fineDoc Number:	EDCS-<XXXXX>
Last Revision Date:	December 17, 2010
Created by:	Louis Watta
Template Ver. Number:	EDCS-XXXX Rev X

CAE
Nexus 1000V 1.4 LACP Offload White Paper

TABLE OF CONTENTS

1 Introduction

The Cisco Nexus 1000V, is a Cisco developed server virtualization switching architecture for VMware ESX environments. The Nexus 1000V enables policy based virtual machine (VM) connectivity, mobility of security and network properties, and a non-disruptive operational model for both Server and Network administrators.

Offering a set of network features, management tools and diagnostic capabilities consistent with the customer's existing physical Cisco network infrastructure and enhanced for the virtual world, the Nexus 1000V allows customers to accelerate their adoption of VMs through the unification & simplification of the physical and virtual networks. The n1000v also secures & simplifies the deployment & movement of VM's to increase service velocity while maintaining and enforcing security policy.

1.1 White Paper

The purpose of this white paper is to walk the user through enabling, configuring, and verifying LACP Offload.

1.2 Assumptions

The assumptions of this white paper are that the reader has

- Installed VMware VC 4.0U1/U2 or VC 4.1
- Installed Cisco Nexus 1000V 1.4 on an ESX VM
- At least 2 ESX 4.0U2/U1 or 4.1 boxes with VEM module already loaded
- Created a Nexus 1000V Distributed Virtual Switch (DVS) under vCenter
- Added the ESX boxes to the Nexus 1000V DVS

2 LACP Offload

New in this release of Cisco Nexus 1000V is a feature called LACP Offload. LACP is typically run on the control plane of a switch. In the context of the Nexus 1000V the control plane is run on the VSM. Since the VSM and the VEM do not reside on the same physical entity, there might be time windows where the VEM is disconnected from the VSM. In these scenarios, the VEM is considered to be in *headless* mode. When VEMs are in headless mode, LACP cannot be negotiated with the upstream switch as the VEM is not capable of running the protocol to negotiate LACP.

LACP Offload is an attempt to address the above situation by moving the LACP negotiation process onto the VEM module. With LACP Offload enabled the VEM can now negotiate LACP without being connected to the VSM. This eliminates several scenarios where the VEM and VSM can become isolated because an LACP port-channel cannot be brought up.

By default LACP Offload is disabled in this release. To enable and test LACP offload follow the steps below.

2.1 Sample Topology

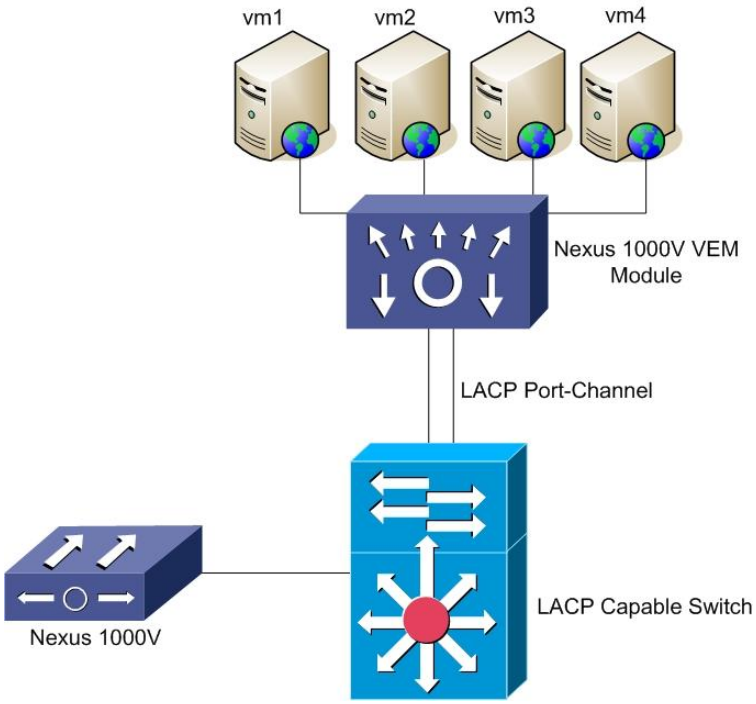

2.2 Create LACP port-channel on Nexus 1000V and Upstream Switch

LACP port-channels on Nexus 1000V are supported when all the connections of the VEM module connect to

- The same switch
- Multiple Switches that can support VPC, VSS, Stacking

This test plan uses the simplest configuration. All connections from a VEM module to the same Catalyst 6500.

