0.320
		L3:Inter-Vlan	0	0
PE dual active links to aggregation failure and recovery <b>Multicast Traffic</b>	shut	Multicast L2:Intra-Vlan	0.210	0.443
	no shut	Multicast L2:Intra-Vlan	3.269	0.193

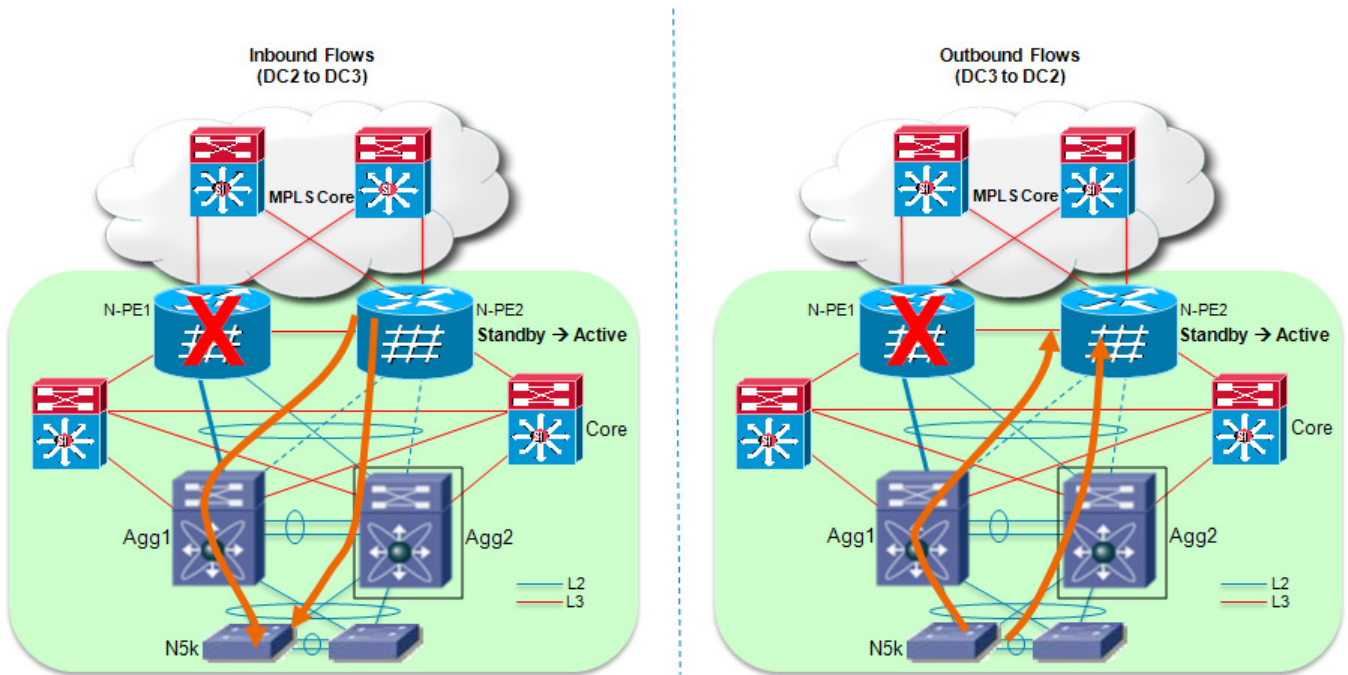
**Table 5-12** Test 5 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
PE dual active links to aggregation failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.693	0.668
		L2:Intra-Vlan-350-999-1100-1199	0.693	0.668
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	0.693	0.668
		L3:Inter-Vlan	0	0
	no shut	L2:Intra-Vlan-100-349	1.994	1.145
		L2:Intra-Vlan-350-999-1100-1199	2.877	2.303
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	1.992	1.142
		L3:Inter-Vlan	0	0

## Test 6: Active PE Router Failure and Recovery

The failure of the active PE router is a second scenario that will force the standby PE to transition to active state.

Figure 5-8 Active PE Router Failure Scenario



The recovery for inbound traffic is similar to what discussed in the previous test case. The main difference is that the remote PE routers will start flooding traffic directed to DC3 not because of the reception of a MAC withdrawal notification, but as a consequence of the fact the PWs connected to the failed router are brought down. The end result is the same, with traffic being received by the newly activated PE router, which will then send it toward the aggregation layer switches as soon as the physical links are transitioned from Hot-Standby to Active state.

Outbound traffic recovery once again is dependant on the activation of the standby links only.

