

What You Make Possible



Deploying and Troubleshooting the Nexus 1000v Virtual Switch on vSphere

BRKVIR-3013





Agenda

- Implement
 - The Deployment
- Optimise and Sustain
 - HA
 - Best Practices
 - Tools
- Operate
 - Troubleshooting

Session Prerequisites



Prerequisites

- Understanding of VMware ESX and vCenter Server
- Cisco NXOS
 - Understand the CLI
 - General switching concepts
- Cisco Nexus 1000V concepts
 - Understand VSM and VEM
 - Port-profile concepts

Current Releases



Current Nexus 1000V Releases

- Release 4.2(1)SV2(1.1) **NEW!**
- Release 4.2(1)SV1(5.1) (aka 1.5.1)
 - 1.5.1, 1.5.1a
 - Dropped support for ESX/ESXi 4.0
- 4.2(1)SV1(4b) (aka 1.4)
 - 1.4, 1.4a, and 1.4b
 - 1.4a first version to support ESXi 5


2.1 Features

- Split VSMS across data centres
- VEMs in remote branch offices
- Cisco TrustSec SXP support
- vCenter Web Plug-in
- vTracker
- Updated Installer
 - Installs the VSMS and the VEMs!

The Deployment




Deployment Complete!



Where do I start?



From the Beginning!



The Checklist

- Deploy VSM
 - Where will the VSM live
 - Which Control Mode
- Establish SVS connection
- Deploy VEM
- Create Port-Profile
- Add Uplinks

Virtual Supervisor Module




Virtual Supervisor Module (VSM)


- VSM is a Virtual Machine
 - On ESX/ESXi
 - On Nexus 1010
- Control plane for the Nexus 1000V solution
- Responsible for
 - VMware vCenter communication
 - Programming and managing Virtual Ethernet Modules (VEM)
- 1 VSM HA pair can manage 64 VEMs
- Nexus 1000V can coexist with VMware vSwitch and DVS

1000v Deployment Scenarios

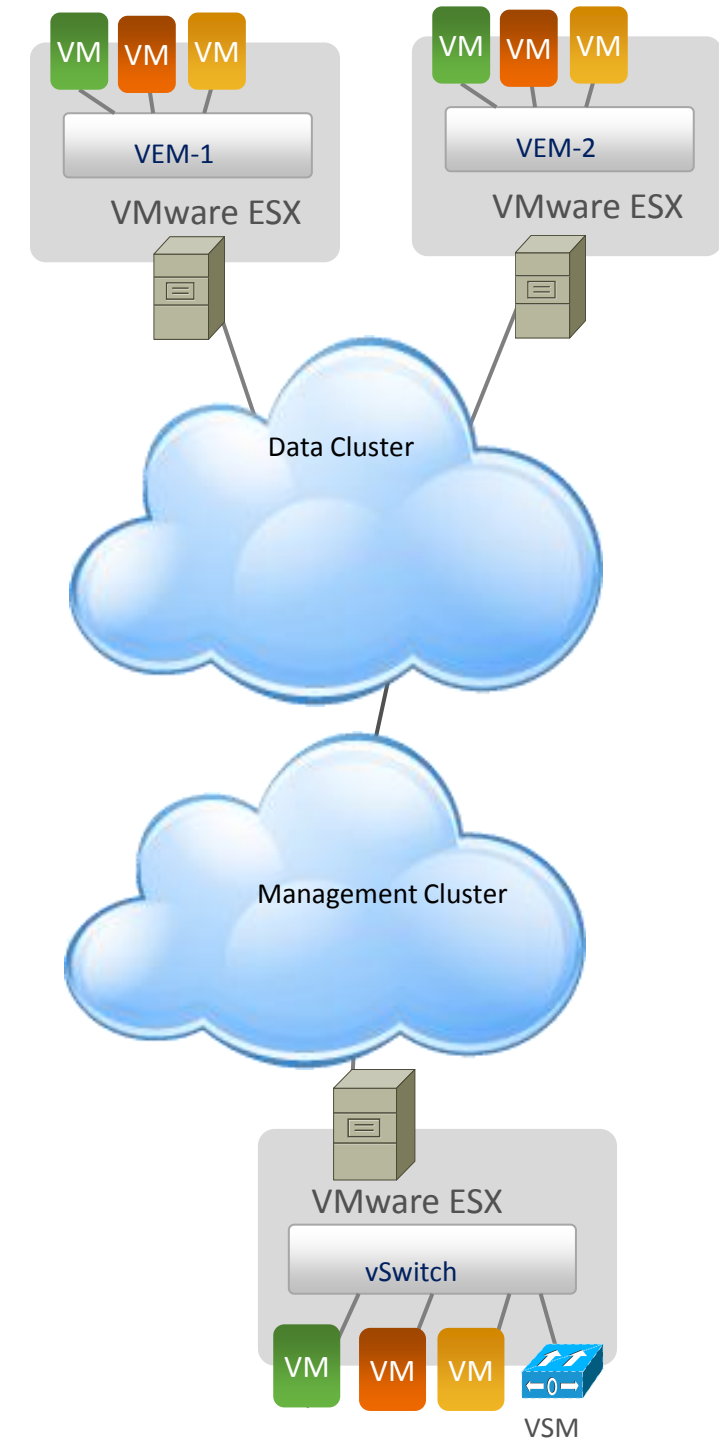
VSMs on VEM



VSMs on vSwitch

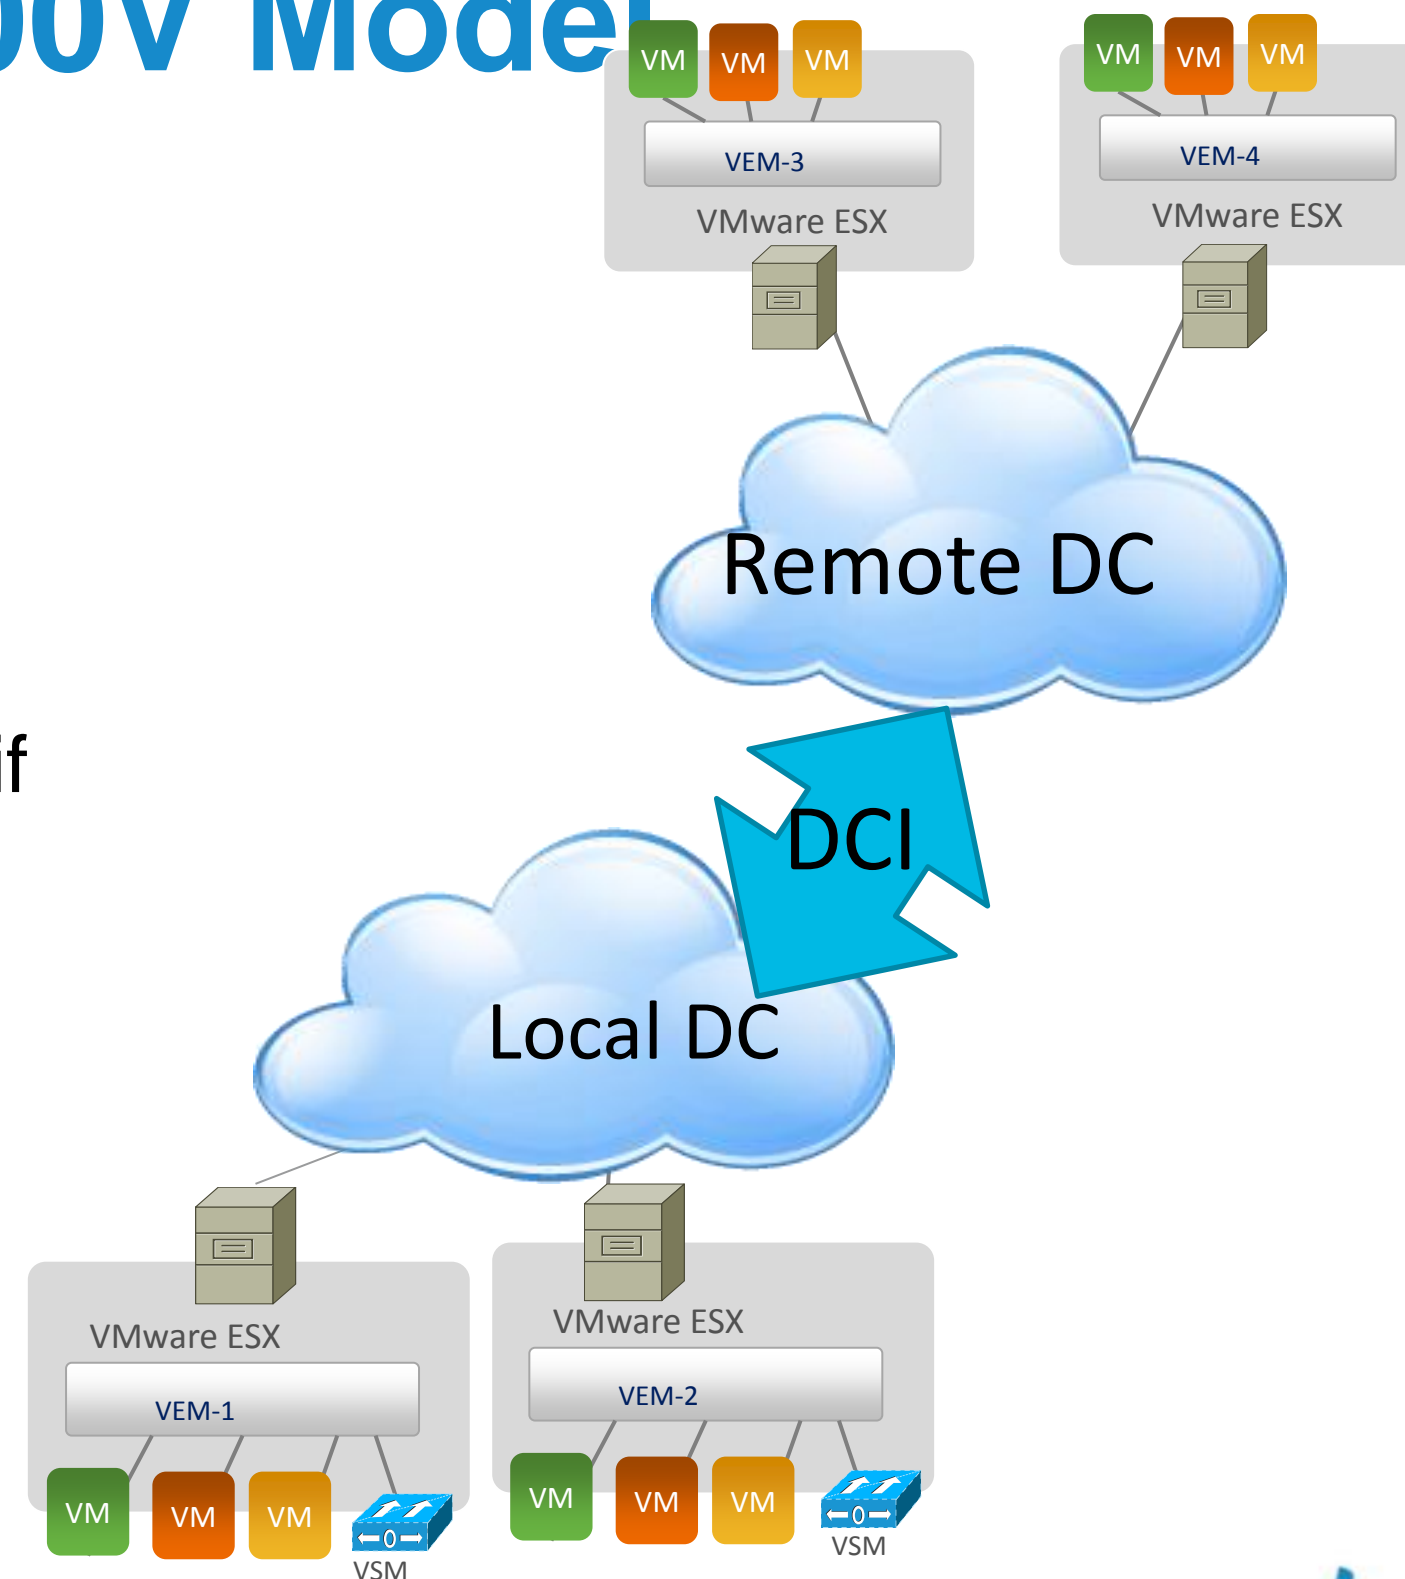


VSMs in Mgmt Cluster



Stretched Nexus 1000V Model

- VSMs and VEMs spread across Datacentres
- VSMs need to be in the same DC location
 - Prevents an active/active scenario if DC link goes down
 - L3 control is preferred
- Lowers supported configuration limits
 - Less VEMs and VLANs
 - Check release



VSM Control Modes

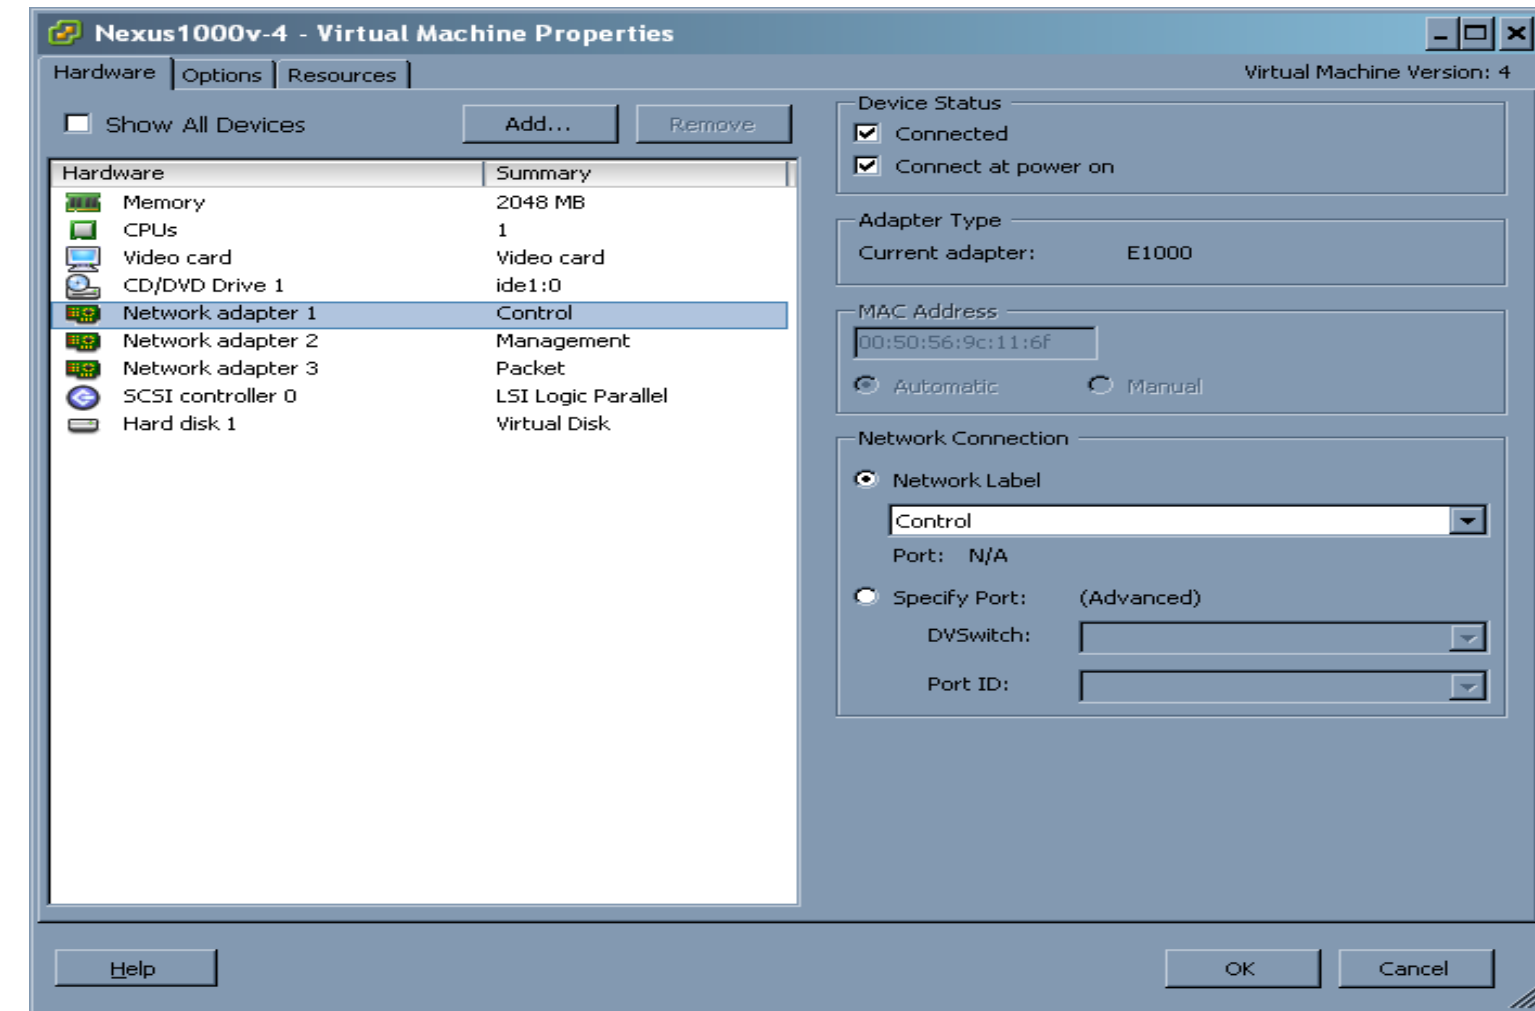
- L2 mode
 - Requires L2 connectivity through Control interface to all VEM modules
- L3 Mode
 - Default mode
 - Requires an IP address be assigned to the VEM
 - L3 uses UDP port 4785 for both source and destination
 - Uses Mgmt or Control interface of the VSM
 - VSM mgmt 0 is default interface for L3
 - Can also use “control 0”
 - Ties to control adapter of the VM (Adapter 1)

VSM Installation


- Initial configuration from Installer or via ISO image
- Installer Improvements
 - Standalone Java application now
 - Installs VSM in L2 or L3
 - Installs primary and secondary VSMs
 - Add ESX/ESXi hosts to Nexus 1000V

VSM Virtual Machine Requirements

- 3 network interfaces
- Adapter 1 is the **Control** interface
 - Heartbeat between VSMs and VEM(L2)
 - Heartbeat between VSMs
 - Control 0
- Adapter 2 is the **Management** interface
 - VSM terminal connectivity
 - Connectivity to VMware vCenter
 - Backup Heartbeat for VSM HA
 - Mgmt 0
- Adapter 3 is the **Packet** interface
 - Passes CDP and IGMP information
- 2GB of memory RAM reserved




The VSM



The Checklist

- ✓ Deploy VSM
- Establish SVS connection
- Deploy VEM
- Create Port-Profile
- Add Uplinks

VSM to VMware vCenter Communication



- VSM connects to vCenter using SSL connection.
- Management Interface
- Self-Signed certificate used for this connection
- VSM configures vCenter using its API
 - VSM creates N1KV Port-Groups in vCenter
 - VSM also stores opaque data in vCenter
 - VSM pulls information from vCenter (DC, DVS, VM, ...)

