



VPLS and EoMPLS Based DCI Solution with nV Edge and vPC

July 11, 2012

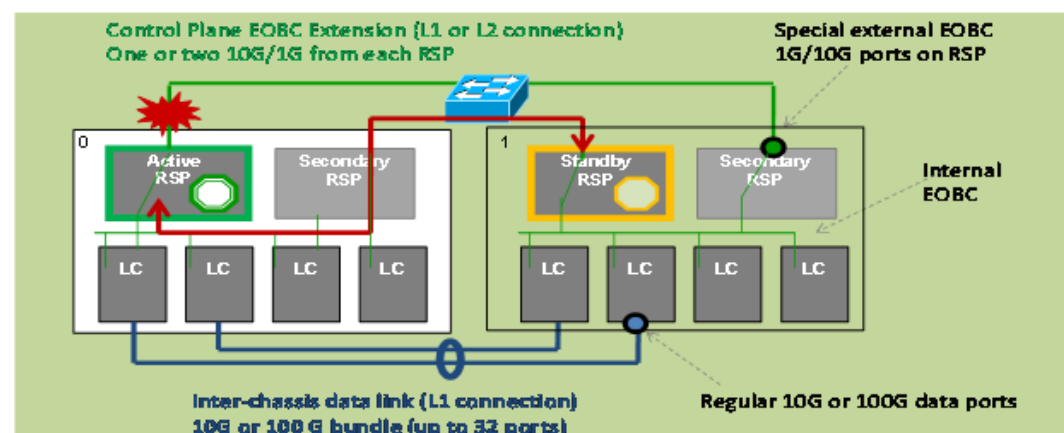


VPLS and EoMPLS Based DCI Solution with nV Edge and vPC

Data Center Interconnect (DCI) is an important and integral part of the Virtual Multi-Services Data Center (VMDC) solution architecture. The VMDC 2.x release utilizes MPLS-based DCI solutions – EoMPLS and VPLS – in which the ASR 9000 PE leverages Multi Spanning Tree Access Gateways (MST-AG) and Multi Chassis Link Aggregation Group (MC-LAG) for Attachment Circuit (AC) redundancy, and the Nexus 7000 uses Virtual Port Channels (vPC) for redundancy. Since there is only one active AC between the PE and CPE, both MST-AG and MC-LAG only offer an Active/Standby failover solution. To address this problem, the ASR 9000 platform has developed a new feature called Network Virtualization (nV). In nV, the Edge/Cluster, in which a pair of ASR 9000 routers forms a virtual chassis, acts as a single unit, and provides Active/Active AC failover between the PE and CPE. The purpose of this document is to outline the ASR 9000 nV Edge system and perform Proof of Concept (PoC) testing showing how the nV Edge can be used in conjunction with vPC in the Nexus 7000 for both AC and Psuedowire (PW) redundancy in the VMDC DCI solution.

The ASR 9000 nV Edge is a virtual chassis solution where a pair of ASR 9000 routers acts as a single device by maintaining a single control plane, management plane, and a fully distributed data plane across two physical chassis. It extends the control plane through special 1/10G Ethernet Out-of-Band Channel (EOBC) ports in RSP 440, and the data plane through regular data ports. For redundancy, at least two control plane links and two data plane links are required. [Figure 1](#) shows how control and data planes are extended in the nV Edge system. The control plane is extended by interconnecting one or more EOBC ports in RSP 440 between two chassis, and the data plane is extended by interconnecting one or more 10G or 100G regular data ports between two chassis.

Figure 1 *nV Edge System Overview*




ASR 9000 Edge DCI Solution Benefits

The ASR 9000 nV Edge is a very simple solution offering the following advantages over other solutions:

1. Active/Active regular LAG for AC redundancy.
2. Single PW between ASR 9000 nV edge PEs.
3. LAG ensures link/node failure, in which PW is totally unaware of any failure and thus provides faster convergence in the range of 50 msec.
4. Provides PW independent convergence.

Figure 2 shows the advantages of the nV Edge system over MC-LAG which is considered as one of the best solutions. MC-LAG shows only one AC that is active at any given time between the CPE and the PE. Also, only one PW is active, out of the four available among the PEs. In contrast, the nV Edge system eliminates AC and PW redundancy requirements because it requires only one AC and one PW. All links within a bundle-interface between the CPE and PE, as well as all links between the PEs are always active and redundancy/availability is provided via Link Aggregation Control Protocol (LACP).