Create the LACP port-channel on the upstream switch

```
cae-cat6k-1#show run int gig 4/7

!
interface GigabitEthernet4/7
 description esx-186
 switchport
 switchport mode trunk
 no ip address
 spanning-tree portfast
 spanning-tree bpdufilter enable
 spanning-tree bpduguard enable
 channel-protocol lacp
 channel-group 180 mode active
```

```
cae-cat6k-1#show run int gig 4/9

!
interface GigabitEthernet4/9
 description esx-186
 switchport
 switchport mode trunk
 no ip address
 spanning-tree portfast
 spanning-tree bpdufilter enable
 spanning-tree bpduguard enable
 channel-protocol lacp
 channel-group 180 mode active

cae-cat6k-1#show running int port-channel 180

!
interface Port-channel180
 switchport
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no ip address
```

2.3 Enable LACP Offload

New in this release of Nexus 1000V is the **feature** command. For the Nexus 1000V to support LACP port-channels you need to first enable the LACP feature

```
BL-beta# show feature
Feature Name          Instance  State
--------------------  --------  --------
dhcp-snooping         1         disabled
http-server           1         enabled
ippool                1         disabled
lacp                  1         disabled
lisp                  1         disabled
lisphelper            1         disabled
netflow               1         disabled
private-vlan          1         disabled
sshServer             1         enabled
tacacs                1         disabled
telnetServer          1         disabled

BL-beta(config)# feature lacp

BL-beta(config)# show feature
Feature Name          Instance  State
--------------------  --------  --------
dhcp-snooping         1         disabled
http-server           1         enabled
ippool                1         disabled
lacp                  1         enabled
lisp                  1         disabled
```

```
lisphelper            1         disabled
netflow               1         disabled
```

Once the LACP feature is turned on, “LACP Offload” has to be enabled. This WILL require a reboot of the the VSM.

```
n1kv-bl(config)# lacp offload
Please do a "copy running startup" to ensure the new setting takes effect on
next reboot  LACP Offload Status can be verified using "show lacp offload
status" Change in LACP Offload Status takes effect only on the next VSM Reboot
This can potentially cause modules with LACP uplinks to flap

n1kv-bl(config)# copy running-config startup-config
[########################################] 100%
n1kv-bl(config)# reload
This command will reboot the system. (y/n)?  [n] y
2010 Sep 29 16:40:37 n1kv-bl %PLATFORM-2-PFM_SYSTEM_RESET: Manual system
restart from Command Line Interface
```

When the VSM comes back up verify that LACP offload is enabled and create an uplink port-profile with “channel-group auto mode on” on the VSM.

```
n1kv-bl# show lacp offload status
    Current Status         : Enabled
    Running Config Status  : Enabled
    Saved Config Status    : Enabled

BL-beta(config)# port-profile type ethernet uplink-vpc
BL-beta(config-port-prof)#   vmware port-group
BL-beta(config-port-prof)#   switchport mode trunk
BL-beta(config-port-prof)#   channel-group auto mode active
BL-beta(config-port-prof)#   no shutdown
BL-beta(config-port-prof)#   system vlan 2,10
BL-beta(config-port-prof)#   state enabled
```

Add the ESX host to the Nexus 1000V that corresponds to the physical switch configuration you setup in the first step. A “show int brief” on the VSM should show the module with an LACP port-channel.

```
n1kv-bl# show int brief

--------------------------------------------------------------------------------
Port     VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0    --           up     172.18.217.242                          1000     1500

--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Port
Interface                                                                    Ch #
--------------------------------------------------------------------------------
Eth5/2        1      eth  trunk  up      none                      1000      2

--------------------------------------------------------------------------------
Port-channel VLAN  Type Mode   Status  Reason                   Speed  Protocol
Interface
--------------------------------------------------------------------------------
```

```
Po2          1     eth  trunk  up      none                       a-1000(D)  lacp
------------------------------------------------------------------------------
```

Note above how module 5 has one link as part of the port-channel and that Po2 is up and an LACP port-channel.