Create SVS connection

Specify SVS domain Information

```
n1000v# svb-domain
  domain id 888
  control vlan 1
  packet vlan 1
  svb mode L3 interface control0
```

Specify SVS connection Information

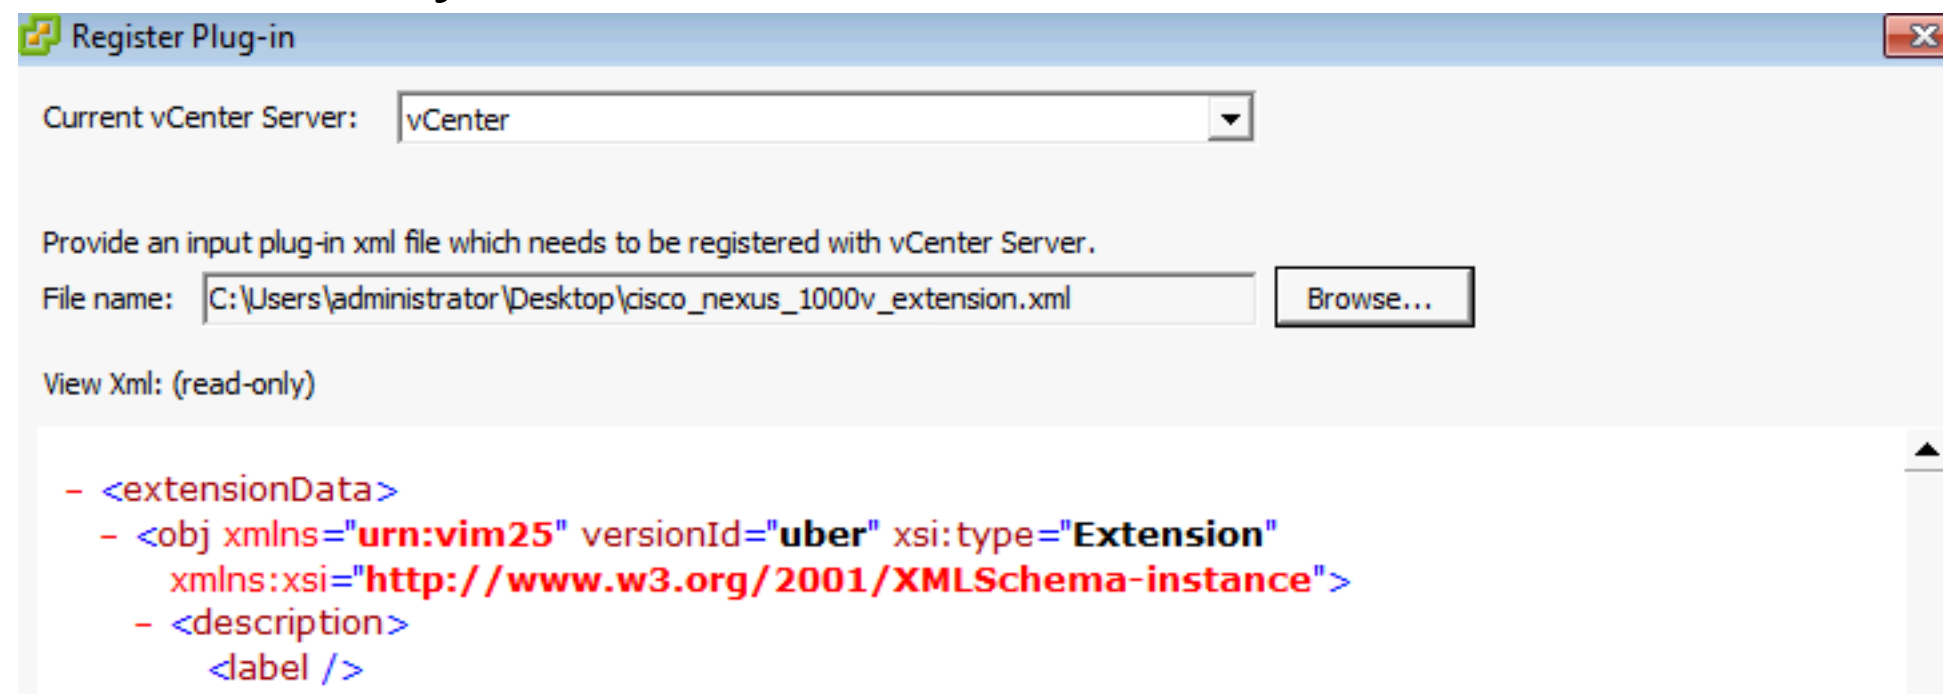
```
n1000v# svb connection VC-test
  protocol vmware-vim
  remote ip address 172.18.217.241 port 80
  vmware dvs
  connect
```

Connectivity Error – Extension Key

- Below error means wrong key or key is not registered

```
n1000v(config-svs-conn)# connect
ERROR: [VMware vCenter Server 4.0.0. build-162856] Extension key
was not registered before its use
```

- Register the Key in VMware vCenter



Connectivity Error – Connection Refused

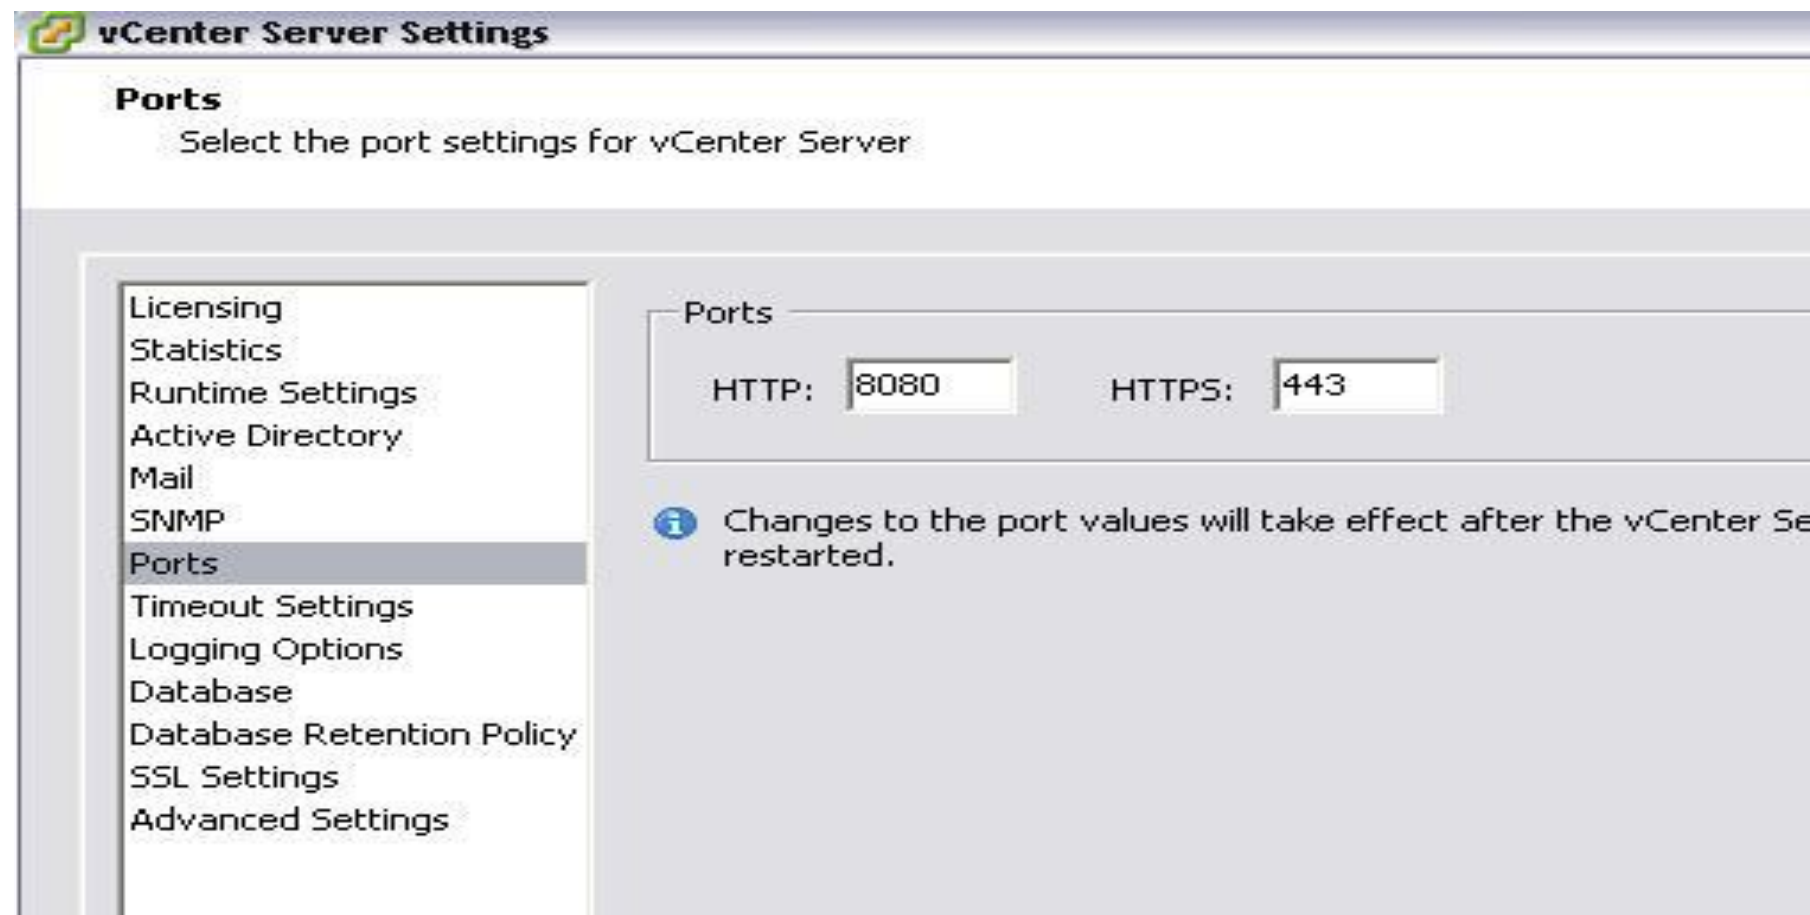
- Below error could indicate port mismatch

```
n1000v(config-svs-conn)# connect  
ERROR: [VMWARE-VIM] Operation could not be completed due to connection  
failure.Connection refused. connect failed in tcp_connect()
```

- Default port for communication is port 80
- Admins change default port for various reasons

What Port is VMware configured to Use?

- Verify Port number in vCenter
- Administration->vCenter Server Settings



Change the Port on the VSM

- Make sure VSM SVS port matches vCenter http port

```
n1000v# show svcs connections
connection vcenter:
  ip address: 172.18.217.41
  remote port: 80
  protocol: vmware-vim https
```

- To change the port

```
n1000v(config)# svcs connection vcenter
n1000v(config-svs-conn)# remote port 8080
```

Verify VSM to vCenter Connectivity

- Verify SVS connection settings

```
- n1000v# show svcs connections
```

```
connection VC-test:  
  ip address: 172.18.217.241  
  protocol: vmware-vim https  
  certificate: default
```

```
  datacenter name: Harrison
```

```
  DVS uuid: 72 f7 01 50 b2 01 7b 8b-55 68 cf df 10 5a db 55
```

```
  config status: Enabled
```

```
  operational status: Connected
```

- If Datacentre is underneath a folder and spaces

```
n1000v(config-svs-conn)# vmware dvs datacenter-name ?
```


```
  LINE  Datacenter name in VC with path (e.g. DCName, DCFolder/DC Name)
```

Backing up the VSM

- A running-config is not always enough to restore
- VSM on ESX/ESXi
 - We now support clone to a template
 - You can restore from a template and saved-config
- VSM on Nexus 1010
 - You can now export a VSM
 - Import the saved VSM to restore
- VSM on ESX/ESXi Snapshots
 - Not officially supported
 - Useful for upgrades

The VSM

Domain-cfg (Domain ID, Control VLAN, Packet VLAN)
Switchname
VSM image version
System profiles [System VLANS, profile names]
VSM IP address
VSM MAC address




The Checklist

- ✓ Deploy VSM
- ✓ Establish SVS connection
- Deploy VEM
- Create Port-Profile
- Add Uplinks

What is a VEM?

- Enhances VMware switch architecture
- Enables advanced switching capability on the hypervisor
- Provides each VM with dedicated “switch ports”



VEM Deployment

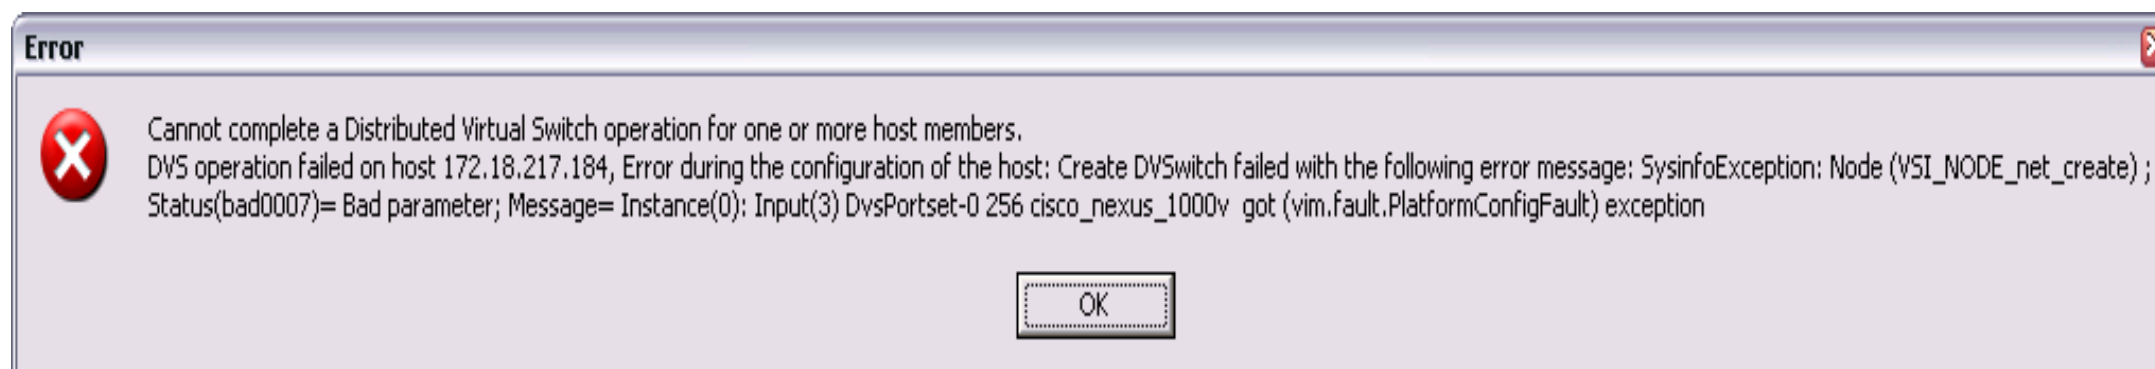
- Recommending L3 Control
- L3 control requires a VMKernel NIC on ESX
 - We need an L3 interface to forward control traffic
- We recommend creating VMK interface explicitly for the VEM
 - An existing VMK interface can be used in test or small environments
- VMK interface must be moved to the VEM

VEM Installation - VUM

- VMware Update Manager(VUM) does all the work
- Requires HTTP server on the VSM
- Turn off the following VMware cluster settings
 - HA, DRS, and DPM
- Logs on vCenter Server in
 - C:/Documents and Settings/Application Data/All Users/VMware/VUM/logs
- VEM modules get stored in
 - C:\ProgramData\VMware\VMware Update Manager\Data\hostupdate
 - CISCO directory comes from VSM
 - CSCO directory comes from VUM portal
 - Uninstalling VUM does not clean out the above directories

VEM Installation DVS Error

- DVS operation failed error
 - VUM is not installed or configured
 - VUM could not find the right VEM version
 - Check the VUM logs
 - Cluster HA, DRS, DPM was not disabled
 - Manual installation of VEM was not performed



VEM Installation - Manual

- With ESXi 5 use esxcli

- Can be run remote or locally on the host

- Local

```
esxi5.0# esxcli software vib install -v http://<VSM-IP-Address>/cisco/vibs/VEM/4.1.0/VEM-4.1.0-patch01/cross_cisco-vem-v140-4.2.1.1.5.1.0-3.0.1.vib
```

- Remote

```
linux1# esxcli --server <server> software vib install -v http://<VSM-IP-Address>/cisco/vibs/VEM/4.1.0/VEM-4.1.0-patch01/cross_cisco-vem-v140-4.2.1.1.5.1.0-3.0.1.vib
```

- ESX/ESXi 4.1

- Esxupdate (local)

```
esx4.1# esxupdate -b cross_cisco-vem-v140-4.2.1.1.5.1.0-2.0.1.vib update
```

- Vihostupdate (remote)

```
linux1# vihostupdate -install -bundle cisco-vem-v140-4.2.1.1.5.1.0-2.0.1.zip --server <server>
```


VEM Manual Installation Issues

- Dependency error

```
[root@cae-esx-180 ~]# esxupdate -b ./cross_cisco-vem-v100-4.0.4.1.1.27-0.4.2-release.vib u
pdate
cross_cisco-vem-v100-4.0.4.1.1.27.. ##### [100%]

Unpacking cross_cisco-vem-v100-es.. ##### [100%]

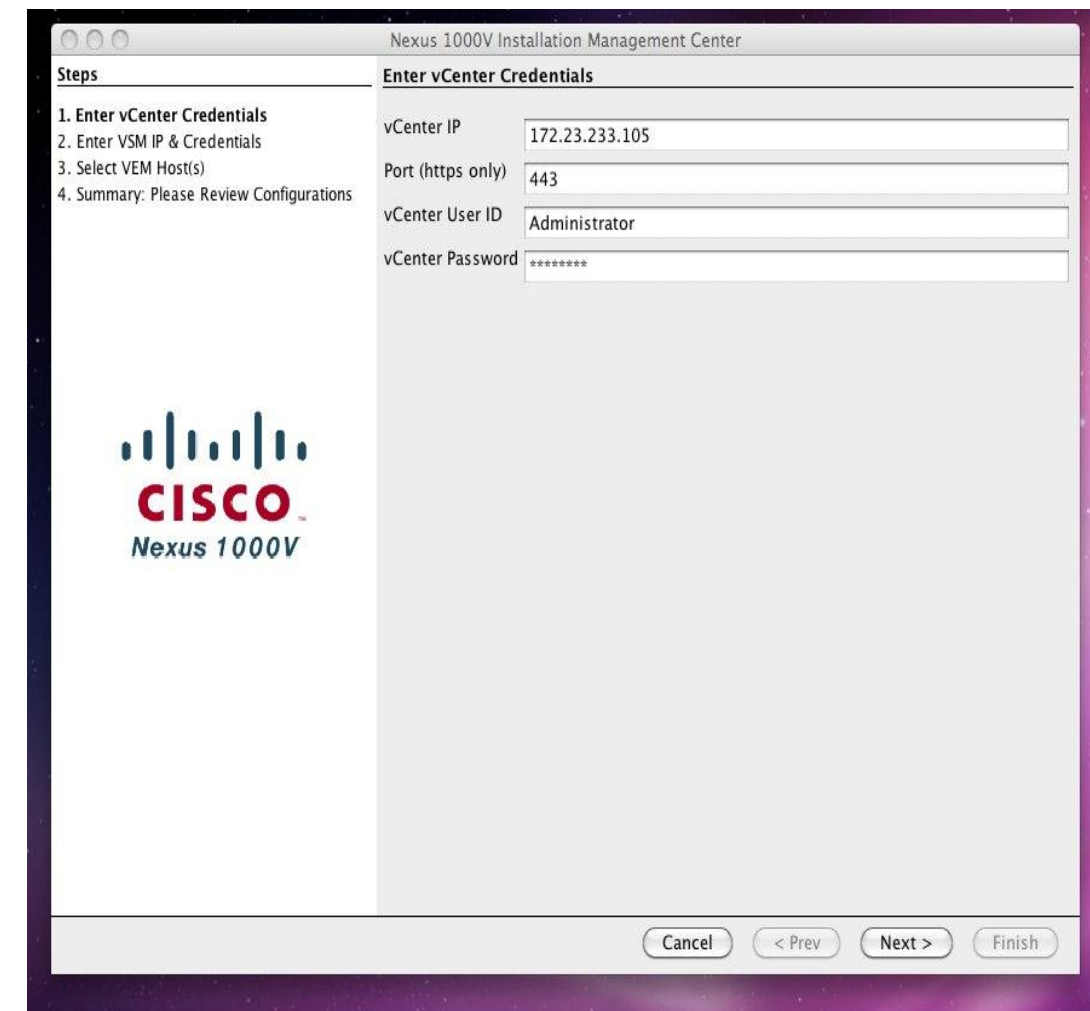
The following problems were encountered trying to resolve dependencies:
  No VIB provides 'vmknexus1kvapi-0-4' (required by cross_cisco-
  vem-v100-esx_4.0.4.1.1.27-0.4.2)
  Requested VIB cross_cisco-vem-v100-esx_4.0.4.1.1.27-0.4.2 conflicts with the
  host
```

- Verify VEM VIB version to ESX/ESXi version
- Compatibility matrix will identify right VEM version

VEM Installation – Nexus 1000V Installer App

- Installation App is now standalone Java application
- Requires administrator privileges to the ESXi host
- Allows the network admin to directly install the VEM
- Need to start it with VEM option

```
java.exe -jar Nexus1000V-install.jar VEM
```




The screenshot shows the 'Nexus 1000V Installation Management Center' application window. On the left, a 'Steps' pane lists the installation process: 1. Enter vCenter Credentials (current step), 2. Enter VSM IP & Credentials, 3. Select VEM Host(s), and 4. Summary: Please Review Configurations. The main area is titled 'Enter vCenter Credentials' and contains four input fields: 'vCenter IP' with the value '172.23.233.105', 'Port (https only)' with '443', 'vCenter User ID' with 'Administrator', and 'vCenter Password' with '*****'. The Cisco logo and 'Nexus 1000V' text are visible on the left side of the main area. At the bottom, there are navigation buttons: 'Cancel', '< Prev', 'Next >', and 'Finish'.