Table 5-13 Test 6 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE router failure and recovery <b>Unicast Traffic</b>	reload	L2:Intra-Vlan-100-349	0.973	0.727
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.972	0.726
		L3:Inter-Vlan	0	0
	restore	L2:Intra-Vlan-100-349	0.598	0.327
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.598	0.324
		L3:Inter-Vlan	0.098	0
Active PE router failure and recovery <b>Multicast Traffic</b>	reload	Multicast L2: Intra-Vlan	0	0.075
	restore	Multicast L2:Intra-Vlan	4.986	0.209

**Table 5-14** Test 6 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE router failure and recovery Unicast Traffic	reload	L2:Intra-Vlan-100-349	2.756	2.068
		L2:Intra-Vlan-350-999-1100-1199	2.985	2.518
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	2.755	2.068
		L3:Inter-Vlan	0	0
	restore	L2:Intra-Vlan-100-349	1.667	1.039
		L2:Intra-Vlan-350-999-1100-1199	2.816	2.334
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	1.666	1.937
		L3:Inter-Vlan	0	0

## Test 7: Active PE Router Core Link Failure and Recovery

The failure of one of the routed link connecting a PE router to the MPLS core can be recovered by routing the VPLS traffic on the alternate L3 path available.

For inbound flows, the re-routing happens in the MPLS core (since traffic is destined to the PE loopback interface used to establish the LDP session with the remote PEs).

For outbound flows, the re-routing happens locally on the PE affected by the link failure, since traffic is destined to the loopback interface of a remote PE device.

In both cases, only the flows that were sent via the failed link are affected, whereas the traffic originally routed via the remaining L3 uplink continues to flow undisturbed.