Add another link to the port-channel from the ESX host to have two valid links in the port-channel. This is done through VMware vCenter. After adding the second link you can see with "show int br" that we have two links and that Po2 is still up. You can also check the port-channel directly to make sure both links are up and active.

```
n1kv-bl# show int brief

--------------------------------------------------------------------------------
Port     VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0    --           up     172.18.217.242                          1000     1500

--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Port
Interface                                                                    Ch #
--------------------------------------------------------------------------------
Eth5/1        1       eth  trunk  up      none                      1000       2
Eth5/2        1       eth  trunk  up      none                      1000       2

--------------------------------------------------------------------------------
Port-channel VLAN  Type Mode   Status  Reason                   Speed  Protocol
Interface
--------------------------------------------------------------------------------
Po2          1     eth  trunk  up      none                       a-1000(D)  lacp

n1kv-bl# 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)
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
2     Po2(SU)     Eth      LACP      Eth5/1(P)    Eth5/2(P)
```

And on the upstream switch verify the LACP port-channel.

```
cae-cat6k-1#show lacp 180 neighbor
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 180 neighbors

Partner's information:

          Partner Partner   LACP Partner  Partner   Partner  Partner     Partner
Port      Flags   State     Port Priority Admin Key Oper Key Port Number Port State
Gi4/7     SA      bndl      32768         0x0       0x1      0x502       0x3D
Gi4/9     SA      bndl      32768         0x0       0x1      0x501       0x3D
```

Lastly you can check the on the VEM to make sure the LACP port-channel is up.

```
n1kv-bl# module vem 5 execute vemcmd show pc
pce_ind    chan   pc_ltl   pce_in_pc   LACP   SG_ID   NumVethsPinned   mbrs
-------    ----   ------   ---------   ----   -----   --------------   ----
      0       2      305           0      Y      0*                4    18,17,
```

2.4 Test that LACP offload is correctly working

At this point the ESX host should have 2 nics in the uplink port-profile and in an LACP port-channel to the Nexus 1000V. The benefit of LACP offload is that the VEM can negotiate LACP with the upstream switch if the VSM is offline.

In this test we will power off the VSM and reboot the ESX host. The LACP port-channel from the ESX host to the upstream switch should come up even with the VSM offline.

Step 1: Copy run start on the VSM

```
n1kv-bl# copy run start
[########################################] 100%
```

Step2: Shutdown the VSM from the vCenter console

Step3: Reboot the ESX host or run “vem restart”
After reboot or “vem restart” the port-channel should come up. Verify on the ESX host with “vemcmd show pc”

```
[root@cae-esx-186 ~]# vemcmd show pc
pce_ind    chan   pc_ltl   pce_in_pc   LACP   SG_ID   NumVethsPinned   mbrs
-------    ----   ------   ---------   ----   -----   --------------   ----
      0       2      305           0      Y      0*                4    17,18,
```

Also verify on the upstream switch the port-channel is up and active.

```
cae-cat6k-1#show lacp neighbor
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 180 neighbors

Partner's information:

          Partner Partner   LACP Partner  Partner   Partner  Partner     Partner
Port      Flags   State     Port Priority Admin Key Oper Key Port Number Port State
Gi4/7     SA      bndl      32768         0x0       0x1      0x502       0x3D
Gi4/9     SA      bndl      32768         0x0       0x1      0x501       0x3D
```

You should also note that on the upstream switch that while the port-channel is up there is no CDP information about the VEM present. CDP information comes from the control plane and with the VSM down there is no control plane.

Step 4: Power on the VSM
When you power on the VSM there should be no change in the config on the VEM since the LACP port-channel is already up. Once the VSM is up CDP data should get propagated and you should see ports from the VEM show up on the upstream switch with a "show cdp neighbor".

```
cae-cat6k-1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
RTP-UCS-N5K-1(FOX1009009B)
                 Ten 2/1            158          S I      N5K-C5020PEth 1/19
RTP-UCS-N5K-2(FOX1009009B)
                 Ten 2/2            150          S I      N5K-C5020PEth 1/19
cae-rtp10-gw1.cisco.com
                 Gig 4/1            167         R S I     WS-C3560E-Gig 0/10
n1kv-bl.cisco.com(Nexus-Switch)
                 Gig 4/7            159          S I      Nexus-SwitEth 5/2
n1kv-bl.cisco.com(Nexus-Switch)
                 Gig 4/9            159          S I      Nexus-SwitEth 5/1
n1kv-bl.cisco.com(Nexus-Switch)
                 Gig 3/25           143          S I      Nexus-SwitEth 4/2
rtp-6100-2-A(SSI12480256)
                 Ten 2/3            177          S I      N10-S6100 Eth 1/17
rtp-6100-2-B(SSI12490560)
                 Ten 2/4            152          S I      N10-S6100 Eth 1/17
QCI1342900P/00000C
```

Above you can see that ports 4/7 and 4/9 are now visible.