VEM Installation – ESXi Stateless

- VMware introduced Stateless ESXi with version 5
- ESXi PXE boots
- No information is stored on local disks
 - No place to install the VEM and store opaque data
- VEM module has to be built into the boot image
- Possible using VMware Powershell
- Instructions are in Install and Upgrade Guide
 - http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_sv_1_5_1/install_upgrade/vsm_vem/guide/n1000v_installupgrade.html

The VSM



VMK Interface

- ESX/ESXi VMK interfaces are special interfaces
- Need to be created by VMware Admin
- Verify interface exists

```
~ # esxcfg-vmknic -l
Interface  Port  Group/DVPort  IP Family  IP Address  Netmask
Broadcast  MAC Address  MTU  TSO MSS  Enabled  Type
vmk0      VMkernel  IPv4      172.18.217.180  255.255.255.0
172.18.217.255  00:18:fe:72:f2:aa 1500  65535  true  STATIC
vmk2      34      IPv4      192.168.11.180  255.255.255.0
192.168.11.255  00:50:56:78:6f:52 1500  65535  true  STATIC
```

- Can you ping the control interface of the VSM?

```
~ # vmkping 192.168.11.10
PING 192.168.11.10 (192.168.11.10): 56 data bytes
64 bytes from 192.168.11.10: icmp_seq=0 ttl=255 time=0.570 ms
64 bytes from 192.168.11.10: icmp_seq=1 ttl=255 time=0.420 ms
```


The Checklist

- ✓ Deploy VSM
- ✓ Establish SVS connection
- ✓ Deploy VEM
- Create Port-Profile
 - Port-Channel Type
 - Which System Vlans
- Add Uplinks

Port Channels

- 3 load balancing modes
- LACP Port-channels
 - Upstream switch support and configuration
- VPC – MAC Pinning
 - Works with any upstream switch
 - Allows for pinning of veths (VM) to specific links.
- VPC – Host Mode CDP/Manual
 - NIC association is either Manual or CDP
 - Multiple connections per physical switch require a port-channel

Cisco Nexus 1000V System VLANs

- System VLANs enable interface connectivity before an interface is programmed
- Address chicken and egg issue
 - VEM needs to be programmed, but it needs a working network for this to happen
- Port profiles that contain system VLANs are “system port profiles”
 - Allowed 32 port-profiles with system VLAN
- System port-profiles become part of the opaque data
 - VEM will load system port-profiles and pass traffic even if VSM is not up
- System vlans must be set on egress and ingress port-profiles

Create Uplink Port Profile

- Uplinks are usually a trunk
- Specify the Port-Channelling mode for uplinks
- Ensure that all system vlans required for intial communications are specified
 - VSM Control
 - VSM Packet
 - VSM Management
 - vCenter/Host Management

```
port-profile type ethernet uplink-pinning
```

```
vmware port-group
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan all
```

```
channel-group auto mode on mac-pinning
```

```
no shutdown
```

```
system vlan 2,10,150-151
```

Create Veth Port-Profiles

Example: Management Port-profile

```
n1000v# show run port-profile Management
port-profile type vethernet Management
  vmware port-group
  switchport mode access
  switchport access vlan 2
  no shutdown
system vlan 2
```

Example: VM Port-profile

```
n1000v# show run port-profile VM-150
port-profile type vethernet VM-150
  vmware port-group
  switchport mode access
  switchport access vlan 150
  no shutdown
```

Create Veth Port-Profiles

Example: Vmkernel Port-Profile for L3 control

```
port-profile type vethernet L3control
  capability l3control
  vmware port-group
  switchport mode access
  switchport access vlan 702
  no shutdown
  system vlan 702
  state enabled
```

Specifies that vmk's assigned to this port-profile will be used for VSM-VEM communication

System VLAN required

How do I Recover?

- Vemcmd allows you to set system vlan on LTLs on ESX host
- From ESX/ESXi console

```
[root@cae-esx-180 ~]# vemcmd show port
```

LTL	IfIndex	Vlan	Bndl	SG_ID	Pinned_SGID	Type	Admin	State	CBL	Mode	Name
8	0	3969	0	32	32	VIRT	UP	UP	4	Access	l20
9	0	3969	0	32	32	VIRT	UP	UP	4	Access	l21
10	0	150	0	32	0	VIRT	UP	UP	4	Access	l22
11	0	3968	0	32	32	VIRT	UP	UP	4	Access	l23
12	0	151	0	32	0	VIRT	UP	UP	4	Access	l24
13	0	1	0	32	32	VIRT	UP	UP	0	Access	l25
14	0	3967	0	32	32	VIRT	UP	UP	4	Access	l26
15	0	3967	0	32	32	VIRT	UP	UP	4	Access	l27
16	1a040000	1 T	304	0	32	PHYS	UP	UP	4	Trunk	vmnic0
48	1b040000	10	0	32	0	VIRT	UP	UP	4	Access	vmk0
49	1b040010	2	0	32	0	VIRT	UP	UP	4	Access	vswif0
304	16000002	1 T	0	32	32	VIRT	UP	UP	4	Trunk	

```
[root@cae-esx-180 ~]# vemcmd set system-vlan 2 ltl 49
```

The Checklist

- ✓ Deploy VSM
- ✓ Establish SVS connection
- ✓ Deploy VEM
- ✓ Create Port-Profile
- Add Uplinks and VM's

Add Host and Nic

Select Hosts and Physical Adapters
Select hosts and physical adapters to add to this vSphere distributed switch.


Select Host and Physical Adapters
[Network Connectivity](#)
Virtual Machine Networking
Ready to Complete

Settings... View Incompatible Hosts...

Host/Physical adapters	In use by switch	Settings	Uplink port group
<input checked="" type="checkbox"/> 10.67.82.211		View Details...	
Select physical adapters			
<input checked="" type="checkbox"/> vmnic2	--	View Details...	Uplink
<input type="checkbox"/> vmnic3	vSwitch0	View Details...	Select an uplink port gr...

The Magic!

Domain-cfg (Domain ID, Control VLAN, Packet VLAN)
 Switchname
 VSM image version
 System profiles [System VLANS, profile names]
 VSM IP address
 VSM MAC address



VEM Programing

Vemcmd Show card
 SVS information
 Domain
 Card Slot number
 Control IP address
 Control Mac address
 Control Mode

VEM Seeding

- How does VEM know VSM information?
- Opaque data – copied to VEM to seed during install.
- Opaque data consists of:
- Domain-cfg (Domain ID, Control VLAN, Packet VLAN)
 - Switchname
 - VSM image version
 - System profiles [System VLANs, profile names]
 - VSM IP address
 - VSM MAC address

Checking Opaque Data

- VSM stores opaque-data in vCenter as persistent data for its DVS.
- vCenter downloads this information to ESX for VEM to use, whenever a host is added to N1KV-DVS

Checking opaque-data in VSM

```
switch-cp# show svcs domain
```

```
SVS domain config:
```

```
Domain id: 100  
Control vlan: 150  
Packet vlan: 150  
Status: Config push
```

Checking opaque-data in VEM


```
[root@sfish-30-119 sbin]#  
/usr/lib/ext/cisco/nexus/vem/sbin/vemcmd show card
```

```
Switch name: switch-cp  
Card domain: 100  
Card slot: 2  
Card control VLAN: 150  
Card packet VLAN: 150
```

Checking opaque-data in vCenter

1. https://vc_ip_address/mob/.
2. Content → rootFolder (group-dx) → childEntity (dataCenter-n) → networkFolder (group-n6) → childEntity (group-n) → childEntity (dvs-n) → config → VendorSpecificConfig

Deployment Complete!



Plan **B**



=



Nexus 1000V High Availability



VSM – VSM Heartbeat

Statistics & Errors

n1000v-MV# show system internal redundancy info

My CP:

```
slot: 0
domain: 184
role: primary
status: RDN_ST_AC
state: RDN_DRV_ST_AC_SB
intr: enabled
power_off_reqs: 0
reset_reqs: 1
```

Active VSM

Other CP:

```
slot: 1
status: RDN_ST_SB
active: true
ver_rcvd: true
degraded_mode: false
```

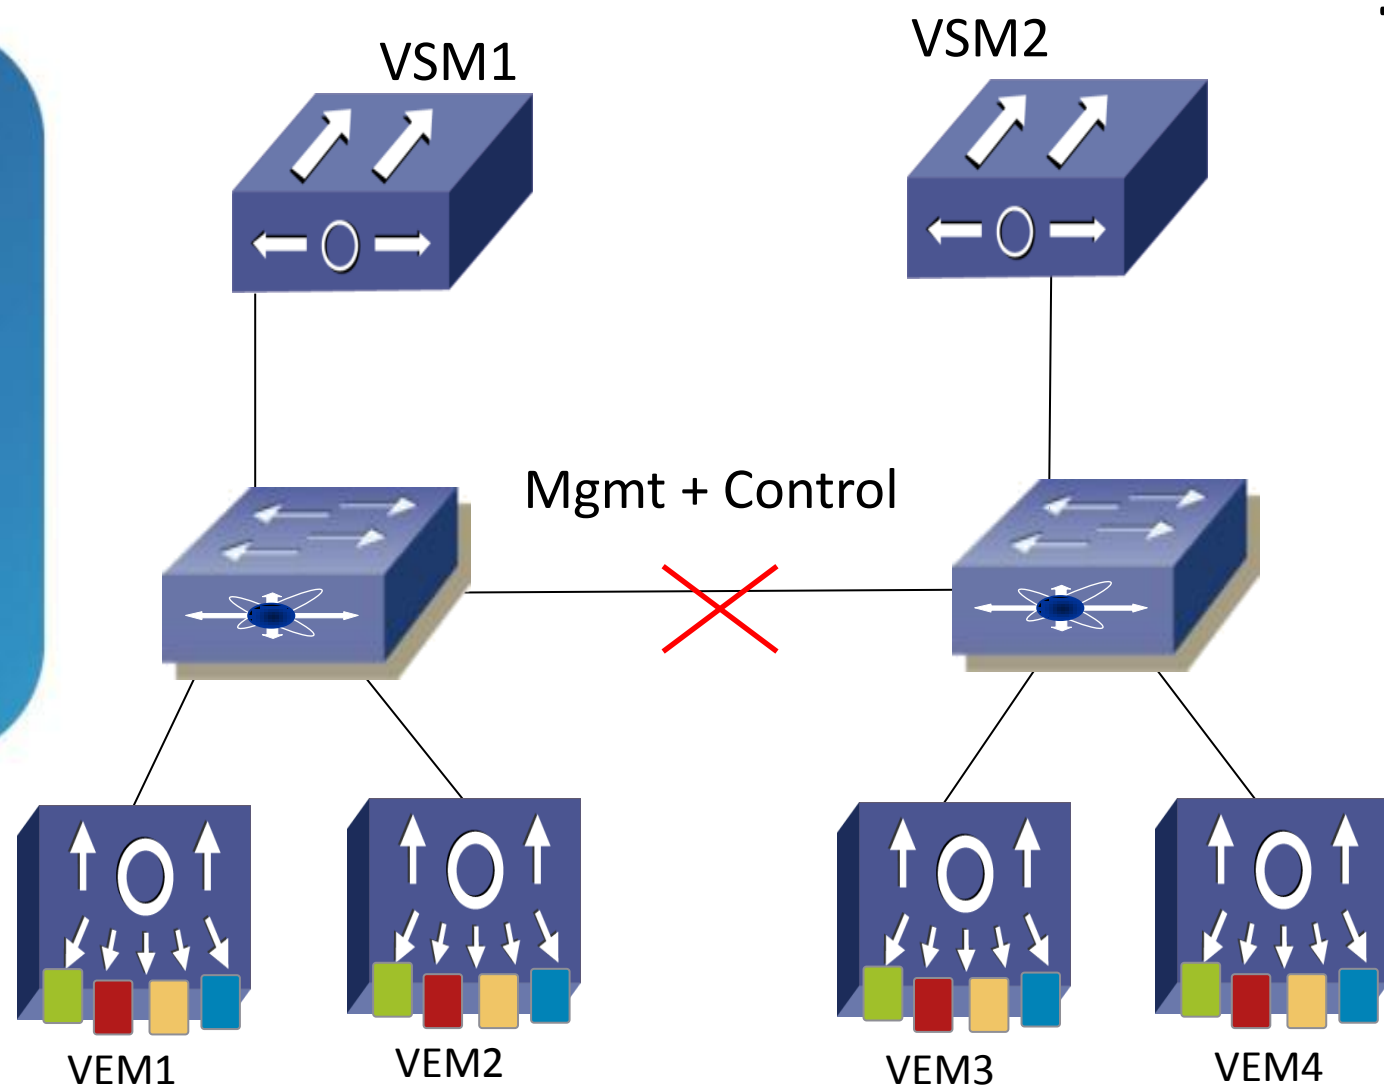
Standby VSM

Redun Device 0:

```
name: ha0
```

```
pdev: bc1bb000
alarm: false
mac: 00:50:56:8e:5e:f5
tx_set_ver_req_pkts: 13
tx_set_ver_rsp_pkts: 2
tx_heartbeat_req_pkts: 168155
tx_heartbeat_rsp_pkts: 318
rx_set_ver_req_pkts: 2
rx_set_ver_rsp_pkts: 1
rx_heartbeat_req_pkts: 318
rx_heartbeat_rsp_pkts: 168148
rx_drops_wrong_domain: 0
rx_drops_wrong_slot: 0
rx_drops_short_pkt: 0
rx_drops_queue_full: 0
rx_drops_inactive_cp: 0
rx_drops_bad_src: 0
rx_drops_not_ready: 0
rx_unknown_pkts: 0
```


Fail Scenario 1 – Split Brain



Failed Interface

#3 & #4

- Split Brain

Effect:


- VSM1 stays active, keeping VEM1 & 2
- VSM2 becomes Active, taking over VEM3 & 4
- If changes are made on VSM2 redundancy manager will notice and make it primary when links come back up
- Important note, just creating and changing configs on VSM2 is not enough. We need to syncs to vCenter.

VSM High Availability

- Behaviour was changed in 1.4a
- No longer continuous rebooting of standby VSM
 - Single reboot and it stays in warm standby state
- Management interface is truly used for backup heartbeat
 - State changes to “degraded” when control 0 fails
- Use “show system internal redundancy trace” for state changes

```
n1kv-13# show system internal redundancy trace
1 0s START_THREAD      ST_NP   ST_NP   ST_INVALID
2 0s CP_STATUS_CHG      ST_INIT ST_NP   ST_INIT
3 5s DEGRADED_MODE     ST_INIT ST_NP   ST_INIT
4 5s STATE_TRANS       ST_INIT ST_NP   ST_INIT   EV_OS_NP   ST_AC_NP
5 0s CP_STATUS_CHG      ST_AC   ST_NP   ST_AC_NP
6 4s SET_VER_RCVD      ST_AC   ST_NP   ST_AC_NP
7 0s STATE_TRANS       ST_AC   ST_INIT ST_AC_NP   EV_OS_INIT ST_AC_INIT
8 0s STATE_TRANS       ST_AC   ST_SB   ST_AC_INIT EV_OS_SB   ST_AC_SB
```


Fail Scenario 2



Failed Interface #3 - Split DVS

Effect:

- VEM3,4 become headless
- VSMs use mgmt for heartbeat
- VSM1 remains primary
- VSM2 does not reboot. Stays secondary

Exit:

- VEM3,4 reconnect
- VSM use control for heartbeat
- VSM1 remains primary
- VSM2 reboots, joins as secondary

Tools & Best Practices



Tools

- NXOS show commands on VSM
 - Show vtracker - **new command**
- Cisco VEM commands on ESX
 - /usr/sbin/vem*
- Networking view in vCenter
 - vCenter 5.0 Web Plugin – **new plugin**

Best Practices

- VSM
 - Use L3 Control with Control 0
 - VSM to VSM/VEM Latency 5-10ms
 - Back up config!!
- VEM
 - Check version
 - Hard code VM
- Port-Channels
 - Use LACP with upstream Port-channel
 - Use Mac-pinning with no Port-channel
- Spanning-Tree and BPDU
 - Ports must be in portfast/edge mode
 - BPDU Guard and Filter recommended

VEM-VSM Connectivity Troubleshooting



Symptoms

- VEM adds in vCenter but does not show up on VSM “show module”
- With L2 most of the time its a Control VLAN issue
 - Verify Control VLAN connectivity
- With L3 its usually an IP routing problem
 - If you can ping VMK interface the VEM should connect to VSM
 - Troubleshoot as you would all VMware L3 issues


```
Cisco-Live-N1K# show module
```

Mod	Ports	Module-Type	Model	Status
1	0	Virtual Supervisor Module	Nexus1000V	active *
2	0	Virtual Supervisor Module	Nexus1000V	ha-standby
3	248	Virtual Ethernet Module	NA	ok
5	248	Virtual Ethernet Module	NA	ok

Mod	Sw	Hw
1	4.2(1)SV2(1.1)	0.0
2	4.2(1)SV2(1.1)	0.0
3	4.2(1)SV2(1.1)	VMware ESXi 5.0.0 Releasebuild-623860 (3.0)
5	4.2(1)SV2(1.1)	VMware ESXi 4.1.0 Releasebuild-260247 (2.0)

Mod	MAC-Address(es)	Serial-Num
1	00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8	NA
2	00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8	NA
3	02-00-0c-00-03-00 to 02-00-0c-00-03-80	NA
5	02-00-0c-00-05-00 to 02-00-0c-00-05-80	NA

Mod	Server-IP	Server-UUID	Server-Name
1	10.67.82.88	NA	NA
2	10.67.82.88	NA	NA
3	10.67.82.63	202da395-c666-e111-911b-503de59e5fcc	10.67.82.63
5	10.67.82.211	00000000-0000-0000-1200-00840000000a	10.67.82.211

```
* this terminal session
```

Check the
present modules

```
Cisco-Live-N1K# sh module vem missing
```

```
Mod  Server-IP          Server-UUID          Server-Name
---  -
4    10.67.82.54        02000000-8000-0000-0000-000000000000c  vsphere-54
```

Confirm missing
vem

```
Cisco-Live-N1K# show module vem counters
```

```
-----
Mod  InNR  OutMI  InMI  OutHBeats  InHBeats  InsCnt  RemCnt  Crit Tx Errs
-----
3    1     1     1     32439     29404     1       0       0     0
4    202   2     2     28219     27219     2       2       0     0
5    5555  3     3     28716     1773      3       2       0     0
```

```
Cisco-Live-N1K# [16D[J
```

```
Cisco-Live-N1K# show module vem counters
```

```
-----
Mod  InNR  OutMI  InMI  OutHBeats  InHBeats  InsCnt  RemCnt  Crit Tx Errs
-----
3    1     1     1     32454     29419     1       0       0     0
4    205   2     2     28234     27219     2       2       0     0
5    5555  3     3     28731     1788      3       2       0     0
```

OutHbeats being
sent

InHbeats not
being received

```
Cisco-Live-N1K# sh svcs domain
```

```
SVS domain config:
```

```
Domain id: 888
```

```
Control vlan: NA
```

```
Packet vlan: NA
```

```
L2/L3 Control mode: L3
```

```
L3 control interface: control0
```

```
Status: Config push to VC successful.
```

```
Control type multicast: No
```

Check svcs
domain

```
Cisco-Live-N1K# show int control 0
```

```
control0 is up
```

```
Hardware: Ethernet, address: 0050.5680.1bfc (bia 0050.5680.1bfc)
```

```
Internet Address is 192.168.2.88/24
```

```
MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,  
reliability 255/255, txload 1/255, rxload 1/255
```

```
Encapsulation ARPA
```

```
full-duplex, 1000 Mb/s
```

```
Auto-Negotiation is turned on
```

```
1 minute input rate 3512 bits/sec, 4 packets/sec
```

```
1 minute output rate 64 bits/sec, 0 packets/sec
```

```
Rx
```

```
35468 input packets 20746 unicast packets 5100 multicast packets
```

```
9622 broadcast packets 3761565 bytes
```

```
Tx
```

```
1148 output packets 29 unicast packets 570 multicast packets
```

```
549 broadcast packets 185532 bytes
```

Check the L3
Control Interface
on VSM

The screenshot displays the vCenter console interface for an ESXi host with IP 10.67.82.54. The left-hand navigation pane shows a tree structure with 'Cisco Live' selected, containing sub-items like 'Live' (with IP addresses 10.67.82.211, 10.67.82.54, 10.67.82.63, and Cisco-Live-N1K-1/2) and 'VXI' (with DC-Solutions and various servers like Ale, anam, CIAC, etc.).

The main console area is titled '10.67.82.54 VMware ESXi, 5.0.0, 623860'. It features several tabs: Summary, Virtual Machines, Performance, Configuration, Tasks & Events, Alarms, Permissions, Maps, and Storage View. The 'Configuration' tab is active, showing a 'view:' dropdown set to 'vSphere Standard Switch' and 'vSphere Distributed Switch'.