Figure 2 **Advantage of the ASR 9000 nV Edge Over MC-LAG**



As long as there is one active link within a link-bundle the corresponding PW or AC remains up. Also, individual link failure does not affect respective AC or PW. Though not shown in Figure 2, nV Edge offers similar advantages over the MST-AG solution.

VMDC DCI Solution with the ASR 9000 nV Edge System

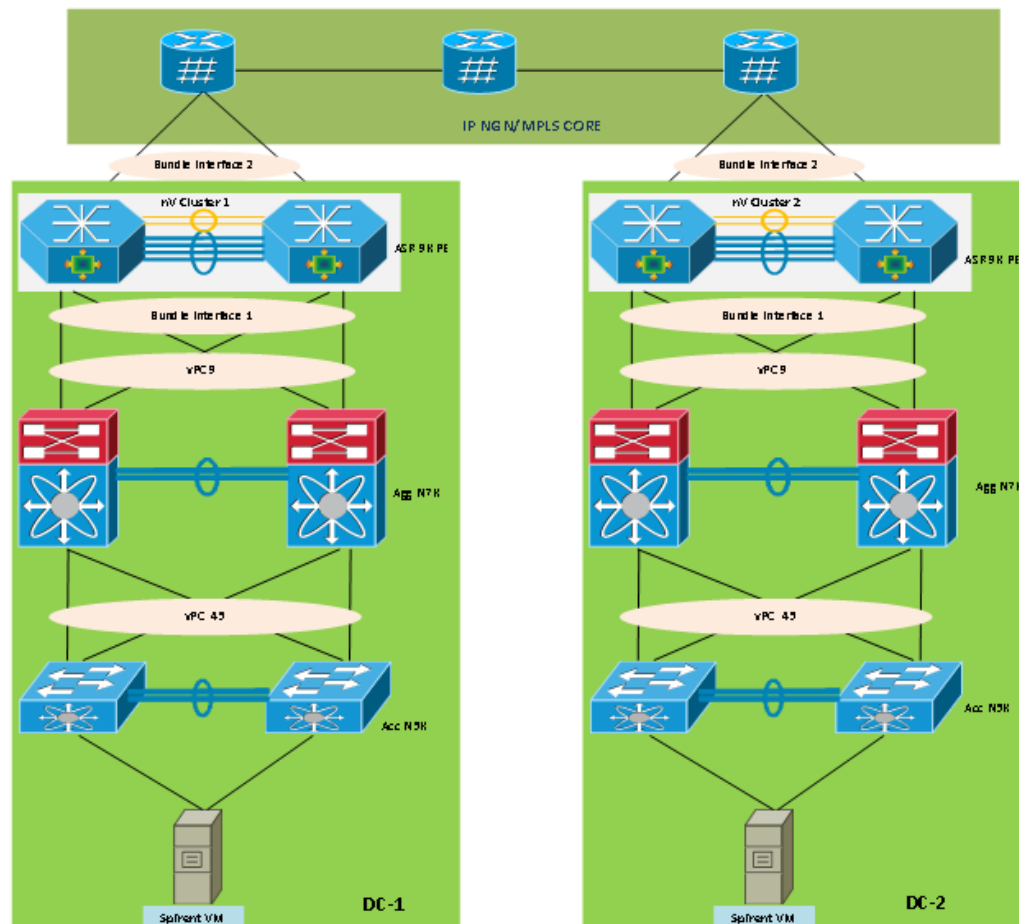
This PoC is based on VMDC 2.2 DCI architecture, where both EoMPLS and VPLS are used to extend VLANs between SP-SP and SP-ENT data centers. This PoC leverages the ASR 9000 nV Edge for PW and AC redundancy. In contrast, VMDC 2.2 uses MST-AG and MC-LAG for PW and AC redundancy. For detailed information, including configurations, for EoMPLS, VPLS, MST-AG, and MC-LAG, see the VMDC 2.2 DCI Design and Implementation Guide (DIG), available at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_vmvc.html#~overview

Test Bed Configuration

This PoC consists of two data centers, DC-1 and DC-2. These data centers are based on VMDC 2.2 DCI architecture with some modifications, as shown in Figure 3. VMDC 2.2 includes a pair of core Nexus 7000 switches, which are omitted from this PoC because core switches only perform Layer 3 (L3) functionality in VMDC 2.2 and are not necessary in Layer 2 DCI solutions. Additionally, a pair of Nexus 5000 switches is added in the access layer to connect UCS C-Series servers. Nexus 5000 switches are not used in the VMDC 2.2 access layer. The Nexus 1000V switch is also not used in this PoC.

