



# Nexus 1000V Troubleshooting

**Louis Watta – Customer Assurance Engineer**

**Cisco Confidential**



## Nexus 1000V Public Webinar Series

Date	Business Sessions
22-Mar	Nexus 1000V Family Overview and Update
5-Apr	Virtual Network Services (vPath, NAM, vWAAS)
19-Apr	Virtual Security Gateway Introduction
3-May	Journey to the Cloud w/ N1KV: vCloud Director & Long Distance vMotion
17-May	Secure VDI with Nexus1000V & VSG

Date	Technical Sessions
29-Mar	Nexus 1000V New Features and Installation Overview
12-Apr	Nexus1010 Installation & Upgrade
26-Apr	Virtual Security Gateway Overview
10-May	Nexus 1000V Key Features
24-May	Nexus 1000V Troubleshooting

# Virtual Network Services

## Today's Agenda

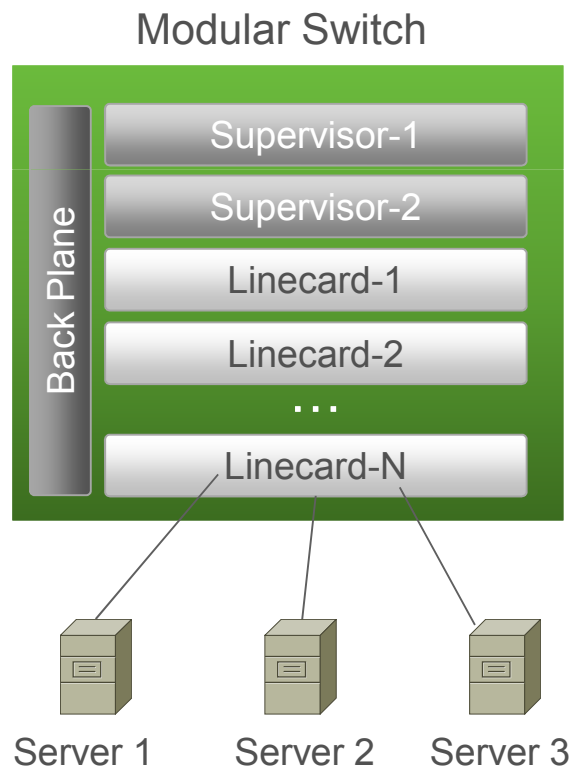
- Nexus 1000V Architecture – Bob Vizza
  - Virtual Supervisors and Ethernet Modules
  - Virtual Service Datapath (vPath)
- Nexus 1000V Troubleshooting – Louis Watta
- Q &A as we go...

# Nexus 1000V Architecture Overview




# Nexus 1000V Architecture

## Comparison to a Physical Switch




# Nexus 1000V Architecture

## Moving to a Virtual Environment




# Nexus 1000V Architecture

Supervisors → Virtual Supervisor Modules (VSMs)



# Nexus 1000V Architecture

Linecards → Virtual Ethernet Modules (VEMs)






# Nexus 1000V Architecture

VSM + VEMs = Nexus 1000V Virtual Chassis


- 64 VEMs per 1000V (connected by L2 or L3)
- 200+ vEth ports per VEM
- 2K vEths per 1000V
- Multiple 1000Vs can be created per vCenter



VSM: Virtual Supervisor Module  
VEM: Virtual Ethernet Module

# Nexus 1000V Architecture

Customer Request: Host VSMs on a Physical Appliance



- 200+ vEth ports per VEM
- 64 VEMs per 1000V
- 2K vEths per 1000V
- Multiple 1000Vs can be created per vCenter