The 'Networking' section is expanded to show 'Cisco-Live-N1K'. A table lists network interfaces:

Interface Name	Virtual Machines
L3-Control-vmk	0
VMkernel Ports (1)	0
vmk2 : 192.168.2.54	0
N1K-Control	0
N1K-Management	0
N1K-Packet	0
Unused_Or_Quarantine_V...	0

The 'L3-Control-vmk' row is highlighted with a blue selection box. To the right, a 'Uplink' configuration panel is partially visible, showing a dropdown menu with a plus sign for adding more uplinks.

Check Host settings in vCenter

```
Cisco-Live-N1K# show run port-profile L3-Control
```

```
port-profile type vethernet L3-Control-vmk
```

```
capability l3control
```

```
vmware port-group
```

```
switchport mode access
```

```
switchport access vlan 702
```

```
no shutdown
```

```
system vlan 702
```

```
state enabled
```

Check vmk L3
port-profile

```
Cisco-Live-N1K# show run port-profile Uplink
```

```
port-profile type ethernet Uplink
```

```
vmware port-group
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1-3967,4048-4093
```

```
channel-group auto mode on mac-pinning
```

```
no shutdown
```

```
system vlan 174,701-703
```

```
state enabled
```

Check uplink
port-profile

```
Cisco-Live-N1K# ping 192.168.2.54 vrf default
```

```
PING 192.168.2.54 (192.168.2.54): 56 data bytes
```

```
64 bytes from 192.168.2.54: icmp_seq=0 ttl=63 time=0.661 ms
```

```
64 bytes from 192.168.2.54: icmp_seq=1 ttl=63 time=0.471 ms
```

```
64 bytes from 192.168.2.54: icmp_seq=2 ttl=63 time=0.461 ms
```

```
64 bytes from 192.168.2.54: icmp_seq=3 ttl=63 time=0.548 ms
```

```
64 bytes from 192.168.2.54: icmp_seq=4 ttl=63 time=0.37 ms
```

Ping vmk0

```
gchami-mac:~ gchami$ ssh root@10.67.82.54
```

```
~ # vem status
```

```
VEM modules are loaded
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	4	64	1500	vmnic0
DVS Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
Cisco-Live-N1K	256	13	256	1500	vmnic1

```
VEM Agent (vemdpa) is running
```

Check vem status

```
~ # vem version
```

```
Running esx version -623860 x86_64
```

```
VEM Version: 4.2.1.2.1.1.0-3.0.2
```

```
VSM Version: 4.2(1)SV2(1.1)
```

```
System Version: VMware ESXi 5.0.0 Releasebuild-623860
```

Check vem version

```
~ # vemcmd show port
```

LTL	VSM Port	Admin	Link	State	PC-LTL	SGID	Vem Port	Type
18	Eth4/2	UP	UP	FWD	305	1	vmnic1	
49	Veth6	UP	UP	FWD	0	1	vmk2	
305	Po3	UP	UP	FWD	0			

```
~ #
```

Check vmk port


```
~ # vemcmd show card
Card UUID type 2: 02000000-8000-0000-0000-000000000000c
Card name: vsphere-54
Switch name: Cisco-Live-N1K
Switch alias: DvsPortset-0
Switch uuid: 29 b2 00 50 a5 ba 10 84-23 4a 1c 86 0b 5d 02 e3
Card domain: 888
Card slot: 4
VEM Tunnel Mode: L3 Mode
L3 Ctrl Index: 49
L3 Ctrl VLAN: 702
VEM Control (AIPC) MAC: 00:02:3d:13:78:03
VEM Packet (Inband) MAC: 00:02:3d:23:78:03
VEM Control Agent (DPA) MAC: 00:02:3d:43:78:03
VEM SPAN MAC: 00:02:3d:33:78:03
Primary VSM MAC : 00:50:56:80:1b:fc
Primary VSM PKT MAC : 00:50:56:80:1b:fe
Primary VSM MGMT MAC : 00:50:56:80:1b:fd
Standby VSM CTRL MAC : ff:ff:ff:ff:ff:ff
Management IPv4 address: 10.67.82.54
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
Primary L3 Control IPv4 address: 192.168.2.88
Secondary VSM MAC : 00:00:00:00:00:00
Secondary L3 Control IPv4 address: 0.0.0.0
Upgrade : Default
Max physical ports: 32
Max virtual ports: 216
Card control VLAN: 1
Card packet VLAN: 1
Control type multicast: No
Card Headless Mode : Yes
<output removed>
Licensed: Yes
```

Check opaque
data

```

~ # vmkping 192.168.2.88
PING 192.168.2.88 (192.168.2.88): 56 data bytes

~ # vmkping 192.168.2.63
PING 192.168.2.63 (192.168.2.63): 56 data bytes
64 bytes from 192.168.2.63: icmp_seq=0 ttl=64 time=0.604 ms
64 bytes from 192.168.2.63: icmp_seq=1 ttl=64 time=0.360 ms
64 bytes from 192.168.2.63: icmp_seq=2 ttl=64 time=0.219 ms

--- 192.168.2.63 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.219/0.394/0.604 ms

```

Check vmk connectivity

```

Cisco-Live-N1K#
Cisco-Live-N1K#Cisco-Live-N1K# 2013 Jan 17 07:20:41 Cisco-Live-N1K %VEM_MGR-2-
VEM_MGR_DETECTED: Host vsphere-54 detected as module 4
2013 Jan 17 07:20:41 Cisco-Live-N1K %VEM_MGR-2-MOD_ONLINE: Module 4 is online
2013 Jan 17 07:20:45 Cisco-Live-N1K %VEM_MGR-2-VEM_MGR_REMOVE_NO_HB: Removing VEM 5
(heartbeats lost)
2013 Jan 17 07:20:50 Cisco-Live-N1K %VEM_MGR-2-MOD_OFFLINE: Module 5 is offline

```

After Ping

```

Cisco-Live-N1K#
Cisco-Live-N1K#Cisco-Live-N1K# show module vem missing
Mod  Server-IP          Server-UUID          Server-Name
---  -
5    10.67.82.211        00000000-0000-0000-1200-00840000000a  localhost.
Cisco-Live-N1K#

```

Missing module changed



enter
Cisco Live
Live
10.67.82.211
10.67.82.54
10.67.82.63
Cisco-Live-N1K-1
Cisco-Live-N1K-2

10.67.82.211 VMware ESXi, 4.1.0, 260247

Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage Views Hardware

view: vSphere Standard Switch vSphere Distributed Switch

Networking Refresh Properties

Cisco-Live-N1K

- L3-Control-vmk
 - VMkernel Ports (1)
 - vmk0 : 192.168.2.54
- N1K-Control
 - Virtual Machines (1)
- N1K-Management
 - Virtual Machines (0)
- N1K-Packet
 - Virtual Machines (1)
- Unused_Or_Quarantine_V...

Uplink

- UpLink00 (1N1)

Module 4 & 5
have a duplicate
vmk ip address

10.67.82.54 VMware ESXi, 5.0.0, 623860

Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage View

view: vSphere Standard Switch vSphere Distributed Switch

Networking

Cisco-Live-N1K

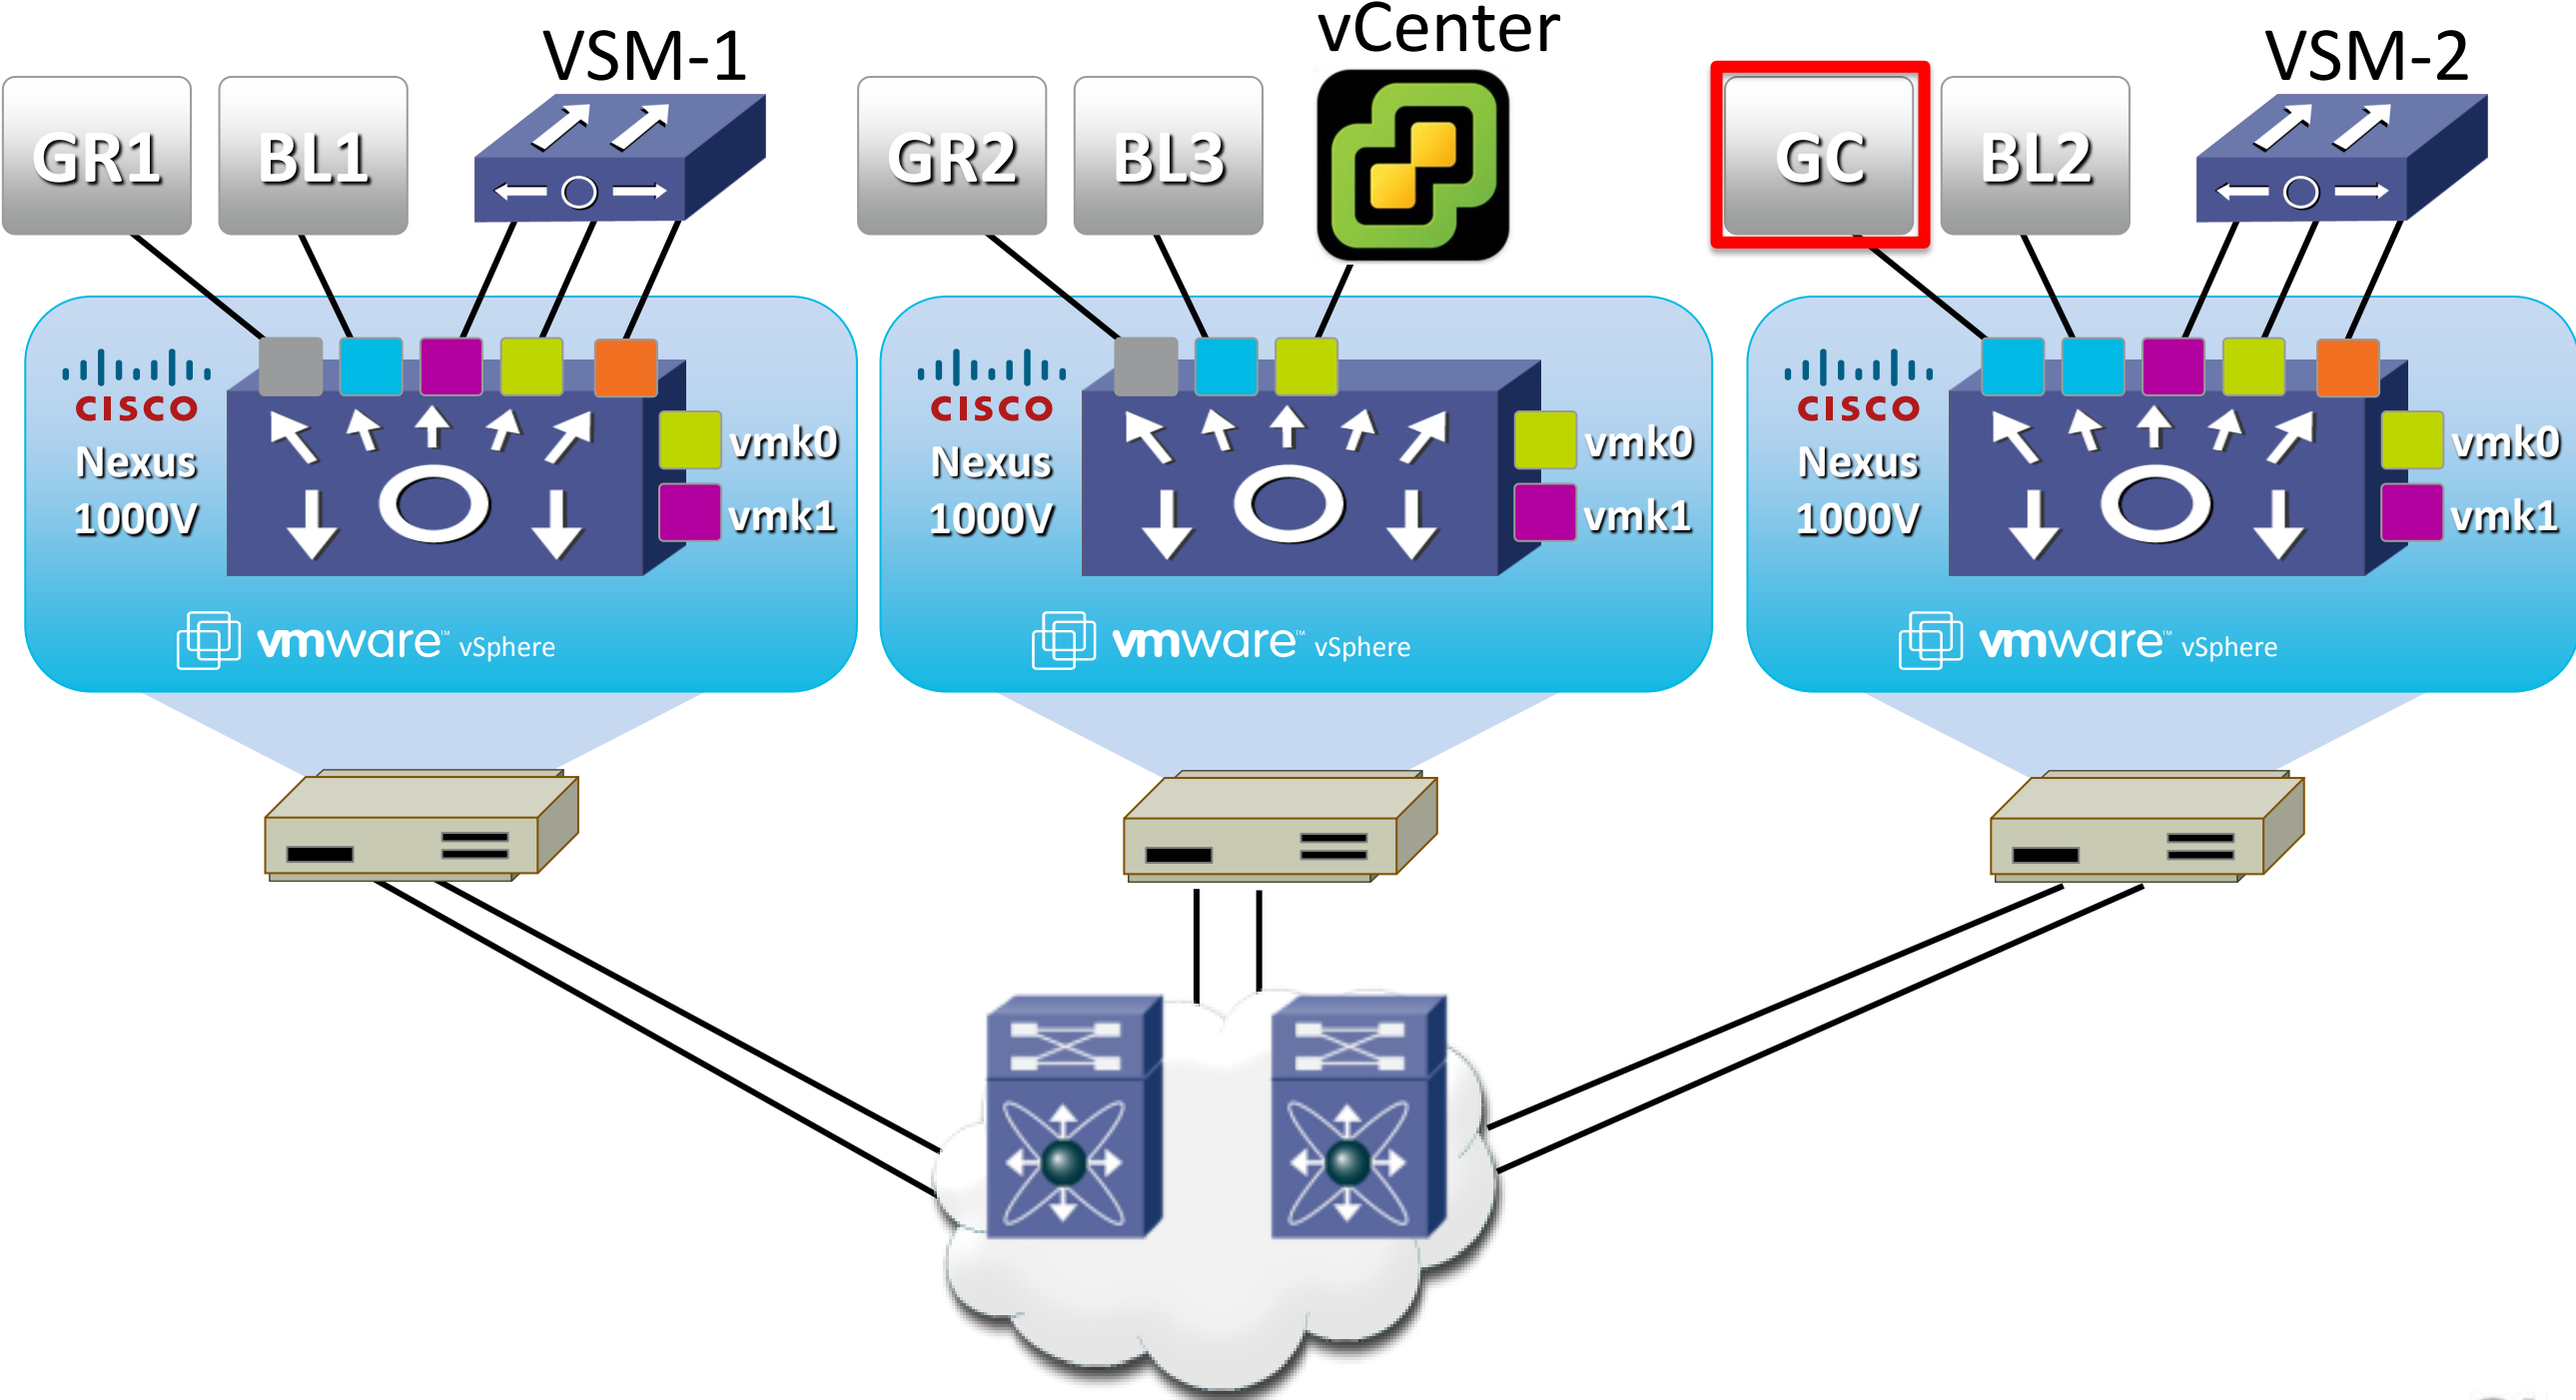
- L3-Control-vmk
 - VMkernel Ports (1)
 - vmk2 : 192.168.2.54
 - Virtual Machines (0)
- N1K-Control
 - Virtual Machines (0)
- N1K-Management
 - Virtual Machines (0)
- N1K-Packet
 - Virtual Machines (0)
- Unused_Or_Quarantine_V...
 - Virtual Machines (0)

Module 4 & 5 have a duplicate vmk ip address

Virtual Machine Connectivity Troubleshooting



Symptom



VM Name	Short Name	IP Address
Gchami-Fileserver	GC	192.168.5.88
VXI-Blue-1	BL1	192.168.5.81
VXI-Blue-2	BL2	192.168.5.82
VXI-Blue-3	BL3	192.168.5.83
VXI-Grey-1	GR1	192.168.6.81
VXI-Grey-2	GR2	192.168.6.82

VLAN	IP address	Default Gateway
705	192.168.5.0/24	192.168.5.254
706	192.168.6.0/24	192.168.6.1


```
gchami-Fileserver-Blue on 10.67.82.63
File View VM
[Icons]

C:\Windows\system32\cmd.exe
Packets: Sent = 4, Received = 1, Lost = 3 (75% loss),
C:\Users\administrator>ping 192.168.5.82
Pinging 192.168.5.82 with 32 bytes of data:
Reply from 192.168.5.82: bytes=32 time<1ms TTL=128
Reply from 192.168.5.82: bytes=32 time<1ms TTL=128
Reply from 192.168.5.82: bytes=32 time<1ms TTL=128
Reply from 192.168.5.82: bytes=32 time<1ms TTL=128
Ping statistics for 192.168.5.82:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\Users\administrator>ping 192.168.5.254
Pinging 192.168.5.254 with 32 bytes of data:
Request timed out.
Reply from 192.168.5.88: Destination host unreachable.
Request timed out.
Request timed out.
Ping statistics for 192.168.5.254:
    Packets: Sent = 4, Received = 1, Lost = 3 (75% loss),
C:\Users\administrator>ping 192.168.5.81
Pinging 192.168.5.81 with 32 bytes of data:
Reply from 192.168.5.88: Destination host unreachable.
Reply from 192.168.5.88: Destination host unreachable.
Reply from 192.168.5.88: Destination host unreachable.
Reply from 192.168.5.88: Destination host unreachable.
Ping statistics for 192.168.5.81:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
C:\Users\administrator>
```

Test 1:
Same VLAN
Same VEM

Test 2:
Default Gateway

Test 3:
Same VLAN
Different VEM