Figure 3 Test Bed Topology



In [Figure 3](#), a vPC is configured between the Nexus 5000 and the Nexus 7000. A vPC is also configured between the aggregation layer Nexus 7000 and the ASR 9000 nV Edge. A port-bundle is configured on the ASR 9000 nV Edge, and sub-interfaces are created to extend VLANs between DC-1 and DC-2 by leveraging either VPLS or EoMPLS. VLANs from the access layer (i.e., from the Nexus 5000) are carried to the ASR 9000 nV Edge PE over Layer 2 (L2) trunks.

ASR 9000 nV Edge Solution Configuration

The ASR 9000 nV Edge system has hardware and software requirements that need to be met before it can be configured. The solution needs IOS XR release 4.2.1, and the nV Edge only supports RSP 440 and second generation line cards. For hardware support details, refer to the following URL:

http://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/interfaces/configuration/guide/hc42clst.html

The following steps are required to convert the two ASR 9000 chassis into an nV Edge system:

-
- Step 1** Upgrade the first chassis (i.e., rack 0) to IOS XR release 4.2.1.
 - Step 2** Configure the nV Edge on rack 0, including inter-chassis data plane ports (configuration details below).
 - Step 3** Upgrade the second chassis (i.e., rack 1) IOS XR release 4.2.1.
 - Step 4** Configure the nV Edge on rack including inter-chassis data plane ports (configuration details below).
 - Step 5** Add control plane links between rack 0 and rack 1.
-

Once these steps are satisfied, rack 0 and rack 1 will detect one another, and consequently, one becomes the primary Designated Shelf Controller (DSC), and the other chassis will reboot itself and join the nV Edge system. It will also upgrade the IOS XR image from the primary DSC if it does not match the primary DSC's image.

ASR 9000 nV Edge Configuration Procedure

The following procedure, using the appropriate corresponding IOS commands, is required to configure one of the chassis. The same commands are also needed to configure the other chassis, however, changing [Step 5](#) to reflect a different serial number and rack number, and [Step 10](#) for the inter-rack connection. Configuration examples are included.

-
- Step 1** Enter admin mode.

```
admin
RP/0/RSP0/CPU0:Cluster-01# admin
```
 - Step 2** Enter configuration mode.

```
configure
RP/0/RSP0/CPU0:Cluster-01(admin)# configure
```
 - Step 3** Enter nV Cluster configuration mode.

```
nv
RP/0/RSP0/CPU0:Cluster-01(admin-config)# nv
```

Step 4 Enter nV cluster /edge configure mode.

edge control

```
RP/0/RSP0/CPU0:Cluster-01(admin-config-nV)# edge control
```

Step 5 Assign a chassis serial number into a rack.

serial <#> rack <#>

```
RP/0/RSP0/CPU0:Cluster-01(admin-config-nV-Edge)# serial FOX1612GFFA rack 1
```

Step 6 Apply and save the configuration.

commit

```
RP/0/RSP0/CPU0:nV-C(admin-config-nV-Edge-Control)# commit
```

Step 7 (Optional) Exit configuration mode.

exit

```
RP/0/RSP0/CPU0:nV-C(admin-config-nV-Edge-Control)# exit
```

Step 8 Enter configuration mode.

configure

```
RP/0/RSP0/CPU0:Cluster-01# configure
```

Step 9 Enter interface configuration mode.

interface

```
RP/0/RSP0/CPU0:nV-Cluster-01(config)# interface tenGigE 0/0/1/3
```

Step 10 Define the interface for nV Edge inter-rack connection for data plane extension.

nv edge interface

```
RP/0/RSP0/CPU0:nV-Cluster-01(config-if)# nv edge interface
```