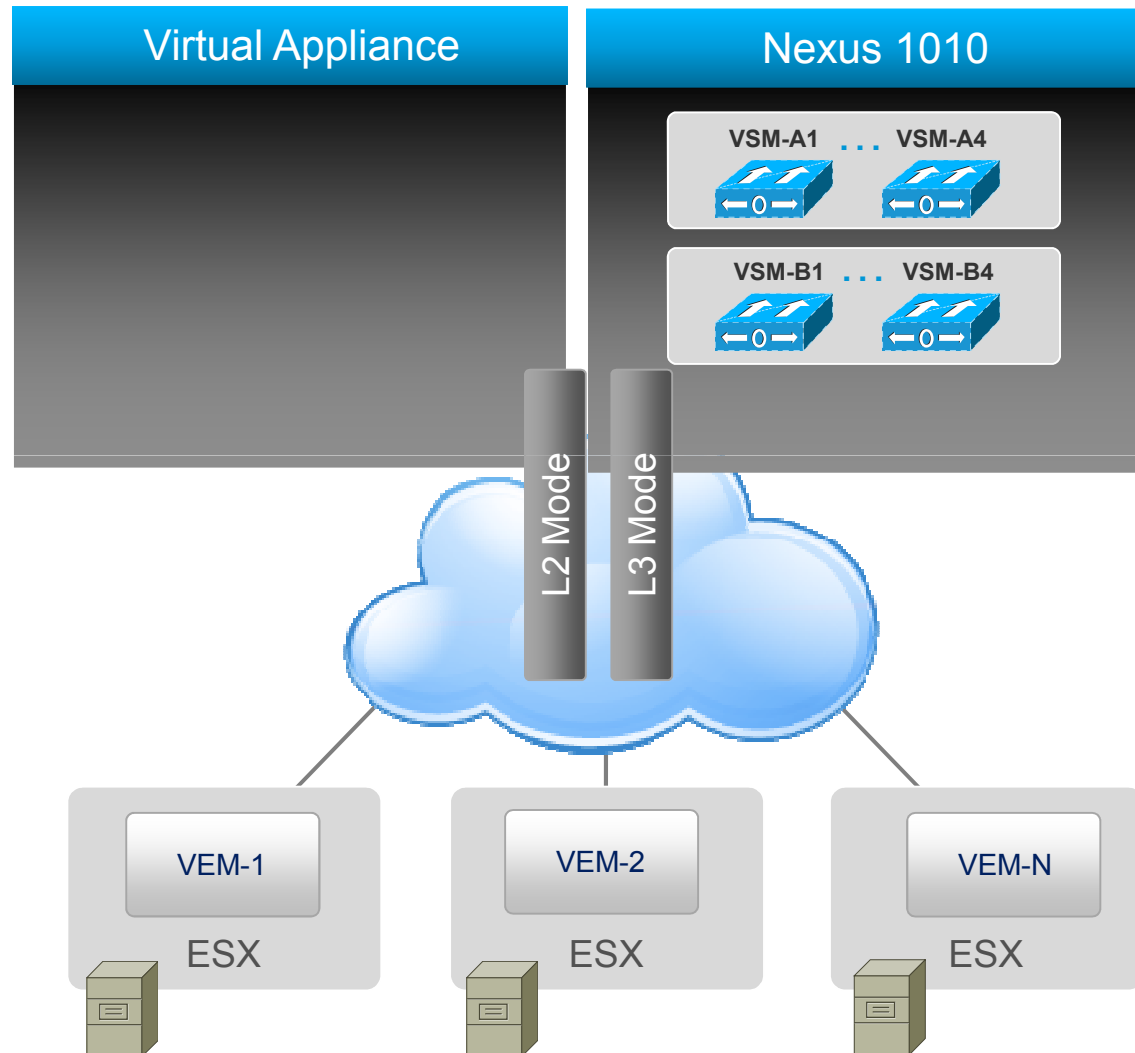
VSM: Virtual Supervisor Module

VEM: Virtual Ethernet Module

# Nexus 1000V Architecture

## VSMs hosted on a Physical Appliance: Nexus 1010

- Up to 4 VSMs per Nexus 1010