```
Cisco-Live-N1K# show int virtual
```

Port	Adapter	Owner	Mod	Host
Veth1	Net Adapter 2	Cisco-Live-N1K-1	5	10.67.82.211
Veth2	vmk0	VMware VMkernel	3	10.67.82.63
Veth3	Net Adapter 1	Cisco-Live-N1K-1	5	10.67.82.211
Veth4	Net Adapter 3	Cisco-Live-N1K-1	5	10.67.82.211
Veth5	vmk0	VMware VMkernel	5	10.67.82.211
Veth6	vmk2	VMware VMkernel	4	10.67.82.54
Veth7	Net Adapter 1	Cisco-Live-N1K-2	5	10.67.82.211
Veth8	Net Adapter 3	Cisco-Live-N1K-2	5	10.67.82.211
Veth9	Net Adapter 1	VXI-Grey-1	5	10.67.82.211
Veth10	Net Adapter 1	VXI-Blue-2	3	10.67.82.63
Veth11	Net Adapter 1	gchami-Fileserver-Blue	3	10.67.82.63
Veth12	Net Adapter 1	VXI-Blue-1	5	10.67.82.211
Veth13	Net Adapter 1	VXI-Blue-3	4	10.67.82.54
Veth14	Net Adapter 1	VXI-Grey-2	4	10.67.82.54
Veth15	vmk4	VMware VMkernel	3	10.67.82.63

Find VM's veth



```
Cisco-Live-N1K# show int virtual
```

Port	Adapter	Owner	Mod	Host
Veth1	Net Adapter 2	Cisco-Live-N1K-1	5	10.67.82.211
Veth2	vmk0	VMware VMkernel	3	10.67.82.63
Veth3	Net Adapter 1	Cisco-Live-N1K-1	5	10.67.82.211
Veth4	Net Adapter 3	Cisco-Live-N1K-1	5	10.67.82.211
Veth5	vmk0	VMware VMkernel	5	10.67.82.211
Veth6	vmk2	VMware VMkernel	4	10.67.82.54
Veth7	Net Adapter 1	Cisco-Live-N1K-2	5	10.67.82.211
Veth8	Net Adapter 3	Cisco-Live-N1K-2	5	10.67.82.211
Veth9	Net Adapter 1	VXI-Grey-1	5	10.67.82.211
Veth10	Net Adapter 1	VXI-Blue-2	3	10.67.82.63
Veth11	Net Adapter 1	gchami-Fileserver-Blue	3	10.67.82.63
Veth12	Net Adapter 1	VXI-Blue-1	5	10.67.82.211
Veth13	Net Adapter 1	VXI-Blue-3	4	10.67.82.54
Veth14	Net Adapter 1	VXI-Grey-2	4	10.67.82.54
Veth15	vmk4	VMware VMkernel	3	10.67.82.63

Find VM's veth

```
Cisco-Live-N1K# show module
```

Mod	Ports	Module-Type	Model	Status
1	0	Virtual Supervisor Module	Nexus1000V	ha-standby
2	0	Virtual Supervisor Module	Nexus1000V	active *
3	248	Virtual Ethernet Module	NA	ok
4	248	Virtual Ethernet Module	NA	ok
5	248	Virtual Ethernet Module	NA	ok

Check Module



```
Cisco-Live-N1K# show run int vethernet 11
```

```
interface Vethernet11
  inherit port-profile blue-vlan-705
  description gchami-Fileserver-Blue, Network Adapter 1
  vmware dvport 480 dvswitch uuid "29 b2 00 50 a5 ba 10 84-23 4a 1c 86 0b 5d 02
e3"
  vmware vm mac 0050.5680.75F1
```

Check Veth
config

```
Cisco-Live-N1K# show int vethernet 11
```

```
Vethernet11 is up
Port description is gchami-Fileserver-Blue, Network Adapter 1
Hardware: Virtual, address: 0050.5680.75f1 (bia 0050.5680.75f1)
Owner is VM "gchami-Fileserver-Blue", adapter is Network Adapter 1
Active on module 3
VMware DVS port 480
Port-Profile is blue-vlan-705
Port mode is access
5 minute input rate 216 bits/second, 0 packets/second
5 minute output rate 112 bits/second, 0 packets/second
Rx
 831 Input Packets 3631715202634248684 Unicast Packets
7306371624041709615 Multicast Packets 8101729934719874174 Broadcast Packets
99641 Bytes
Tx
292057777198 Output Packets 201 Unicast Packets
154618822814 Multicast Packets 2334956331018056756 Broadcast Packets 597 Flo
od Packets
8317696973251616757 Bytes
32491047111387492 Input Packet Drops 451766743649 Output Packet Drops
```

Check Veth
Status

```
Cisco-Live-N1K# show run port-profile blue-vlan-705
```

```
port-profile type vethernet blue-vlan-705  
  vmware port-group  
  switchport mode access  
  switchport access vlan 705  
  no shutdown  
  state enabled
```

Check Veth
port-profile

```
Cisco-Live-N1K# show run port-profile Uplink
```

```
port-profile type ethernet Uplink  
  vmware port-group  
  switchport mode trunk  
  switchport trunk allowed vlan 1-3967,4048-4093  
  channel-group auto mode on mac-pinning  
  no shutdown  
  system vlan 174,701-703  
  state enabled
```

Check Uplink
Port-profile

```

~ # vemcmd show port
LTL    VSM Port  Admin Link  State  PC-LTL  SGID  Vem Port  Type
  21    Eth3/5   UP    UP    FWD    305    4    vmnic4
  22    Eth3/6   UP    UP    FWD    305    5    vmnic5
  49    Veth2     UP    UP    FWD     0     4    vmk0
  50    Veth15    UP    UP    FWD     0     4    vmk4
  53    Veth10    UP    UP    FWD     0     4    VXI-Blue-2.eth0
  54    Veth11    UP    UP    FWD     0     5    gchami-Fileserver-Blue.eth0
 305    Po1      UP    UP    FWD     0
~ #

```

Check local port
on VEM

```

~ # vemcmd show 12 705
Bridge domain 12 brtmax 4096, brtcnt 7, timeout 300
VLAN 705, swbd 705, ""
Flags: P - PVLAN S - Secure D - Drop
      Type          MAC Address  LTL  timeout  Flags  PVLAN
Dynamic 00:50:56:80:2f:69 305    41
Dynamic 50:3d:e5:67:a4:62 305     1
Dynamic 00:50:56:80:1b:f0 305    70
Static  00:50:56:80:75:f1  54     0
Dynamic 00:50:56:80:31:bc 305     6
Static  00:50:56:80:31:bb  53     0
Dynamic 00:50:56:80:31:ba 305   313

```

Check local
mac-address
Table


```

~ # vempkt clear
Cleared log
~ # vempkt capture all-stages vlan 705 ltl 305
Successfully set packet capture specification
~ # vempkt start
Started log
~ # vempkt stop
Will suspend log after next 0 entries
~ # vempkt display brief all | grep 53
Entry  Src-LTL  VLAN  Length  Payload
12     53     705    169    33 33 00 01 00 02 00 50 56 80 31 bb 86 dd
36     53     705     66    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
38     53     705     60    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
39     53     705    404    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
42     53     705     60    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
43     53     705     66    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
48     53     705     60    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
49     53     705    254    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
53     53     705    254    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
56     53     705     83    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
57     53     705    129    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
59     53     705    254    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
65     53     705     98    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
67     53     705     65    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
68     53     705     60    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00
73     53     705     60    50 3d e5 67 a4 62 00 50 56 80 31 bb 08 00

```

Preform a packet capture to ensure traffic is leaving VM & VEM

```
Cisco-Live-N1K# show int virtual pinning module 3
```

```
-----  
Veth      Pinned      Associated PO List of  
          Sub Group id interface      Eth interface(s)  
-----  
Veth2     4           Po1          Eth3/5  
Veth10    4           Po1          Eth3/5  
Veth11    5           Po1          Eth3/6  
Veth15    4           Po1          Eth3/5
```

Trace physical
pinning for Veth
11

```
Cisco-Live-N1K# show cdp neighbors interface ethernet 3/6 detail
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge  
                  S - Switch, H - Host, I - IGMP, r - Repeater,  
                  V - VoIP-Phone, D - Remotely-Managed-Device,  
                  s - Supports-STP-Dispute
```

```
Device-ID          Local Intrfce Hldtme Capability Platform      Port ID  
VXI-Core-N5K-2(SS1145 Eth3/6      131      S I s        N5K-C5548P    Eth1/12
```

```
-----  
Device ID:VXI-Core-N5K-2(SS1145001FX)  
System Name: VXI-Core-N5K-2
```

```
Interface address(es):  
  IPv4 Address: 10.67.82.254
```

```
<output removed>
```

Find upstream
devices

```
VXI-Core-N5K-2# show int ethernet 1/12 brief
```

```
-----  
Ethernet      VLAN   Type Mode   Status Reason           Speed   Port  
Interface                                           Ch #  
-----  
Eth1/12       1      eth  trunk up     none             10G(D) --
```

Check Upstream
interface status

```
VXI-Core-N5K-2# show run int ethernet 1/12
```

```
interface Ethernet1/12  
  switchport mode trunk  
  switchport trunk allowed vlan 1-704,706-3967,4048-4093  
  spanning-tree port type edge trunk
```

Check Upstream
interface
configuraiton

Upstream Interface missing VLAN 705

Wrap Up!

- Plan! Plan! Plan!
- Hidden issues will be discovered
- Keep your tools at hand
- Symptoms help scope



Q & A



Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

www.ciscoliveaustralia.com/portal/login.wv

Cisco *live!*



Appendix



Index

- Appendix 1: Best Practices
- Appendix 2: Tools
- Appendix 3: More HA scenarios
- Appendix 4: VEM-VSM Connectivity Troubleshooting
- Appendix 5: VM Connectivity Troubleshooting
- Appendix 6: Nexus 1010 and 1010-X

Appendix 1: Best Practices



VSM Best Practices

- L3 control will become preferred method
- Use control 0 over mgmt 0
- Primary and Standby VSM in same L2 domain!!!
- VSM on VEM is supported
- VSM primary to secondary latency 5-10ms
- VSM to VEM latency 5-10ms
- Backup your config!!!

VSM and vMotion

- Manual vMotion of VSM is supported
- Not recommended to allow DRS to vMotion Primary and Secondary VSM
- Aggressive DRS vMotion setting can cause VSM to drop packets and loose connectivity to VEM
- Best practice to keep Primary and Secondary VSM outside of DRS

VEM Best Practices

- Control network should have low latency and available bandwidth
- VEM and VSM running on the same versions
- Upstream switch ports configured identically
- On UCS make Service Profile does not contain “Dynamic VNICs”
 - VEM and VM-FEX are mutually exclusive
- Hard code VEM to module number with

```
N1000v-mv# config t
n1000v-mv(config)# vem 12
n1000v-mv(config-vem-slot)# host vmware id 33393138-3335-5553-4537-30354E375832
```

Port Channels – Best Practice

- If the upstream switch can be clustered (VPC, VBS Stack, VSS) use **LACP**
- If you are using LACP also use LACP Offload
- If the upstream switch can NOT be clustered use **MAC-PINNING**
- Create channel-groups in port-profile
 - Let VSM build the port-channel
- All physical switch ports in port-channel configured identical

Spanning-tree and BPDU – Best Practice

- Mandatory Spanning-tree settings per port
 - IOS set STP portfast

```
cat65k-1(config-if)# spanning-tree portfast trunk
```
 - NXOS set port type edge

```
n5k-1(config-if)# spanning-tree port type edge trunk
```
- Highly Recommended Global BPDU Filter/Guard
 - IOS

```
cat65k(config)# spanning-tree portfast bpdufilter  
cat65k(config)# spanning-tree portfast bpduguard
```
 - NXOS

```
n5k-1(config)# spanning-tree port type edge bpduguard default  
n5k-1(config)# spanning-tree port type edge bpdufilter default
```
- BPDU Filter is mandatory for LACP port-channels
- Set per port BPDU Guard when Global is not possible

Appendix 2: Tools



Locations of ESX Tools

- Cisco VEM commands on ESX
 - /usr/lib/ext/cisco/nexus/vem*/sbin
 - Linked in /usr/sbin/
- VEM commands can also be run remote from VSM

```
n1000v# module vem 3 execute vemcmd show port
LTL    IfIndex  Vlan    Bndl  SG_ID Pinned_SGID  Type  Admin State  CBL Mode  Name
8      0      3969    0      2      2      VIRT   UP    UP    4 Access 120
```

- For ESXi use the remote commands from the VSM
 - vemcmd through RCLI does not work for ESXi
 - Enable “unsupported mode”

Google unsupported ESXi

More Tools

- Vemcmd on ESX host
 - Can query and set configuration
 - Improved in 1.4 to allow more “set” commands and many more show commands
- Vem-health on ESX host
 - Will try to point you in right direction if the vem is having issues
- Mping on VSM
 - Command that will broadcast out on the control network looking for VEM modules

vCenter 5.0 Web Plugin

Issues Tasks Events Health **Cisco Nexus 1000V**

Getting Started Dashboard Switch **Hosts/VEM** About

System

Host Name	NX-OS Version	Host IP	License	Host / Module	VMs / Host	VNICs / Host
10.78.0.125	4.2(1)SV2(1.1)	10.78.0.125	licensed	3	1	1

VM Info Port Groups vNICs Uplinks

VMs	VNICs	Adapter	Status	Port Group	VLANs	Host ID
ESX-Host-124	Vethernet1	Net Adapter 1	✓	1812	1812	3

Logout

vTracker

- Introduced in 2.1. Needs to be enabled as a feature
- vTracker provides various intuitive virtualisation perspectives to the network-admin.
- Following different views are available,
 - Upstream view
 - Vlan view
 - Module PNic view
 - VM View
 - VM VNic view
 - VM Info View
 - VMotion view

vTracker Upstream View

- **show vtracker upstream-view [device-id *name* | device-ip *IP address*]**

```
VSM-N1k# show vtracker upstream-view
```

Device-Name Device-IP	Device-Port Local-Port	Server-Name Adapter Status	PC-Type P0-Intf	Veth-interfaces
Upstream-SW-A 172.23.231.27	Gig2/7 Eth3/3	172.23.232.117 vmnic2 up	MacPinn Po1	10-11
Upstream-SW-B 172.23.231.15	Gig3/10 Eth3/4	172.23.232.117 vmnic3 up	MacPinn Po1	9
	Gig3/8 Eth4/3	172.23.232.118 vmnic2 up	Default Po2	1-2
	Gig3/9 Eth4/4	172.23.232.118 vmnic3 up	Default Po2	1-2

```
VSM-N1k#
```

vTracker Module PNic view

- Provides hardware information about the physical network interface cards (pNICs) that are connected to each of the VEM server module managed by Nexus 1000v

```
VSM-N1k# show vtracker module-view pnic
```

Mod	EthIf	Adapter Description	Mac-Address	Driver	DriverVer	FwVer
3	Eth3/8	vmnic7 Intel Corporation 82576	0050.5652.f935	igb Gigabit Network Connection	2.1.11.1	1.4-3
4	Eth4/3	vmnic2 Intel Corporation 82546GB	0050.565e.df74	e1000 Gigabit Ethernet Controller	8.0.3.2-1vmw-NAPI	N/A
4	Eth4/4	vmnic3 Intel Corporation 82546GB	0050.565e.df75	e1000 Gigabit Ethernet Controller	8.0.3.2-1vmw-NAPI	N/A

```
VSM-N1k#
```

vTracker Vlan view

- provides information about all the VMs that are connected to a specific VLAN or a range of VLANs. It is a view from the VLAN perspective.
- show vtracker vlan-view [vlan number/range]**

```
VSM-N1k# show vtracker vlan-view
```

```
* R = Regular Vlan, P = Primary Vlan, C = Community Vlan  
I = Isolated Vlan, U = Invalid
```

VLAN	Type	VethPort	VM Name	Adapter Name	Mod
1	R	-	-	-	-
233	R	-	-	-	-
335	R	-	-	-	-
336	R	-	-	-	-
337	R	-	-	-	-
338	R	-	-	-	-
339	R	Veth3	gentoo-2	Net Adapter 3	3
		Veth4	gentoo-2	Net Adapter 4	3
		Veth5	gentoo-2	Net Adapter 2	3
340	R	-	-	-	-
341	R	-	-	-	-
400	R	Veth1	Fedora-VM2	Net Adapter 1	5
401	R	Veth1	Fedora-VM2	Net Adapter 1	5
402	R	Veth1	Fedora-VM2	Net Adapter 1	5
403	R	-	-	-	-
404	P	Veth6	Fedora-VM1	Net Adapter 1	4
405	C	Veth2	Fedora-VM2	Net Adapter 3	5
406	I	Veth7	Fedora-VM1	Net Adapter 2	4

vTracker VM VNic view

- Provides information about all the vNICs (virtual network interface cards) adapters that are managed by the Cisco Nexus 1000V switch
- **show vtracker vm-view vnic [module *number* | vm *name*]**

```
VSM-N1k# show vtracker vm-view vnic
```

```
* Network: For Access interface - Access vlan, Trunk interface - Native vlan,  
VXLAN interface - Segment Id.
```

Mod	VM-Name HypvPort	VethPort Adapter	Drv Type Mode	Mac-Addr IP-Addr	State	Network	Pinning
3	gentoo-2 1025	Veth3 Adapter 3	Vmxnet3 access	0050.56b5.37de n/a	up	339	Eth3/8
3	gentoo-2 1026	Veth4 Adapter 4	E1000 access	0050.56b5.37df n/a	up	339	Eth3/8
3	gentoo-2 1024	Veth5 Adapter 2	Vmxnet2 access	0050.56b5.37dd n/a	up	339	Eth3/8
4	Fedora-VM1 4258	Veth7 Adapter 2	E1000 pvlan	0050.56bb.4fc1 10.104.249.49	up	406	Eth4/3
5	Fedora-VM2 100	Veth1 Adapter 1	E1000 trunk	0050.56b5.098b n/a	up	1	Po9
5	Fedora-VM2 3232	Veth2 Adapter 3	E1000 pvlan	0050.56b5.098d 10.104.249.60	up	405	Po9

```
VSM-N1k#
```

vTracker VMotion view

- The vMotion view provides information about all the ongoing (if any) as well as previous VM migration events. However, only VMs that are currently being managed by the Cisco Nexus 1000V switch are displayed in the output
- **show vtracker vmotion-view {now | last *number 1-100*}**

```
show vtracker vmotion-view [now | last <1-100>]

VSM-N1k# show vtracker vmotion-view count 20
Note: Command execution is in progress..

Note: VM Migration events are shown only for VMs currently
      managed by Nexus 1000v.