Figure 5-9 Active PE Router Core Link Failure Scenario

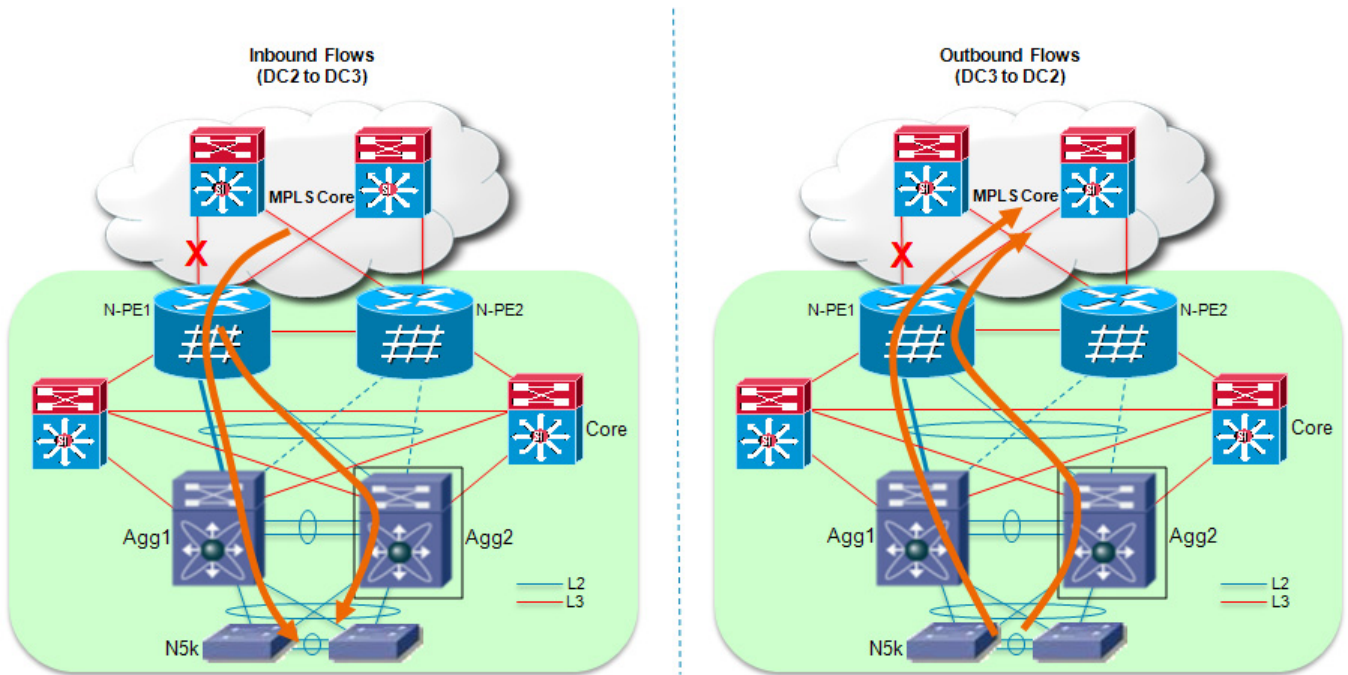


Table 5-15 Test 7 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE router core link failure and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.241	0.246
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.241	0.246
		L3:Inter-Vlan	0.141	0.299
	no shut	L2:Intra-Vlan-100-349	0	0
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0	0
		L3:Inter-Vlan	0	0
Active PE router core link failure and recovery <b>Multicast Traffic</b>	shut	Multicast L2:Intra-Vlan	0.241	1.613
	no shut	Multicast L2:Intra-Vlan	0	1.542

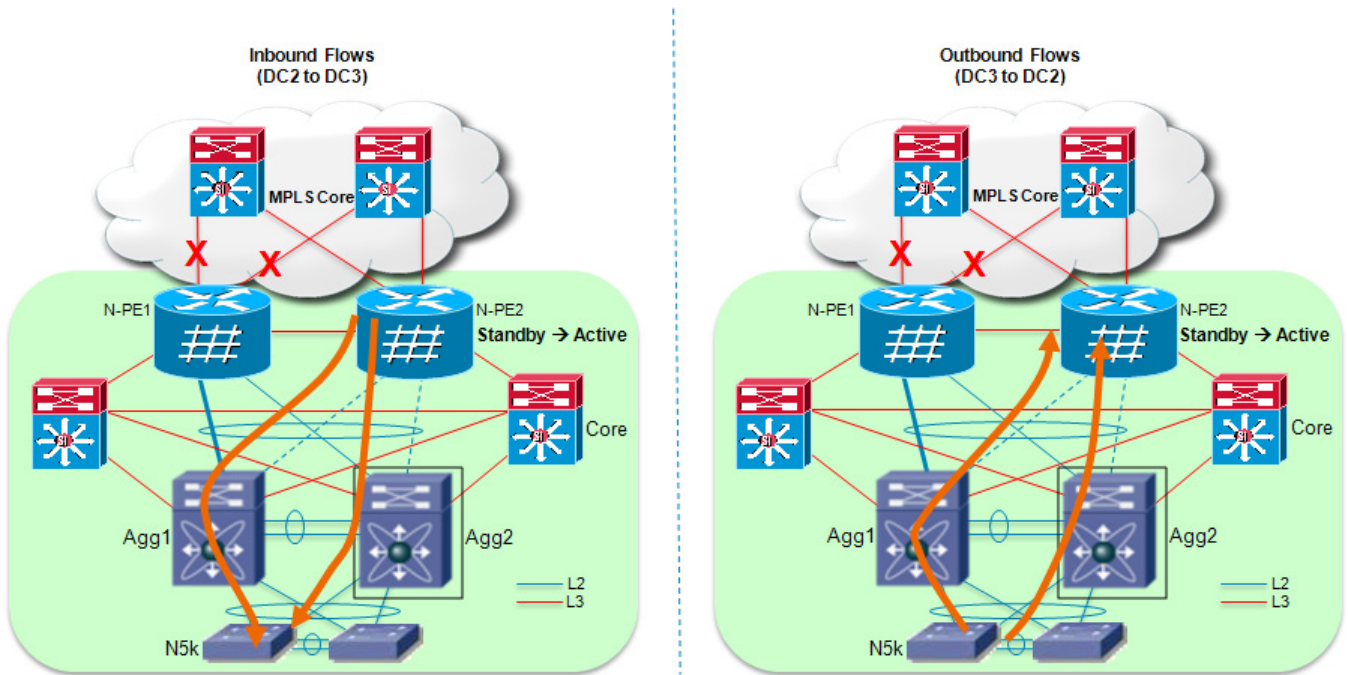
Table 5-16 Test 7 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE router core link failure and recovery Unicast Traffic	shut	L2:Intra-Vlan-100-349	0.243	0.238
		L2:Intra-Vlan-350-999-1100-1199	0.243	0.238
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	0.243	0.238
		L3:Inter-Vlan	0.132	0.297
	no shut	L2:Intra-Vlan-100-349	0	0
		L2:Intra-Vlan-350-999-1100-1199	0	0
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	0	0
		L3:Inter-Vlan	0	0

## Test 8: Active PE Core Isolation and Recovery

If the active PE loses both direct connections to the MPLS core, a “core isolation” recovery is triggered forcing the other PE router to transition to the Active state. The behavior is the result of a specific configuration applied on the active PE router when defining an MC-LAG group: the interfaces used to connect to the core are explicitly defined as “backbone” interfaces. When they both fail, the PE router leverages ICCP to communicate the event to its peer, which will get activated.