- Nexus 1010s deployed in redundant pair

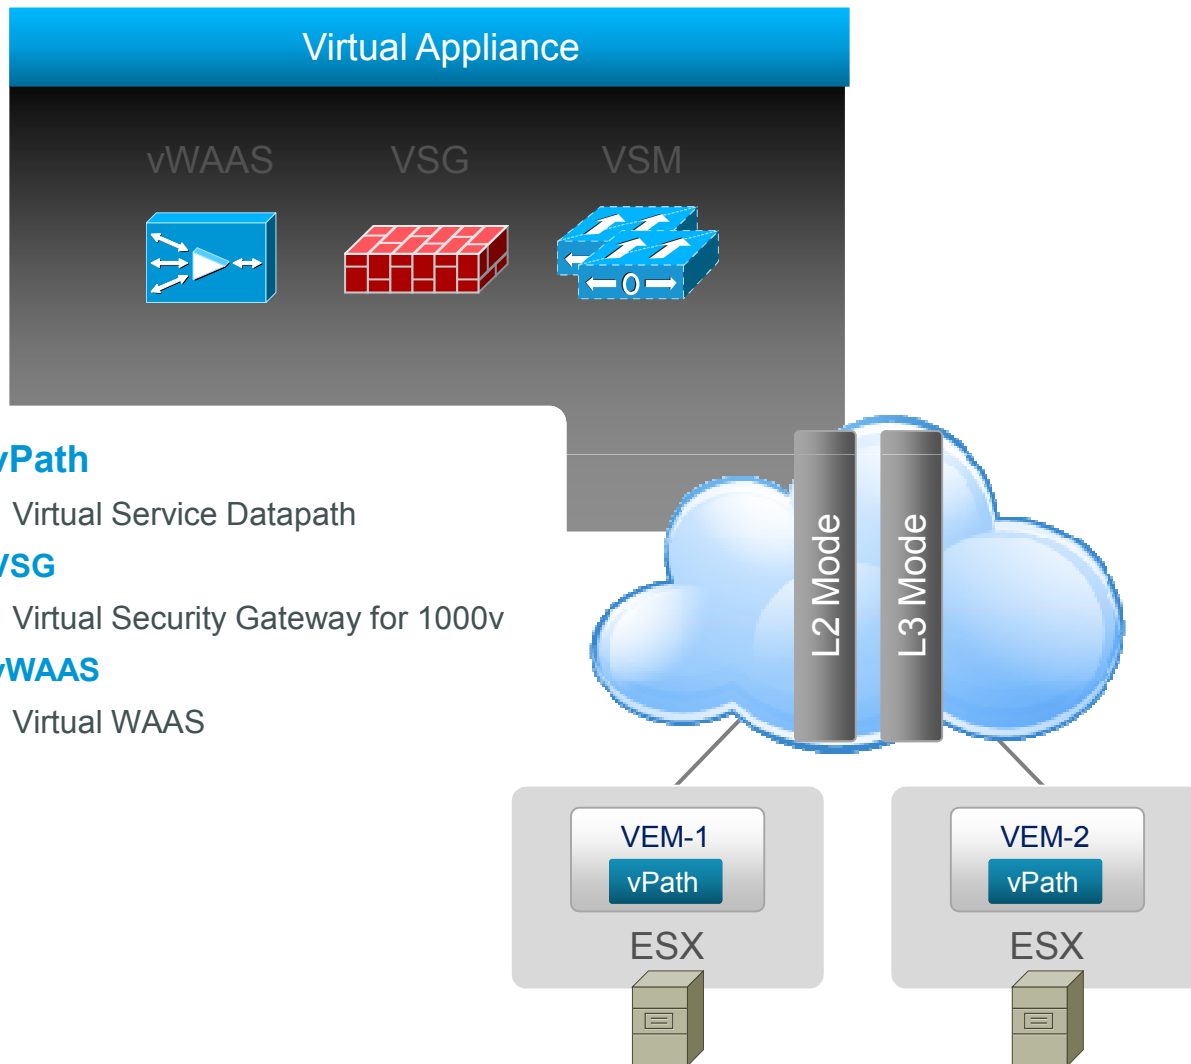


- 200+ vEth ports per VEM
- 64 VEMs per 1000V
- 2K vEths per 1000V
- Multiple 1000Vs can be created per vCenter