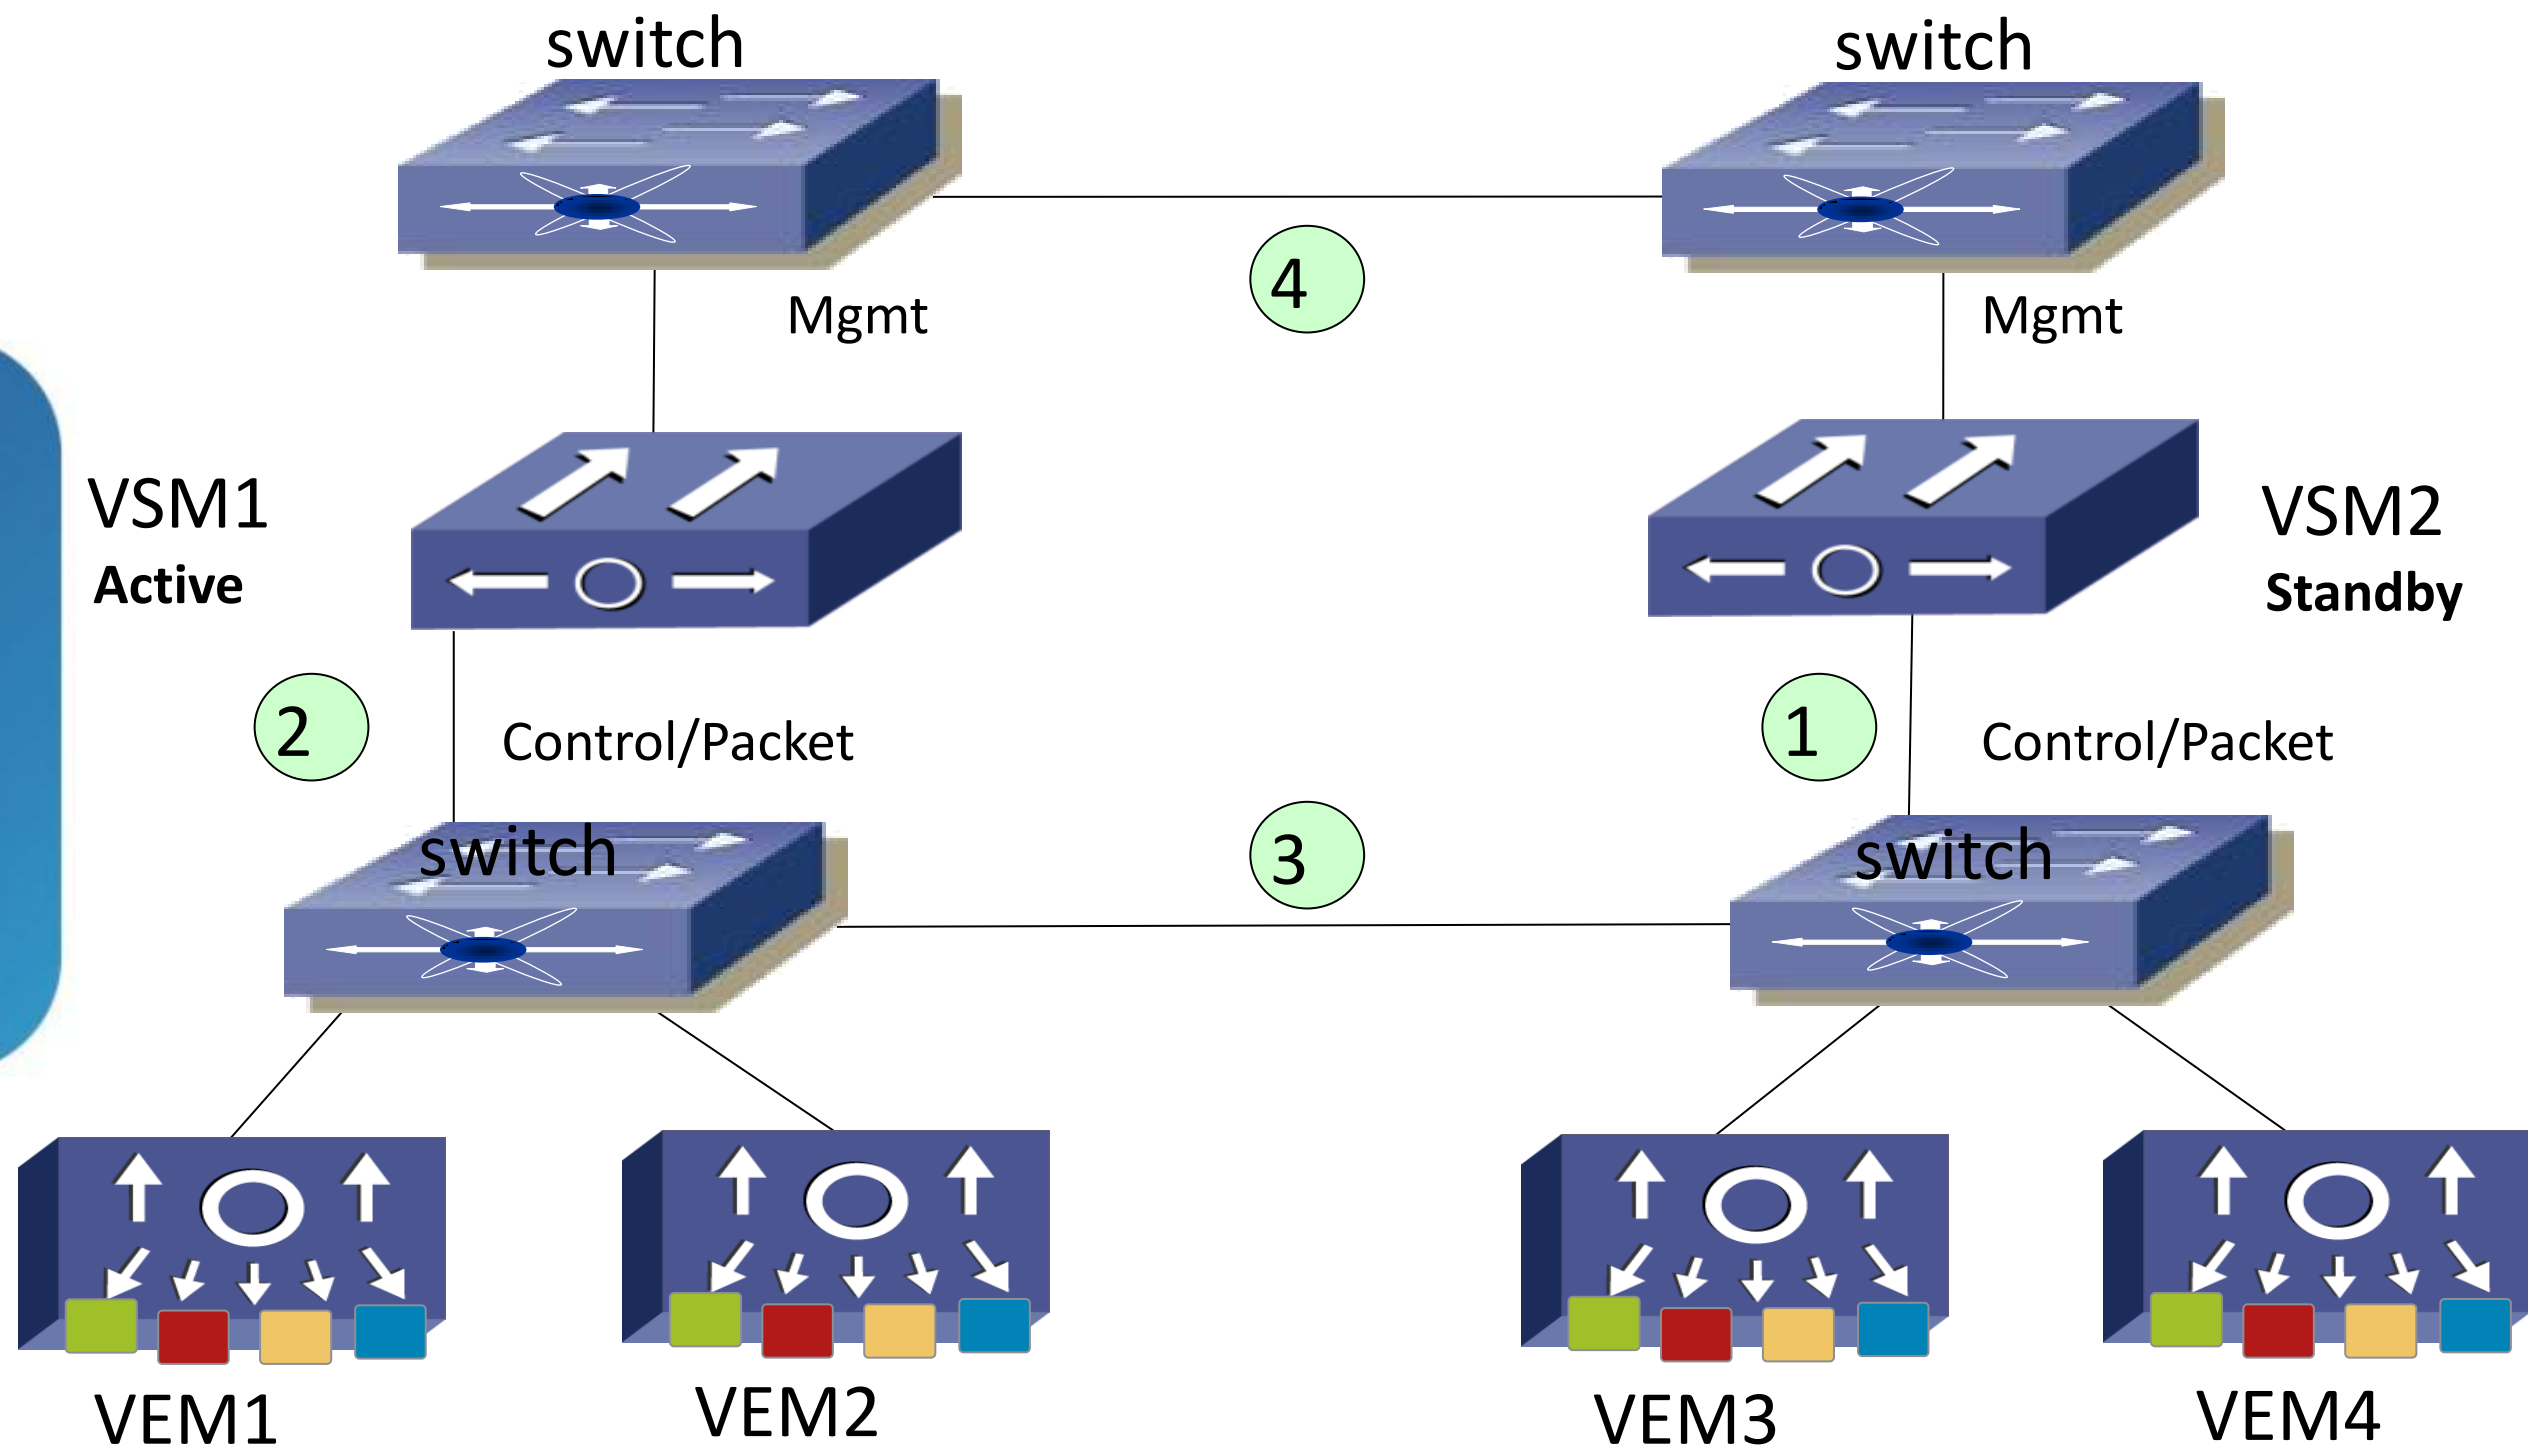
* '-' = Module is offline or no longer attached to Nexus1000v DVS
```

VM-Name	Src Mod	Dst Mod	Start-Time	Completion-Time
rk-ubt-1-0046	6	4	Mon Sep 3 10:42:27 2012	OnGoing
rk-ubt-1-0045	6	4	Mon Sep 3 10:42:27 2012	OnGoing
rk-ubt-1-0031	6	4	Mon Sep 3 10:42:27 2012	Mon Sep 3 10:44:10 2012
rk-ubt-1-0021	6	4	Mon Sep 3 10:42:27 2012	Mon Sep 3 10:43:42 2012

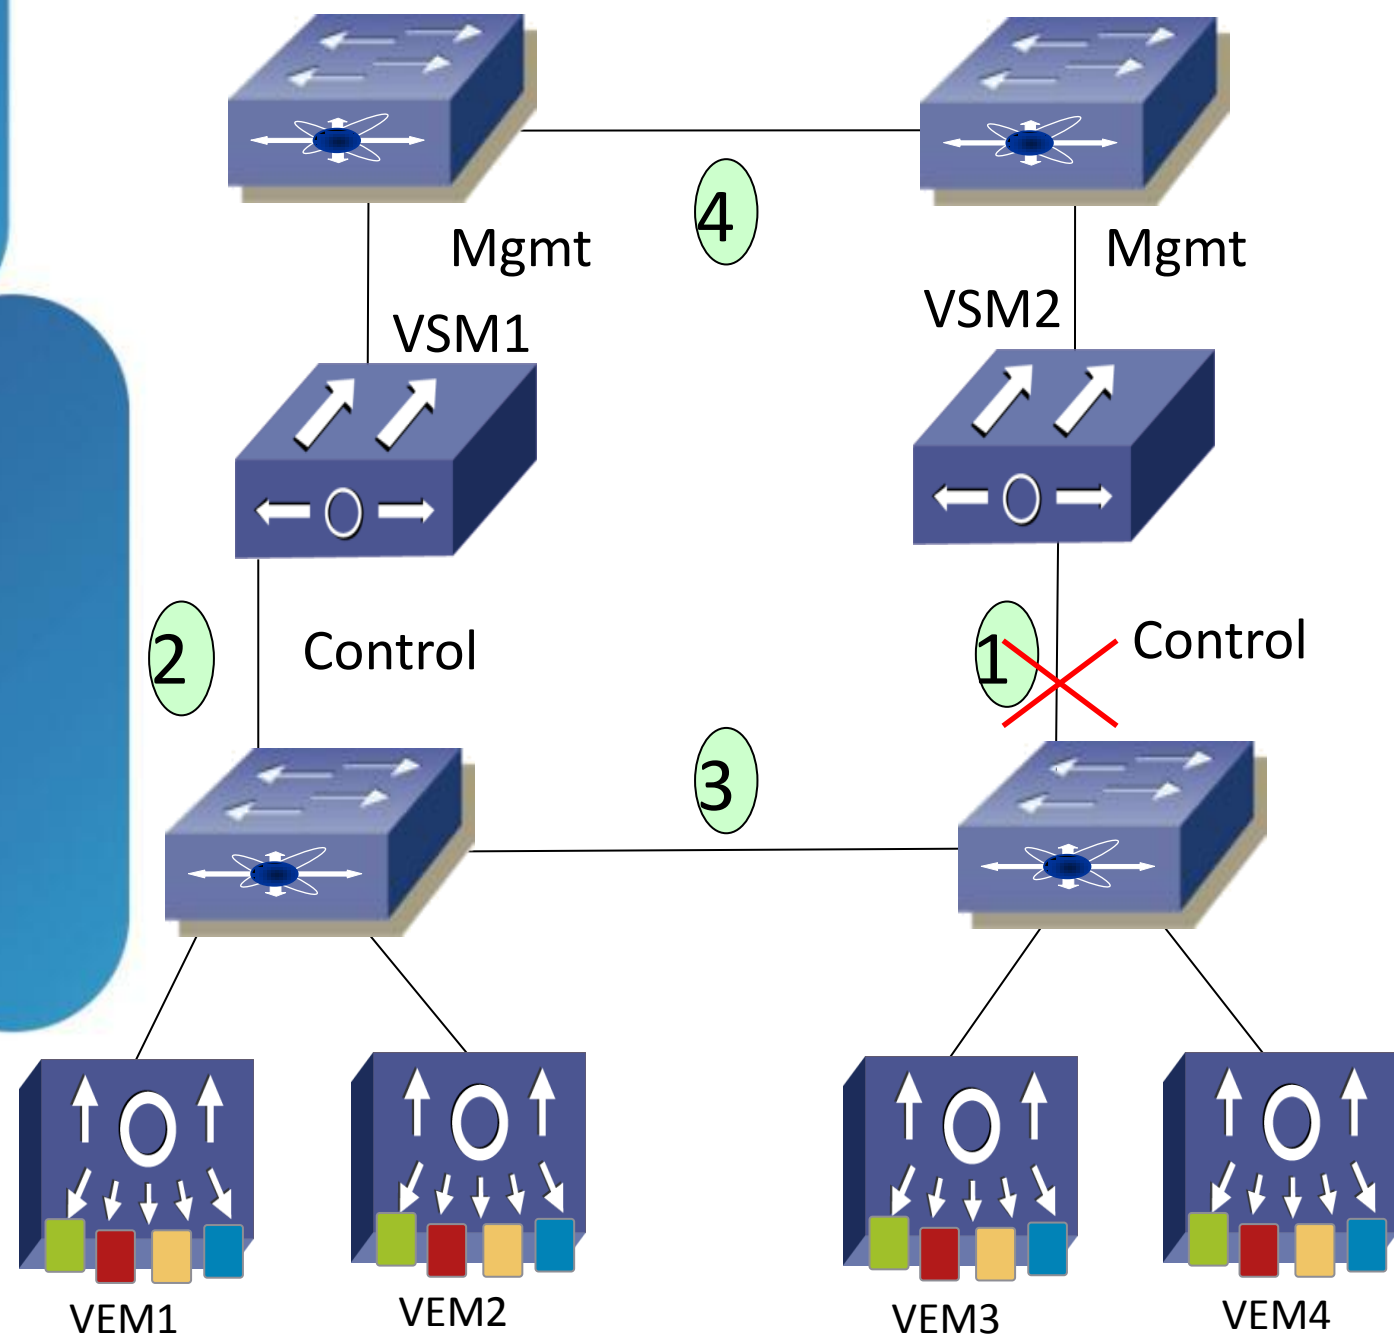
Appendix 3: More HA scenarios



Reference Topology



Fail Scenario 1



Failed Interface

#1- VSM2 Control


Effect: (when communication lost)

- VSMs use mgmt heartbeat
- VSM1 remains primary
- VSM2 remains standby
- VEMs stay connected

Exit: (when communication restored)

- VSMs use control heartbeat
- VSM1 remains primary
- VSM2 remains standby

Fail Scenario 2



Failed Interface #2- VSM1 Control


Effect:

- VEMs 1-4 become headless
- VSMs use mgmt for heartbeat
- VSM1 remains primary
- VSM2 reboots and stays warm standby

Exit:

- VEMs Reconnect
- VSMs use control for heartbeat
- VSM1 remains primary
- VSM2 attaches as secondary

Fail Scenario 3



Failed Interface #3 - Split DVS


Effect:

- VEM3,4 become headless
- VSMs use mgmt for heartbeat
- VSM1 remains primary
- VSM2 does not reboot. Stays secondary

Exit:

- VEM3,4 reconnect
- VSM use control for heartbeat
- VSM1 remains primary
- VSM2 reboots, joins as secondary

Fail Scenario 4



Failed Interface

#2 & #4 - Split Brain


Effect:

- VSM2 becomes Active, taking over VEM1-4
- VSM1 stays active but drops all VEMs
- VSM1 and VSM2 use same IP address

Exit:

- VSM1 reboots
- VSM2 stays primary
- VEM1-4 stay connected to VSM2
- VSM1 joins as standby

Fail Scenario 5



Failed Interface

#1 & #4 - Split Brain

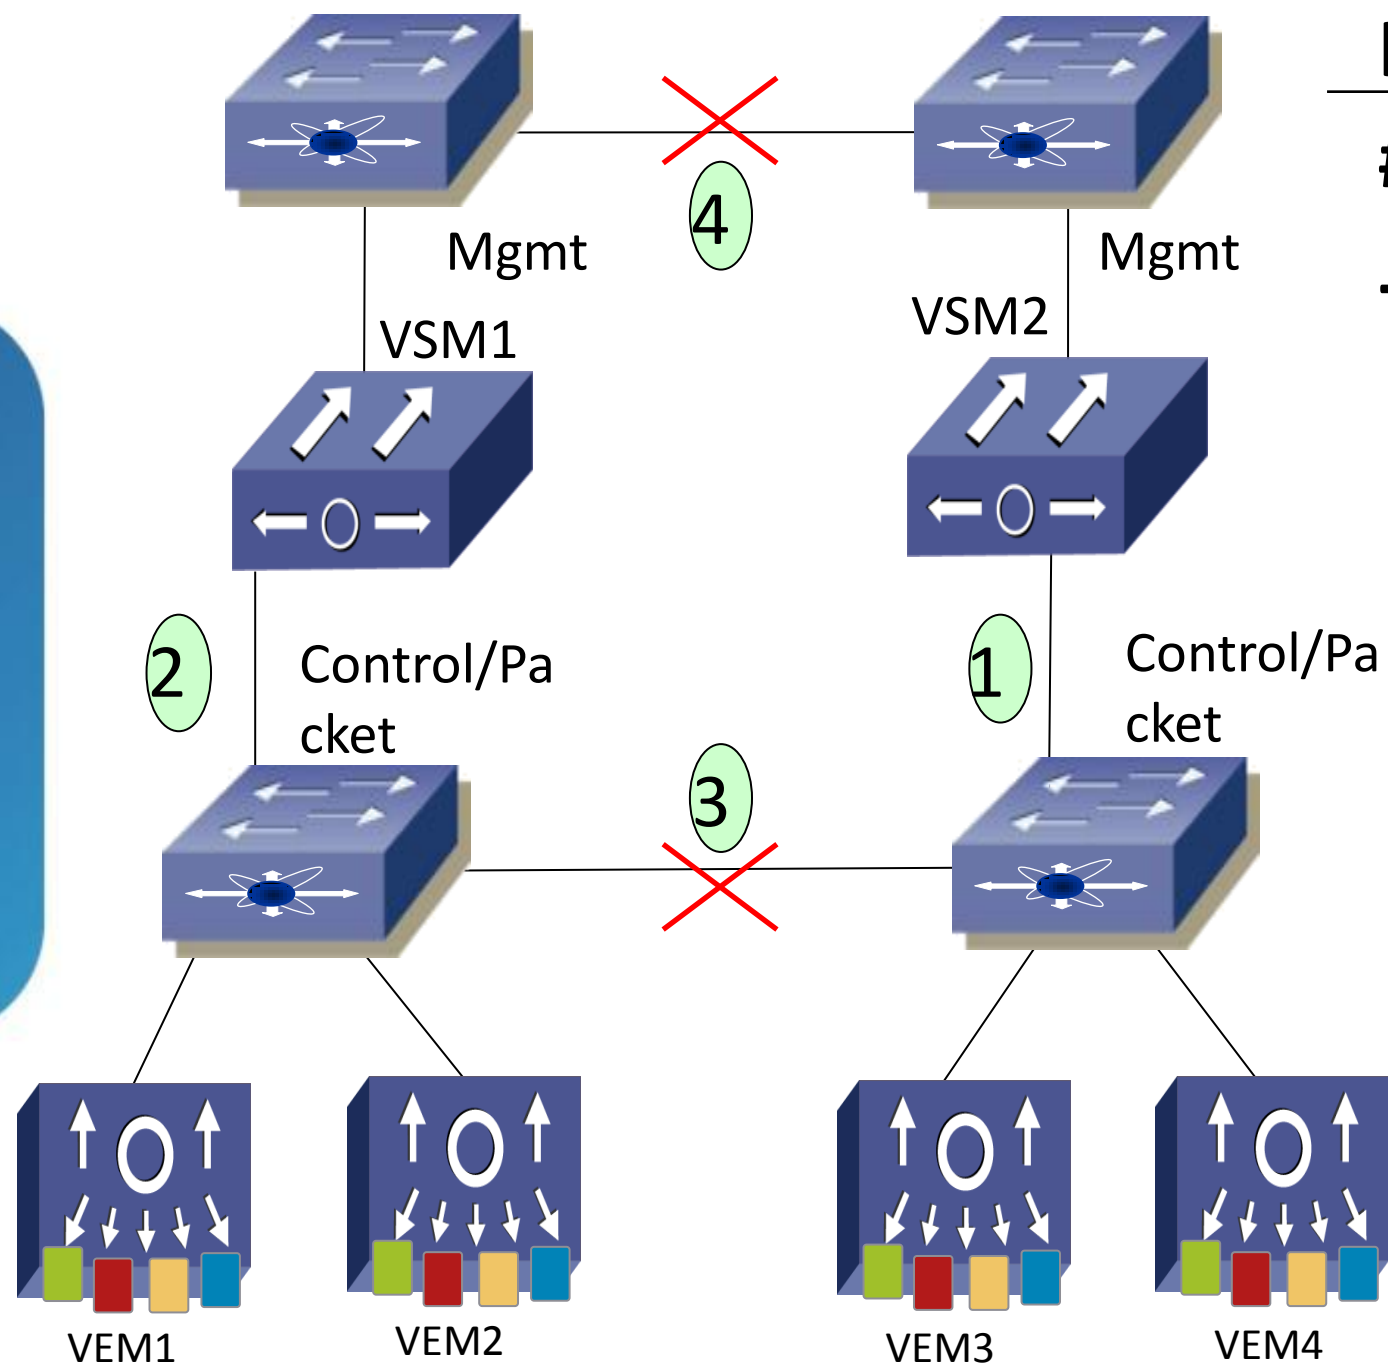
Effect:

- VSM2 becomes Active, but does not see any VEMs
- VSM1 stays active, handling all VEMs
- VSM1 and VSM2 use same IP address

Exit:

- VSM2 reboots
- VSM2 joins as standby
- VEM1-4 never change from VSM1

Fail Scenario 6



Failed Interface

#3 & #4 - Split Brain

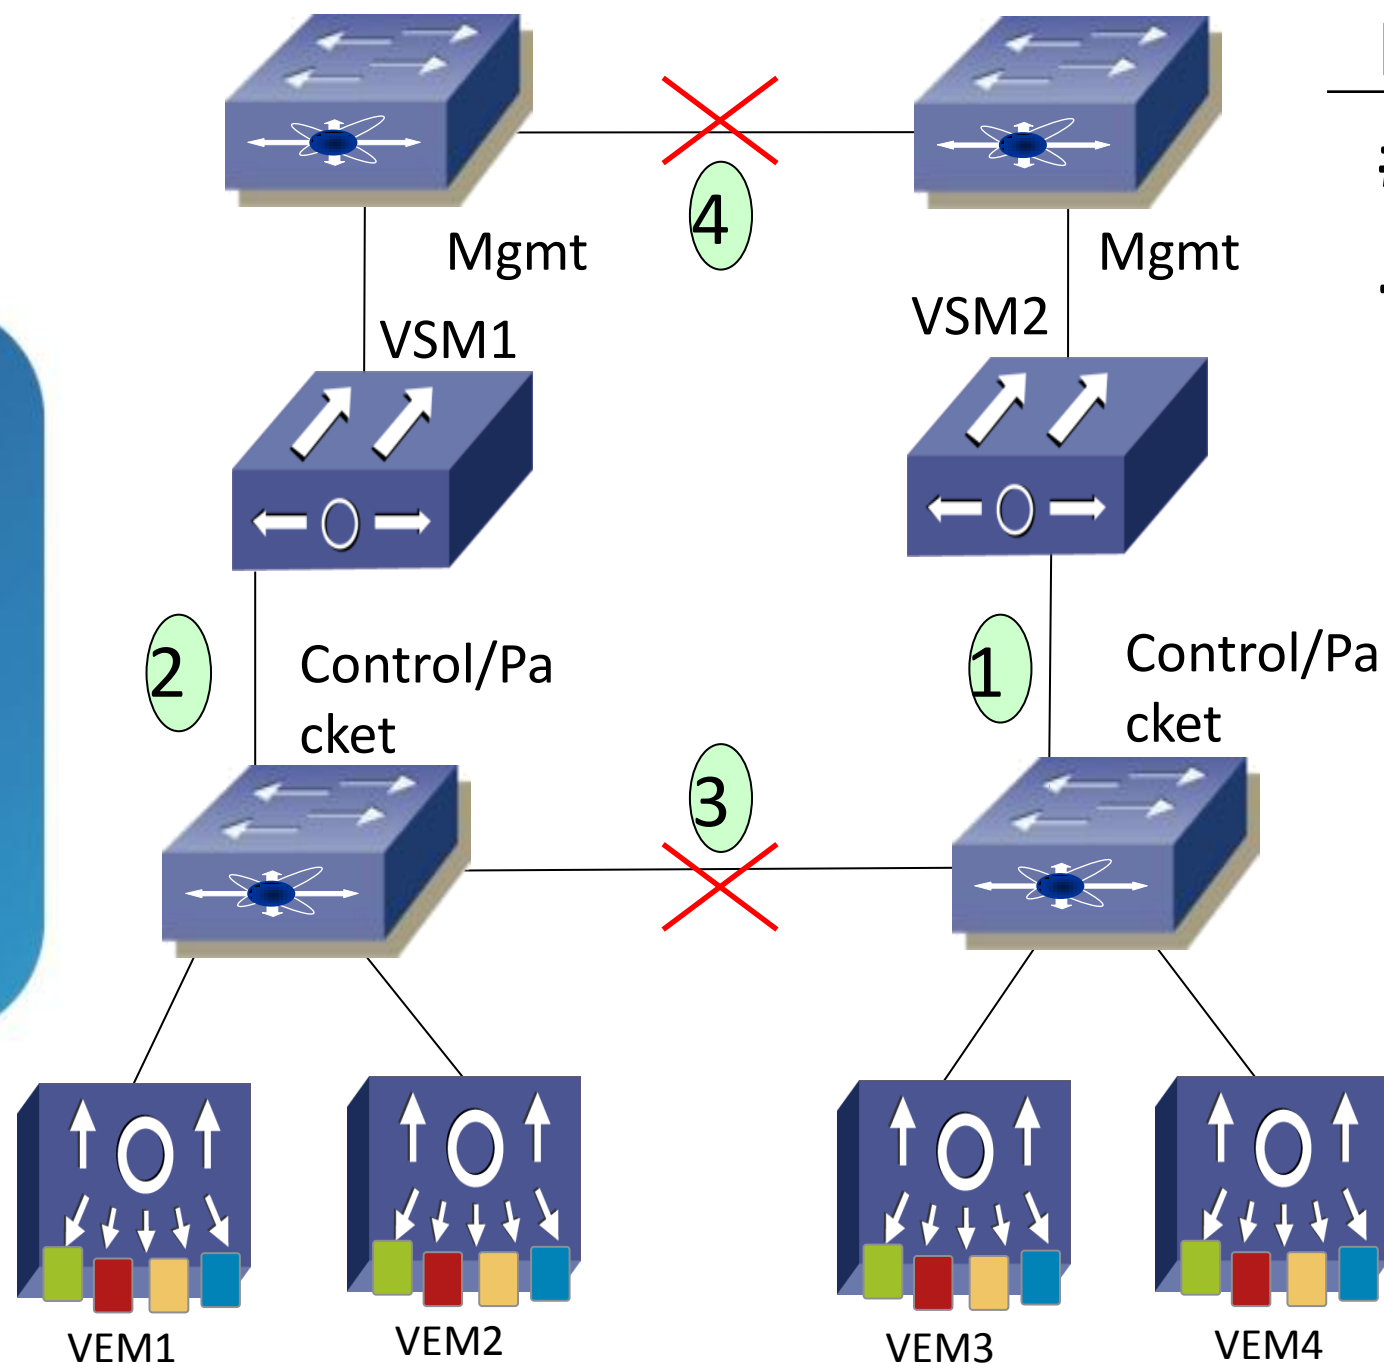
Effect:

- VSM2 becomes Active, handling VEM3 and VEM4
- VSM1 stays active, handling all VEM1 and VEM2
- VSM1 and VSM2 use same IP address

Exit:

- VSM2 reboots
- VEM1-4 connect to VSM1
- VSM2 joins as standby

Fail Scenario 6 – what if changes are made



Failed Interface

#3 & #4 - Split Brain

Effect:

- Same effect as before split-brain
- if changes are made on VSM2 redundancy manager will notice and make it primary when links come back up
- Important note, just creating and changing configs on VSM2 is not enough. We need to syncs to vCenter.


Appendix 4: VEM-VSM Connectivity Troubleshooting



VEM L3 Troubleshooting - Steps

- VMK interface created on ESX/ESXi host
- Can ESX/ESXi host ping VSM control/mgmt interface
- Uplink profile created correctly
- L3 veth port-profile created correctly
- Check the opaque data
- Troubleshoot Layer 3 connectivity

Sample Topology



Verify VSM Settings

- Check SVS domain parameters