Figure 5-10 Active PE Core Isolation Scenario



Similarly to how discuss Test 6 (PE failure scenarios), traffic recovery in both inbound and outbound directions is mainly dependant on how fast the hot-standby links can be activated.

Table 5-17 Test 8 results with 500 VLANs (Unicast and Multicast)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE core isolation and recovery <b>Unicast Traffic</b>	shut	L2:Intra-Vlan-100-349	0.602	0.598
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.602	0.598
		L3:Inter-Vlan	0.132	0.297
	no shut	L2:Intra-Vlan-100-349	0.534	0.349
		L2:Intra-Vlan-1200-1449	0	0
		L3L2: Inter-Vlan	0.534	0.349
		L3:Inter-Vlan	0	0.001
Active PE core isolation and recovery <b>Multicast Traffic</b>	shut	Multicast L2: Intra-Vlan	0.210	0.443
	no shut	Multicast L2:Intra-Vlan	3.269	0.193

**Table 5-18** Test 8 results with 1200 VLANs (Unicast Traffic)

Failure Type	Action	Flows	DC2→DC3	DC3→DC2
Active PE core isolation and recovery Unicast Traffic	shut	L2:Intra-Vlan-100-349	1.936	1.602
		L2:Intra-Vlan-350-999-1100-1199	2.800	2.680
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	1.935	1.600
		L3:Inter-Vlan	0.133	0.282
	no shut	L2:Intra-Vlan-100-349	2.461	1.297
		L2:Intra-Vlan-350-999-1100-1199	2.852	2.324
		L2:Intra-Vlan-1200-1399	0	0
		L3L2:Inter-Vlan	2.460	1.284
		L3:Inter-Vlan	0	0

## Deployment Recommendations

This chapter addresses issues that you should consider when deploying MC-LAG based VPLS solution to interconnect data centers.

### 1. Improve multicast convergence number using “mrouter” command

By default on ASR 9000, ICCP process running between the two POAs only synchronizes IGMP entries from the access side within a data center to mLACP standby peer. Also, mrouter ports are dynamically learned via multicast protocols. This adds delay under failure conditions during mLACP switchover.

Network convergence time for multicast traffic can be improved by configuring “mrouter” on ASR 9000. This command does the following:

- Statically configures an interface as mrouter port which otherwise has to rely on PIM or IGMP for dynamic learning
- Allows synchronization of IGMP entries from the MPLS core to N-PE in mLACP standby mode

This recommendation requires ASR 9000 as the edge router on all data centers.

Below is the relevant configuration. Configuring mrouter is a two-step process:

- Create an IGMP profile
- Configure mrouter under this profile. The second step is to apply this IGMP profile to all neighbors under VFs terminating in remote datacenters.



#### Note

Do not apply this profile to the pseudowire between the NPE's within the same data center.

#### mrouter configuration on ASR 9000 router

```
RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1#conf t
RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1(config)#igmp snooping profile igmp-mrouters
RP/0/RSP0/CPU0:DC3-ASR9(config-igmp-snooping-profile)#mrouter
RP/0/RSP0/CPU0:DC3-ASR9(config-igmp-snooping-profile)#commit
```

```

RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1(config)#l2vpn bridge group group1
RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1(config-l2vpn-bg)#bridge-domain vlan107
RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1(config-l2vpn-bg-bd)#vfi vfi107
RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1(config-l2vpn-bg-bd-vfi)#neighbor 150.3.3.6 pw-id 107
RP/0/RSP0/CPU0:DC3-ASR 9000-N(config-l2vpn-bg-bd-vfi-pw)#igmp snooping profile
igmp-mrouters
RP/0/RSP0/CPU0:DC3-ASR 9000-N(config-l2vpn-bg-bd-vfi-pw)#commit

RP/0/RSP0/CPU0:DC3-ASR 9000-NPE1#sh run l2vpn bridge group group1 bridge-domain vlan107
Thu Mar 10 12:35:21.624 PST
l2vpn
bridge group group1
bridge-domain vlan107
igmp snooping profile igmp-snoop
interface Bundle-Ether31.107
!
vfi vfi107
neighbor 150.2.2.5 pw-id 107
pw-class vpls-pw-class
igmp snooping profile igmp-mrouters
!
neighbor 150.2.2.6 pw-id 107
pw-class vpls-pw-class
igmp snooping profile igmp-mrouters
!
# Neighbor 150.3.3.6 is the N-PE (ASR9000) within the same data center. Hence IGMP mrouter
profile is not configured #

neighbor 150.3.3.6 pw-id 107
pw-class vpls-pw-class
!
neighbor 150.11.11.5 pw-id 107
pw-class vpls-pw-class
igmp snooping profile igmp-mrouters
!
neighbor 150.11.11.6 pw-id 107
pw-class vpls-pw-class
igmp snooping profile igmp-mrouters
!