VSM: Virtual Supervisor Module  
VEM: Virtual Ethernet Module

# Embedding Intelligence for Virtual Services

## vPath – Virtual Service Datapath



### vPath

- Virtual Service Datapath

### VSG

- Virtual Security Gateway for 1000v

### vWAAS

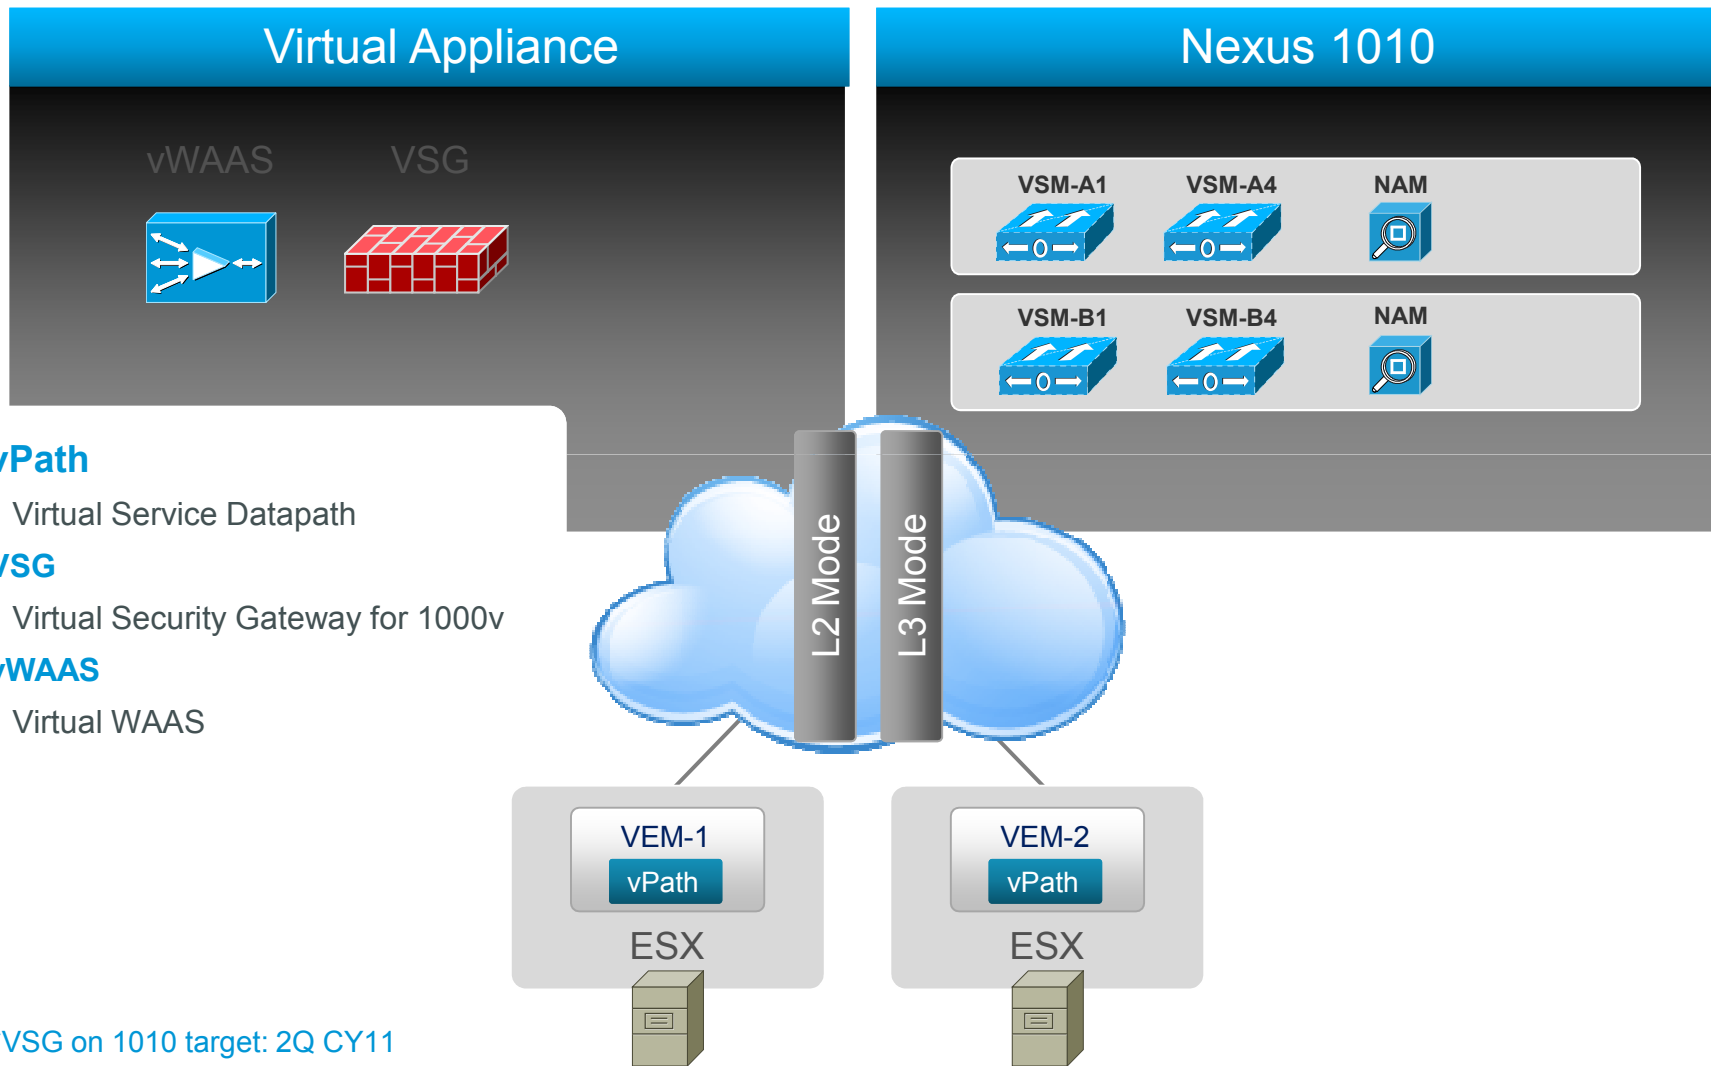
- Virtual WAAS

## vPath

- Traffic Steering
- Fast -Path Offload

• **Nexus 1000V ver 1.4 & above**

# Nexus 1010 – Hosting Platform for Services





# Cisco Nexus 1000V Troubleshooting

Louis Watta  
Customer Assurance Engineer

May 12, 2011



# Agenda

- Useful tools
- Nexus 1000V Releases
- Virtual Supervisor Module (VSM) troubleshooting
- Virtual Ethernet Module (VEM) troubleshooting
- Port-Profiles
- Port Channels and Trunking

# Useful Tools





# Useful Tools

- Vemcmd on ESX host

Can query and set configuration

Can run locally or remote from the 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

- 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

# Nexus 1000V Releases



# Cisco Nexus 1000V Releases

- 4.0(4)SV1(1) – 1.1
- 4.0(4)SV1(2) - 1.2
- 4.0(4)SV1(3) - 1.3
  - Three maintenance releases 1.3a, 1.3b, and 1.3c
- Current Release version 4.2(1)SV1(4)
- All releases work with ESX/ESXi 4.x
- All releases require VMware Enterprise plus license
- Expect VSM/VEM backwards compatibility to work one version back.