```
n1kv-13# show svcs domain
```

```
SVS domain config:
```

```
Domain id: 43
```

```
Control vlan: 1
```

```
Packet vlan: 1
```

```
L2/L3 Control mode: L3
```

```
L3 control interface: control0
```

- Verify control 0

```
n1kv-13# show run int control 0
```

```
interface control0
```

```
ip address 192.168.11.10/24
```

- Verify VRF default

```
n1kv-13# show ip route
```

```
IP Route Table for VRF "default"
```

```
0.0.0.0/0, ubest/mbest: 1/0, pending
```

```
*via 192.168.11.1, control0, [1/0], 4d23h,  
static
```

- Can the VSM ping the VMK interface

```
n1kv-13# ping 192.168.11.180 vrf default
```

```
PING 192.168.11.180 (192.168.11.180): 56 data bytes
```

```
64 bytes from 192.168.11.180: icmp_seq=0 ttl=63 time=1.082 ms
```

```
64 bytes from 192.168.11.180: icmp_seq=1 ttl=63 time=0.841 ms
```

Check the Uplink Port-Profile

- Uplink needs to allow VLANs 10 and 11
- VLANs 10 and 11 need to be system vlans

```
n1kv-13# show run port-profile uplink
port-profile type ethernet uplink
    vmware port-group
    switchport mode trunk
    switchport trunk allowed vlan 10-11,150-152
    channel-group auto mode on mac-pinning
    no shutdown
    system vlan 10-11
    state enabled
```

Check the veth Port-Profile

- VMK interface needs to be migrated to this port-profile
- Allow VLAN 11 and have system VLAN set
- It also must have capability l3control

```
n1kv-l3# show run port-profile L3-control-vlan11
port-profile type vethernet L3-control-vlan11
  capability l3control
  vmware port-group
  switchport mode access
  switchport access vlan 11
  no shutdown
  system vlan 11
  state enabled
```

- Each VMK VLAN needs a new port-profile

Check Opaque Data

- Is the right Opaque data getting pushed to the ESX/ESXi host?

```
# vemcmd show card  
  
Card UUID type 2: 414a3031-3341-5553-4538-31384e375a37  
Card name: cae-esx-186  
Switch name: n1kv-13  
Switch alias: DvsPortset-0  
Switch uuid: 48 68 29 50 e2 ba af 6c-13 72 14 bc 25 cf 3f 86  
Card domain: 43  
Card slot: 4  
VEM Tunnel Mode: L3 Mode  
L3 Ctrl Index: 49  
L3 Ctrl VLAN: 11  
VEM Control (AIPC) MAC: 00:02:3d:10:2b:03  
VEM Packet (Inband) MAC: 00:02:3d:20:2b:03  
VEM Control Agent (DPA) MAC: 00:02:3d:40:2b:03  
VEM SPAN MAC: 00:02:3d:30:2b:03  
Primary VSM MAC : 00:50:56:af:49:30  
Primary VSM PKT MAC : 00:50:56:af:59:a1  
Primary VSM MGMT MAC : 00:50:56:af:47:51  
Standby VSM CTRL MAC : 00:50:56:af:61:e7
```

Should match VLAN
defined veth port-profile

Should match MAC of
control 0 or mgmt 0

Check Counters for Heartbeat on VSM

```
n1kv-13# show module vem counters
```

```
-----  
Mod  InNR  OutMI  InMI  OutHBeats  InHBeats  InsCnt  RemCnt  Crit Tx Errs  
-----  
  3    2    2    2    80768    77667    2      1      0  
  4    1    1    1     529     529    1      0      0  
  5    3    3    3   508890   175821    3      2      0  
  6    2    2    2   508890    77632    2      1      0
```

InNR - NodeID requests received count

OutMI - Module Insert Start requests sent to VEM

InMI - Module Insert Start responses received from VEM

OutHBeats - Number of HBs which have been broadcast by VSM

InHBeats - Number of HBs received from this VEM

View Heartbeat Messages on VEM

- Use vempkt on the ESX/ESXi host
 - vempkt capture ingress/egress vlan 11
 - Let it run
 - vempkt cancel capture all
 - vempkt display detail all
- vempkt can now export to a pcap file
 - vempkt pcap export <filename>
- Look for heartbeat messages from VSM

Verifying Modules with “vem status”

- “vem status -v”
- Reboot the ESX host if not all the modules have loaded.

```
~ # vem status -v  
Package vssnet-esxmn-ga-release  
Version 4.2.1.1.5.1.0-3.0.1  
Build 1  
Date Mon Jan 30 18:38:49 PST 2012
```

```
Number of Passthru NICs are 0  
VEM modules are loaded
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	128	4	128	1500	vmnic0
DVS Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
n1kv-13	256	13	256	1500	vmnic1

```
Number of Passthru NICs are 0  
VEM Agent (vemdpa) is running
```

Appendix 5: VM Connectivity Troubleshooting



VM Connectivity Troubleshooting - Steps

- Verify VSM configuration/Port-Profiles
- Verify VEM configuration
- Check Upstream Device
- Collect VEM statistics
- Vempkt capture

Verify VSM configuration/Ports

```
DCN-Nexus1000V# show int veth 19
Vethernet19 is up
  Port description is gchami-Win7, Network Adapter 1
  Hardware: Virtual, address: 0050.5680.75e2 (bia 0050.5680.75e2)
  Owner is VM "gchami-Win7", adapter is Network Adapter 1
  Active on module 10
  VMware DVS port 901
  Port-Profile is Management-174
  Port mode is access
0 Input Packet Drops 0 Output Packet Drops
<more>

DCN-Nexus1000V# show run int veth 19
interface Vethernet19
  inherit port-profile Management-174
  description gchami-Win7, Network Adapter 1
  vmware dvport 901 dvswitch uuid "bc f0 00 50 6c 88 70 79-7d 88 4b b8 39 c5 6f 9a"
  vmware vm mac 0050.5680.75E2
```

VM(veth) Port-Profile Troubleshooting

- Be careful modifying veths directly
 - Modify the port-profile and not the veth
 - VSM remembers VM veths until they are deleted
 - Changes to a veth will stick around until the VM nic is deleted
- VM to veth mapping does not change until
 - NIC is removed from the VM
 - NIC is reassigned to another port-profile
- Use VMware VMXNET3 NIC type over E1000

Check for recent vMotions

```
Cisco-Live-N1K# show vtracker vmotion-view last 20
```

```
Note: VM Migration events are shown only for VMs currently  
managed by Nexus 1000v.
```

```
* '-' = Module is offline or no longer attached to Nexus1000v DVS
```

```
-----  
VM-Name          Src Dst Start-Time          Completion-Time  
                  Mod Mod  
-----  
gchami-Win7      3  4   Wed Nov 28 07:00:55   Wed Nov 07 12:05:15
```

Uplink(eth) Port-Profile Troubleshooting

- Port-Profiles with multiple NICs need a port-channel
- Causes duplicate packets
- Kicks in déjà vu driver
- Requires extra CPU processing
- Fills the logs
- Example: Eth 6/1 and Eth 6/3 added to below Port-Profile

WRONG

```
port-profile type ethernet uplink-nopc
vmware port-group
switchport mode trunk
switchport trunk allowed vlan 1-3967,4048-4093
no shutdown
system vlan 11
state enabled
```

RIGHT

```
port-profile type ethernet uplink-nopc
vmware port-group
switchport mode trunk
switchport trunk allowed vlan 1-3967,4048-4093
channel-group auto mode on mac-pinning
no shutdown
system vlan 11
state enabled
```

Uplinks with Overlapping VLANs

- VLANs cannot overlap on uplink port-profiles
- Example: Two port-profiles
 - Overlap of VLAN 174
 - Assign 6/1 to uplink1 and 6/3 to uplink2
 - Which uplink will a VM on VLAN 174 use?
 - No way to pin a VM to an uplink port-profile
 - Use VPC MAC Pinning

```
port-profile type ethernet uplink1
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 1,10,11,174
  no shutdown
  state enabled
```

```
port-profile type ethernet uplink2
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 13,14,174
  no shutdown
  state enabled
```

Verify VEM Configuration

```
~#vemcmd show port
```

LTL	VSM Port	Admin	Link	State	PC-LTL	SGID	Vem Port
21	Eth10/5	UP	UP	FWD	305	4	vmnic4
49	Veth56	UP	UP	FWD	0	4	vmk1
50	Veth12	UP	UP	FWD	0	4	test1.eth0
51	Veth51	UP	UP	FWD	0	4	test2.eth0
52	Veth19	UP	UP	FWD	0	4	gchami-Win7.eth0

```
<more>
```

```
~#vemcmd show port vlans
```

LTL	VSM Port	Mode	Native VLAN	VLAN State	Allowed Vlans
21	Eth10/5	T	1	FWD	1,174,700-730
49	Veth56	A	701	FWD	701
50	Veth12	A	705	FWD	705
51	Veth51	A	174	FWD	174
52	Veth19	A	174	FWD	174

Verify VEM Configuration

```
~#vemcmd show bd 174
BD 174, vdc 1, vlan 174, 17 ports
Portlist:
    21  vmnic4
    51  test2.eth0
    52  gchami-Win7.eth0

vemcmd show 12 174
Bridge domain 174 brtmax 4096, brtcnt 44, timeout 300
Flags:  P - PVLAN  S - Secure

```

Type	MAC Address	LTL	timeout	Flags	PVLAN
Static	00:50:56:80:7a:87	56	0		
Static	00:50:56:80:7a:b5	53	0		
Static	00:50:56:80:75:e2	52	0		
Dynamic	00:50:56:80:1b:f5	305	3		
Dynamic	00:50:56:80:1b:fd	305	0		
Dynamic	00:50:56:80:1b:fc	305	0		

Port Channels – How to Tell Pinning

- Can run from the VSM now
- No need to run command on the VEM

```
n1kv-13# show int virtual pinning module 5
```

Veth	Pinned Sub Group id	Associated PO interface	List of Eth interface(s)
Veth19	4	Po5	Et10/4

Check the Neighbouring switch

- Show CDP Neighbour
- Show CDP Neighbour Interface Eth10/4 detail

```
DCN-Nexus1000V# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID                Local Intrfce Hldtme Capability Platform      Port ID
Upstream-N5K-1 (SSI145000MF) Eth10/4      135      S I s      N5K-C5548P      Eth1/11
```

Check Upstream Switch

```
Upstream-N5K-1# show int trunk vlan 174
```

```
<snippet>
```

```
-----  
Port          STP Forwarding  
-----
```

```
Eth1/11      1,174,180,700-729,1002
```

```
<snippet>
```

```
Upstream-N5K-1# show vlan id 174
```

```
<snippet>
```

```
VLAN Name          Status      Ports  
-----  
174Traffic          active     Eth1/11
```

```
<snippet>
```

Check the Port-Channel

- Do not use Network State Tracking(NST) with LACP
- LACP Port-Channel configured on the upstream switches
- Port-profile created with “channel-group auto mode active”
- On the VEM
 - vemcmd show lacp
- On the VSM and Upstream Switch
 - show port-channel summary
 - show lacp counters/neighbour

Are you seeing LACP PDUs?