ASR 9000 nV Edge EoMPLS and VPLS DCI Configuration Example

The following sample configuration is used in VMDC 2.2 DCI nV Edge PoC testing:

```
RP/0/RSP0/CPU0:nV-Cluster-01(admin)#show running-config
```

```
Building configuration...
```

```
!! IOS XR Admin Configuration 4.2.1.32I
```

```
nv
```

```
edge
```

```
control
```

```
serial FOX1612GFFA rack 1
```

```
serial FOX1612GU9N rack 0
```

```
!
```

```
data
```

```
minimum 0 selected-interfaces
```

```
!
```

```
!
```

```
!
```

```
End
```

```
RP/0/RSP0/CPU0:nV-Cluster-01#sh running-config interface tenGigE 0/0/1/3
```

```

interface TenGigE0/0/1/3
nv
  edge
    interface
    !
  !
!

RP/0/RSP0/CPU0:ASR-nV-Cluster-01#sh running-config l2vpn
l2vpn
  xconnect group EoMPLS
    p2p CUST-02
      interface Bundle-Ether1.3
        neighbor 198.18.0.6 pw-id 101
      !
    !
  !
  bridge group bg-1
    bridge-domain bd-1
      interface Bundle-Ether1.1
      !
    vfi vpls-1
      vpn-id 1
      autodiscovery bgp
      rd 109:1
      route-target 109:1
      signaling-protocol ldp
      vpls-id 109:1
      !
    !
    neighbor 198.18.0.6 pw-id 100
    !
  !
!
!
RP/0/RSP0/CPU0:ASR-nV-Cluster-01#sh running-config interface bundle-ether 1.1
interface Bundle-Ether1.1 l2transport
encapsulation dot1q 461

RP/1/RSP0/CPU0:ASR-nV-Cluster-01#sh running-config interface bundle-ether 1.2
interface Bundle-Ether1.2 l2transport
encapsulation dot1q 101
!

```

ASR 9000 nV Edge EoMPLS and VPLS DCI Configuration Verification

The following configuration confirms that the nV Edge is up and in the Active failover state:

```

RP/0/RSP0/CPU0:ASR-nV-Cluster-01(admin)#show dsc
-----
Node   (Seq#)   Role      Serial#   State
-----
0/RSP0/CPU0   ( 0)   ACTIVE   FOX1612GU9N  PRIMARY-DSC
1/RSP0/CPU0   ( 3517533)  ACTIVE   FOX1612GFFA  BACKUP-DSC

```

```

RP/0/RSP0/CPU0:nV-Cluster-01#show platform
Sat Jun  9 00:30:45.851 UTC
Node           Type                               State           Config State

```



```

-----
0/RSP0/CPU0      A9K-RSP440-SE(Active)    IOS XR RUN      PWR,NSHUT,MON
0/0/CPU0         A9K-MOD80-SE             IOS XR RUN      PWR,NSHUT,MON
0/0/0            A9K-MPA-20X1GE           OK               PWR,NSHUT,MON
0/0/1            A9K-MPA-4X10GE           OK               PWR,NSHUT,MON
1/RSP0/CPU0      A9K-RSP440-SE(Active)    IOS XR RUN      PWR,NSHUT,MON
1/0/CPU0         A9K-MOD80-SE             IOS XR RUN      PWR,NSHUT,MON
1/0/0            A9K-MPA-20X1GE           OK               PWR,NSHUT,MON
1/0/1            A9K-MPA-4X10GE           OK               PWR,NSHUT,MON

```

RP/0/RSP0/CPU0:ASR-nV-Cluster-01#show l2vpn bridge-domain bd-name bd-1

Legend: pp = Partially Programmed.

Bridge group: bg-1, bridge-domain: bd-1, id: 0, state: up, ShgId: 0, MSTi: 0

Aging: 300 s, MAC limit: 4000, Action: none, Notification: syslog

Filter MAC addresses: 0

ACs: 1 (1 up), VFIs: 1, PWs: 1 (1 up), PBBs: 0 (0 up)

List of ACs:

BE1.1, state: up, Static MAC addresses: 0

List of Access PWs:

List of VFIs:

VFI vpls-1 (up)

Neighbor 198.18.0.10 pw-id 100, state: up, Static MAC addresses: 0

RP/1/RSP0/CPU0:ASR-nV-Cluster-01#show l2vpn xconnect group EoMPLS

Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

XConnect			Segment 1		Segment 2		
Group	Name	ST	Description	ST	Description		ST
EoMPLS	CUST-02	UP	BE1.2	UP	198.18.0.10	101	UP

Rabiul Hasan**Technical Marketing Engineer (TME), Systems Development Unit, Cisco Systems**

Rabiul Hasan is a TME and DCI Technical Lead in the Systems Development Unit focusing on VMDC data center interconnect (DCI) technologies. Recent VMDC design and architecture efforts include VMDC 2.2 DCI DIG, VMDC Hybrid Cloud Architecture and DIG, and VMDC 2.0 DIG. He has 15 years of experience in Data Communication and Networking. He has an MBA from Duke University and MS in Telecommunications from University of Colorado at Boulder. He also holds numerous professional certifications including double CCIE (#15506), JNCIE(#188), VCP, CISSP, PMP and ITIL.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn is a service mark and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0805R)