```

As shown in [Figure 4-2](#), Cisco 7600 routers deployed as PEs in DC2 do not support syncing of IGMP entries either from access or from core. Due to this, IGMP entries have to be relearned during mLACP switchover. In this scenario, network convergence for multicast traffic depends on the rate at which mrouter ports are dynamically learned which is a factor of PIM and IGMP timers.

**Note**

This issue is going to be fixed in the upcoming 7600 router 15.2(1)S release planned for July.

## 2. Avoid IGMP packets looping under specific mLACP failure conditions

In case of deploying multiple aggregation blocks connected to the pair of PE devices and extending the same set of VLANs, a local pseudowire must be established between the PEs and this may induce IGMP packets looping. To avoid this problem, the recommendation is to configure the “router-guard” command under IGMP profile and assign that IGMP profile to the pseudowire between the NPE’s within the same datacenter.

As shown below, configuring router-guard is a two-step process. First an IGMP profile has to be created, and the router-guard command should be configured under this profile. The second step is to apply this IGMP profile to all neighbors under VFIs between the NPE’s within the same datacenter.

**router-guard configuration on ASR 9000 router**

```

RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#conf t
Mon Mar 14 15:14:01.865 PST
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config)#igmp snooping profile router-guard
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-igmp-snooping-profile)#router-guard
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-igmp-snooping-profile)#commit

RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#conf t
Mon Mar 14 15:14:23.503 PST
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config)#l2vpn bridge group group1 bridge-domain vlan111
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-l2vpn-bg-bd)#vfi vfi111
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-l2vpn-bg-bd-vfi)#neighbor 150.3.3.6 pw-id 111
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-l2vpn-bg-bd-vfi-pw)#igmp snooping profile router-guard
RP/0/RSP0/CPU0:DC3-ASR9K-NPE1(config-l2vpn-bg-bd-vfi-pw)#commit

RP/0/RSP0/CPU0:DC3-ASR9K-NPE1#sh run l2vpn bridge group group1 bridge-domain v$
Mon Mar 14 15:22:52.099 PST
l2vpn
bridge group group1
  bridge-domain vlan111
    igmp snooping profile igmp-snoop
    interface Bundle-Ether31.111
      !
      vfi vfi111
        neighbor 150.2.2.5 pw-id 111
          pw-class vpls-pw-class
          !
        neighbor 150.2.2.6 pw-id 111
          pw-class vpls-pw-class
          !
      !
    !
  !
!
# Neighbor 150.3.3.6 is the N-PE (ASR9000) within the same data center. Hence IGMP profile
with router-guard is configured #

neighbor 150.3.3.6 pw-id 111
  pw-class vpls-pw-class
  igmp snooping profile router-guard
  !
neighbor 150.11.11.5 pw-id 111
  pw-class vpls-pw-class
  !
neighbor 150.11.11.6 pw-id 111
  pw-class vpls-pw-class

```

## Summary

Globalization, security and disaster recovery considerations are driving divergence in business locations across multiple regions. In addition, organizations are looking to distribute workload between computers, share network resources effectively and increase the availability of applications.

As data centers grow in size and complexity, enterprises are adopting server virtualization technologies to achieve increase efficiency and use of resources. Due to the exponential growth, most of these customers are looking at interconnecting more data centers, extending large number of VLANs with high layer 2 traffic capabilities between these data centers.

This design guide describes the deployment of the MC-LAG to VPLS technology on Cisco ASR 9000 routers. While active/standby by nature, an option to provide active/active connectivity was also discussed thus providing greater redundancy and VLAN load sharing between the two POA devices. With the deployment of vPC on Nexus 7000 switches in the aggregation, this solution provides link and chassis level redundancy and faster convergence during link and node failures. In addition, the solution is also fully compatible with the use of VSS technology on Cisco Catalyst 6500s in the aggregation layer.

In summary, the MC-LAG based VPLS solution documented in this design guide provides a high-speed, low latency network with STP isolation between data centers. The solution is extremely flexible and highly scalable and offers key features required for large-scale data center interconnectivity.