Appendix 6: Nexus 1010 and 1010-X



Cisco Nexus 1010 and 1010-X

- Based off UCS C200 M2 server
 - Provide 6 x 1GB network connections
 - 4 distinct topologies
 - 1 flexible topology
- Virtual Service Blade (VSB) Support
 - 1010 supports 6
 - 1010-X supports 10
- Current supported VSBs
 - Nexus 1000V VSM
 - Virtual Security Gateway (VSG)
 - Network Analysis Module (NAM)
 - Data Centre Network Manager (DCNM)

Cisco Nexus 1010

- Must be deployed in pairs
 - No option for a standalone 1010
- Deploy in the Aggregation Layer
- Must be in the same L2 domain for management and control
- Uses same HA mechanism as VSM with domain id and control vlan
 - Do not overlap the domain id between a 1010 and a VSM
- Not supported
 - Primary and Secondary VSM on same 1010
 - Primary VSM on ESX and Secondary VSM on 1010 or vice versa
 - 1010s split between data centres

VSB Import/Export

- Previously no way to save off a VSB
- Works with VSM, NAM, and VSG
- Can Import/Export both primary and secondary
- Export requires that VSB be shutdown
- Images are stored in “export-import/” dir on bootflash
 - Can be manually copied off to remote storage

Using Export to backup a VSM

- Shutdown primary VSM

- Secondary VSM will take over and run Nexus 1000v control plane
- `f340-33-09-n1010-1(config)# virtual-service-blade training`
- `f340-33-09-n1010-1(config-vs-b-config)# shutdown primary`

- Export VSB on 1010

- `f340-33-09-n1010-1(config-vs-b-config)# export primary`
- Note: export started..
- Note: please be patient..
- Note: export completed...

- Verify

- `f340-33-09-n1010-1(config-vs-b-config)# dir bootflash:///export-import/4`
- `147779575 Oct 18 02:47:10 2011 Vdisk4.img.tar.00`

Using Import to Restore a VSB

- Copy the backup to bootflash

- f340-33-09-n1010-1# copy scp://root@172.18.217.165/root/Vdisk4.img.tar.00
bootflash:export-import vrf management

- Import the image

- f340-33-09-n1010-1 (config)# virtual-service-blade training

- f340-33-09-n1010-1 (config-vs-b-config)# import primary
Vdisk4.img.tar.00

- Note: import started..

- Note: please be patient..

- Note: Import cli returns check VSB status for completion

- Verify

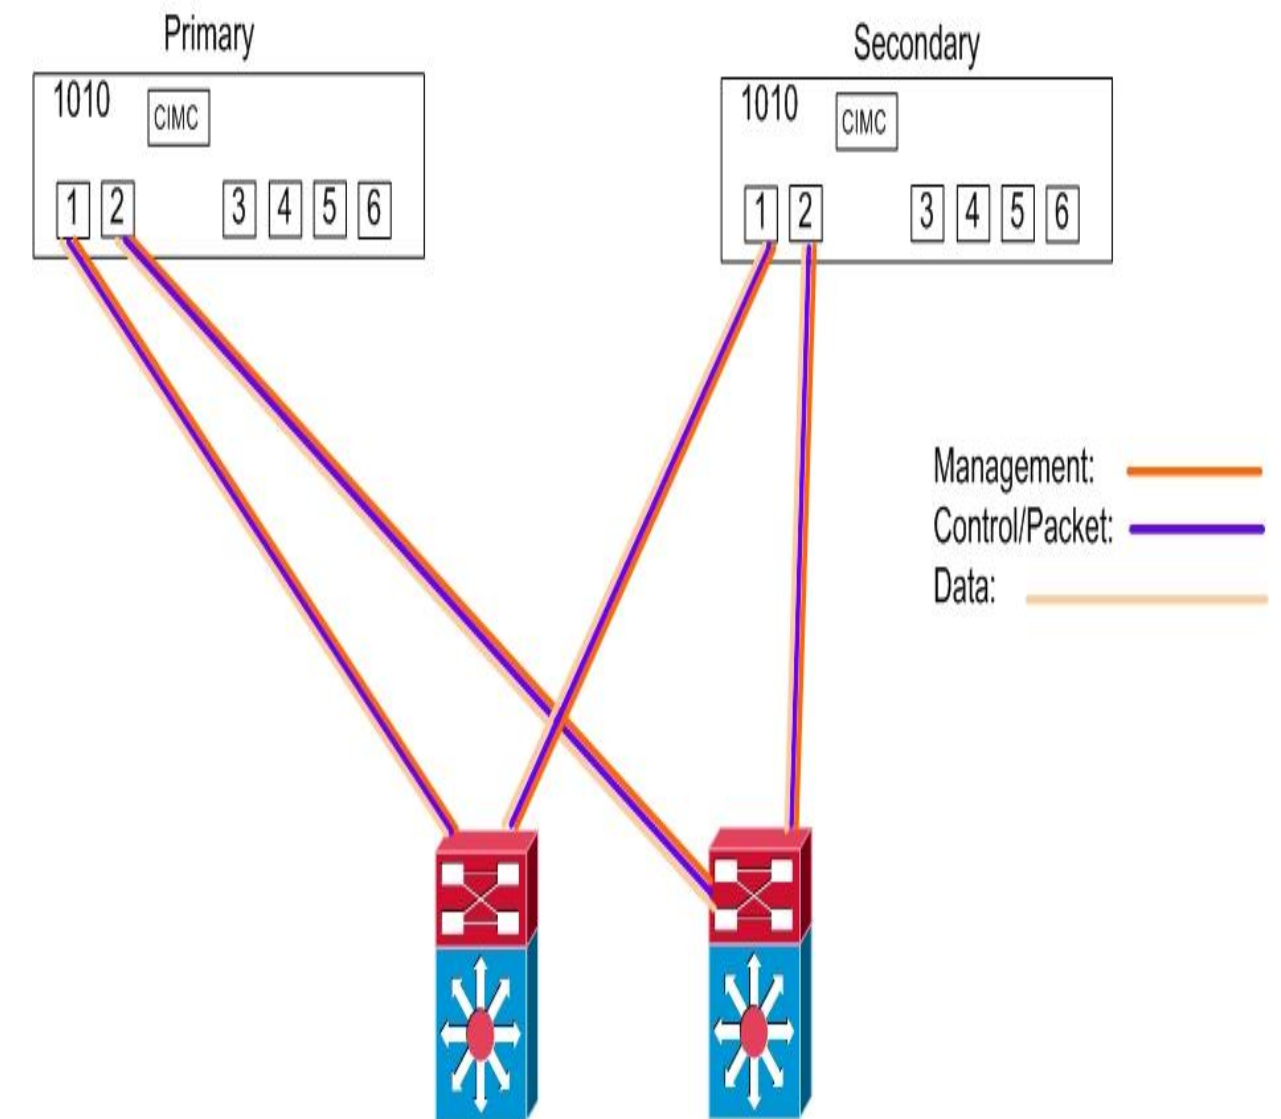
- f340-33-09-n1010-1 (config-vs-b-config)# show virtual-service-blade name
training

Nexus 1010 Network Classes and Topologies

- Network traffic is classed into 3 categories
- Management
 - Carries the mgmt 0 interface of the 1010
 - Carries the mgmt 0 traffic for all VSMs installed
 - No longer required that VSM mgmt and 1010 mgmt be on the same subnet
- Control
 - Carries all the control and packet traffic for the VSMs installed on the 1010
 - Carries control traffic for HA between primary and secondary 1010
- Data
 - Used by Virtual Service Blades (VSB) other than VSM
- 5 Network Topologies you can choose

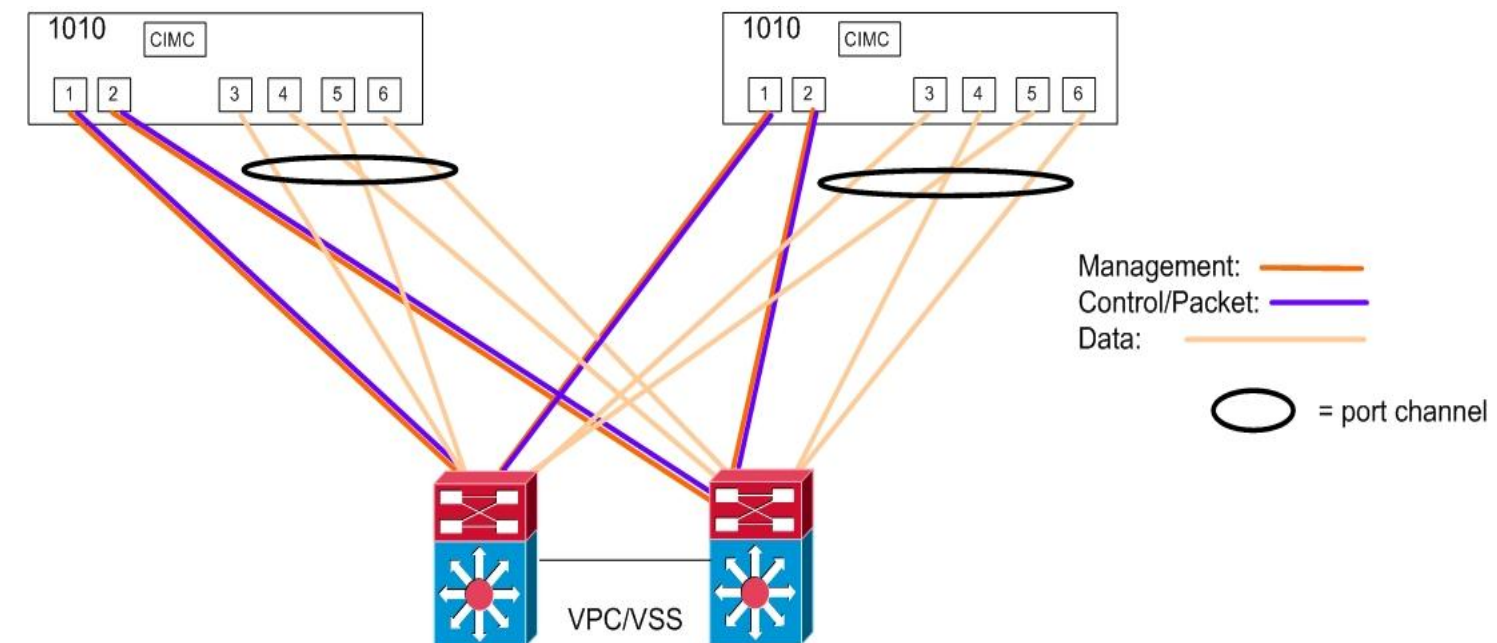

Topology 1 – Single Uplink

- Traffic will only flow over one link (ex eth1)
- Should eth1 fail traffic will fail to eth2
- Should eth2 fail services will fail to secondary 1010
- Port-channel is not supported
- 1GB of throughput
- Best practice
 - eth1 and eth2 connect to different switches for redundancy





Topology 2 – Two uplinks 1

- Management and Control are still combined
 - 1GB throughput
- Data traffic for NAM for best redundancy and throughput
- Without VPC/VSS
 - **Eth3 and Eth5** are paired LACP port-channel
 - **Eth4 and Eth6** are paired LACP port-channel
 - 2 GB of throughput
- With VPC/VSS
 - One big LACP port-channel
 - 4 GB of throughput



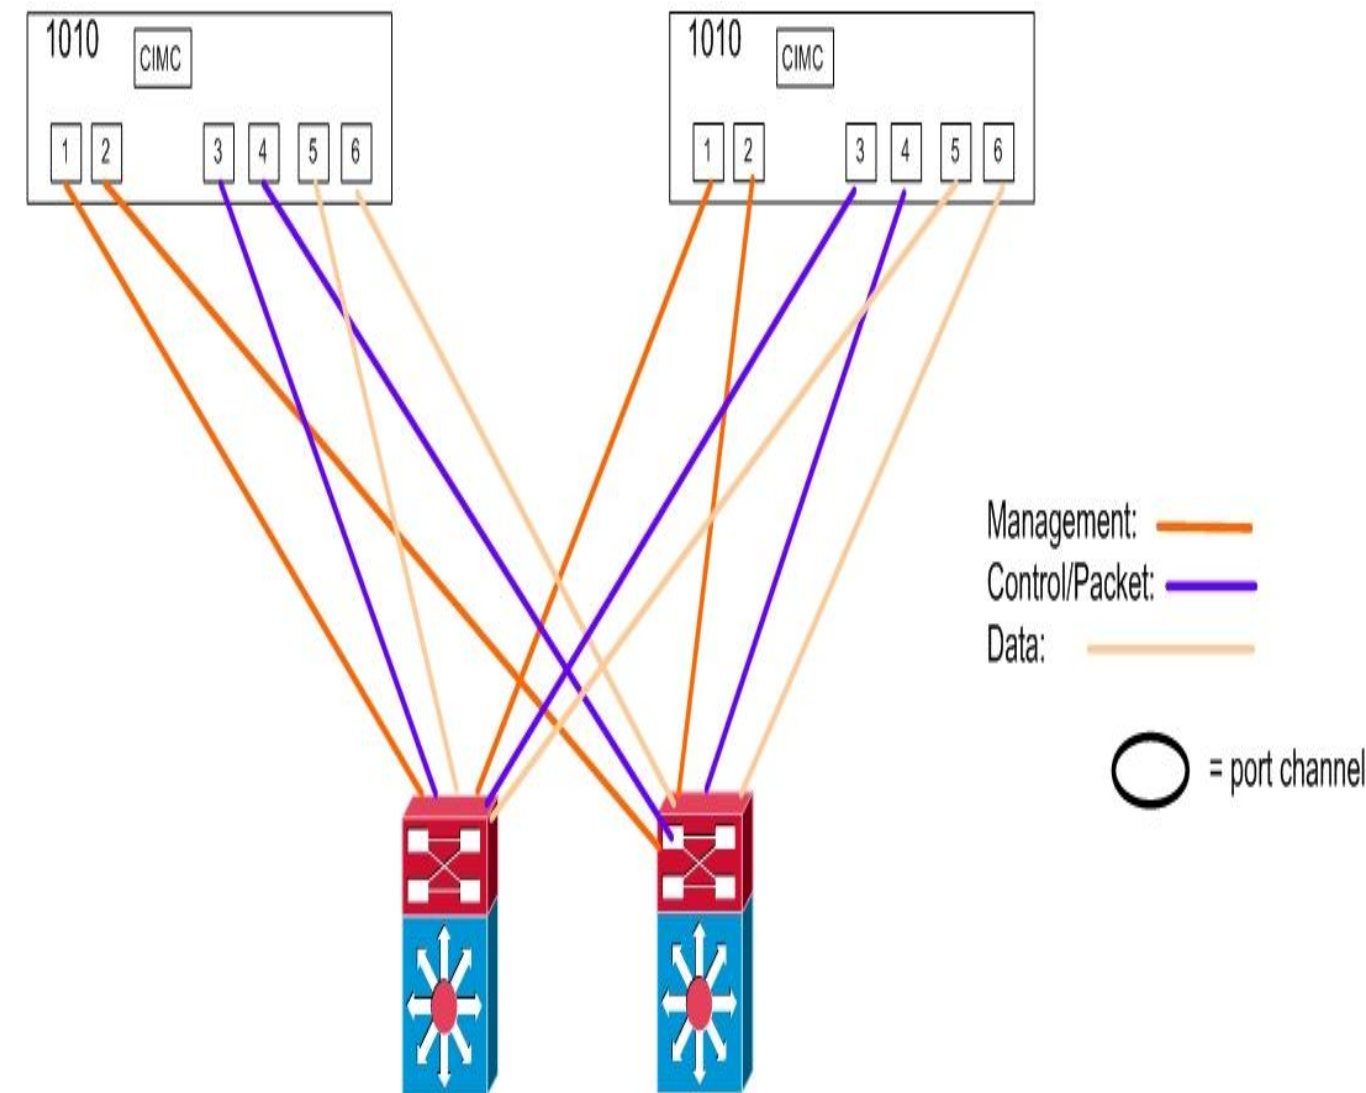
Topology 3 – Two uplinks 2

- Management dedicated interfaces
 - 1GB throughput
- Control and Data traffic combined
- Without VPC/VSS
 - **Eth3 and Eth5** are paired LACP port-channel
 - **Eth4 and Eth6** are paired LACP port-channel
 - 2 GB of throughput
- With VPC/VSS
 - One big LACP port-channel
 - 4 GB of throughput



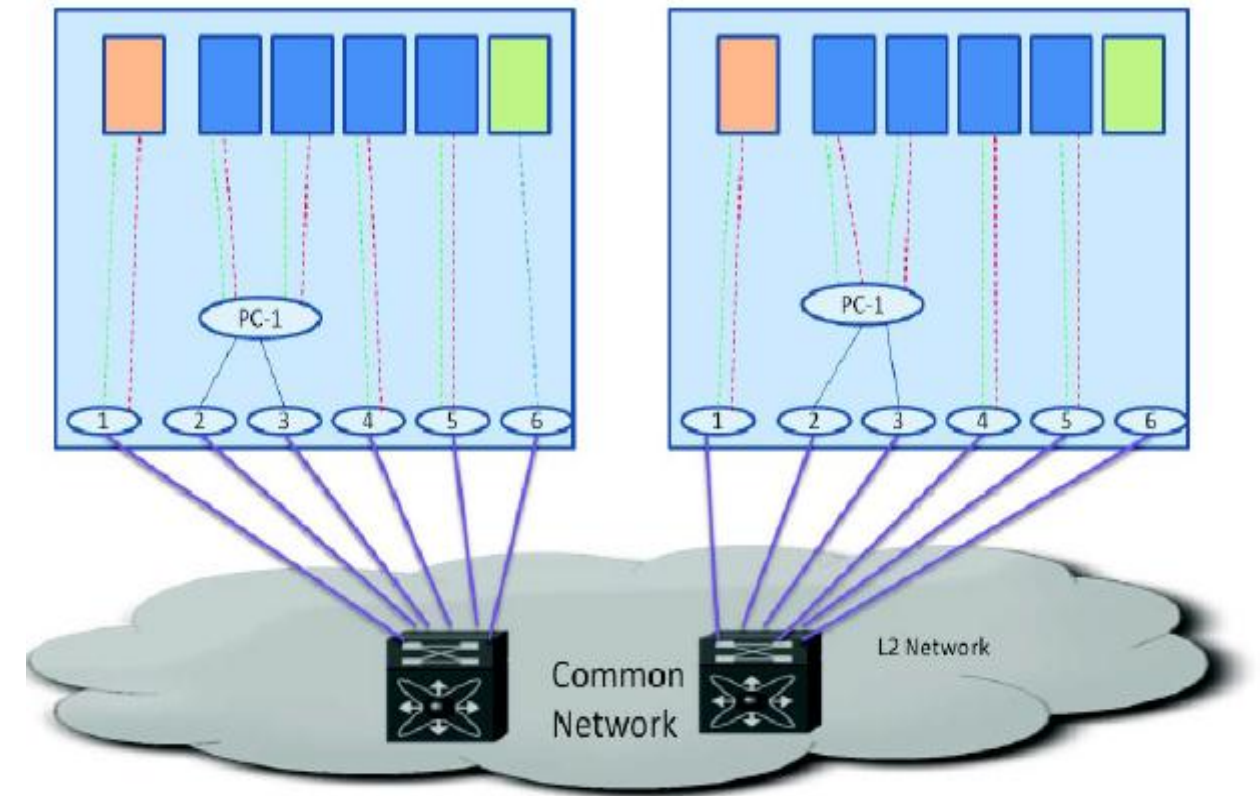
Topology 4 – Three Uplinks

- Management dedicated interfaces
 - Eth1 and Eth2 dedicated for Management
 - 1GB of throughput
- Control dedicated interfaces
 - Eth3 and Eth4 dedicated for Control
 - 1GB of throughput
- Data gets dedicated interfaces
 - Eth5 and Eth6 dedicated for Data
 - 1GB of throughput
- Best for solutions where Control and Data need dedicated links




Topology 5 – Flexible

- Assign a link explicitly to a VSB
 - links per VSB interface
- Share links across VSBs
- Create port-channels
 - Specific to a VSB
 - Shared to multiple VSBs



```
switch(config)#show network summary
```

Port	State		Uplink-Interface		Speed	RefCnt	MTU	Nat-Vlan		
	Oper	Admin	Oper	Admin					Oper	Admin
Gi1	up	up			1000	0	9000			
Gi2	up	up			1000	0	9000			
Gi3	up	up			1000	3	9000			
Gi4	down	up			1000	0	9000			
Gi5	down	up			1000	0	9000			
Gi6	down	up			1000	0	9000			
Po1	up	up			1000	13	9000			
VsbEth6/1	up	up	Gi3	Gi3	1000		9000			
VsbEth6/2	up	up	Gi3	Gi3	1000		9000			
VsbEth6/3	up	up	Gi3	Gi3	1000		9000			
control0	up	up	Po1	Po1	1000		9000			
mgmt0	up	up	Po1	Po1	1000		9000			



Recommendations

- If you are not planning on using vNAM
 - Topology 3 gives best bandwidth and redundancy for control VLAN
 - Negative is that is harder to configure
- Flexible allows any configuration
 - Recommend port-channels
 - Remember VSM latency is key over bandwidth
- Use VPC or VSS upstream if you have it