# 1.4 New Features

- NXOS 4.2 changes
  - “Feature” command for LACP, PVLAN, Netflow, and port-profile-roles
  - Must enable the feature for above commands to become active
- Change to upgrade procedure
  - Now VEMs first then VSMs
- Network State Tracking (Beaconing)
- LACP Offload
- QOS Weighted Fair Queuing
- Port-profile-roles
  - specify which port-profiles a user has access to
- vPath for VSG and vWAAS

# Virtual Supervisor Module Troubleshooting

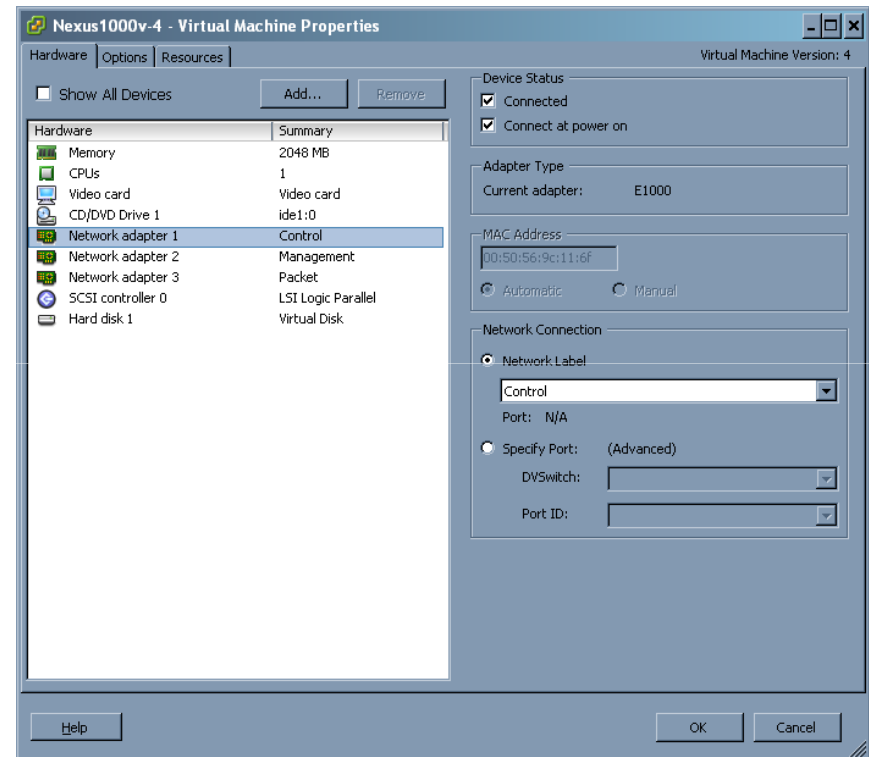


# Virtual Supervisor Module (VSM)

- VSM is a Virtual Machine
  - On ESX
  - 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

# VSM Virtual Machine Requirements

- 3 network interfaces
  - Adapter 1 is always the **Control** interface – Must be on the control VLAN
  - Adapter 2 is always the **Management** interface – Used for VSM connectivity
  - Adapter 3 is always the **Packet** interface – Must be on the packet VLAN
- Need 2GB of memory RAM reserved
  - VSM will do odd things when it is memory starved
- One vCPU



# VSM Required Interfaces


- Management
  - VSM terminal connectivity
  - Connecting to VMware vCenter
  - Backup connectivity for VSM HA
- Control
  - Heartbeat between VSM and VEM
  - Heartbeat and information passing between Active and Standby VSMS
- Packet
  - Passes CDP information
  - Used for ERSPAN sessions
- Multiple VSM installations can use the same VLANs
  - Several VSMS using same Control VLAN
  - SVS domain ID keeps them separate



# VSM Control Modes

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

# 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 VMware APIs:
  - Create N1KV dv-Port-Groups in vCenter
  - Stores DVS data to be passed to ESX hosts which become member of N1KV
  - Get useful information from vCenter (DC, DVS, VM, ...)
  - Maintain periodic connectivity checks

# 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
```

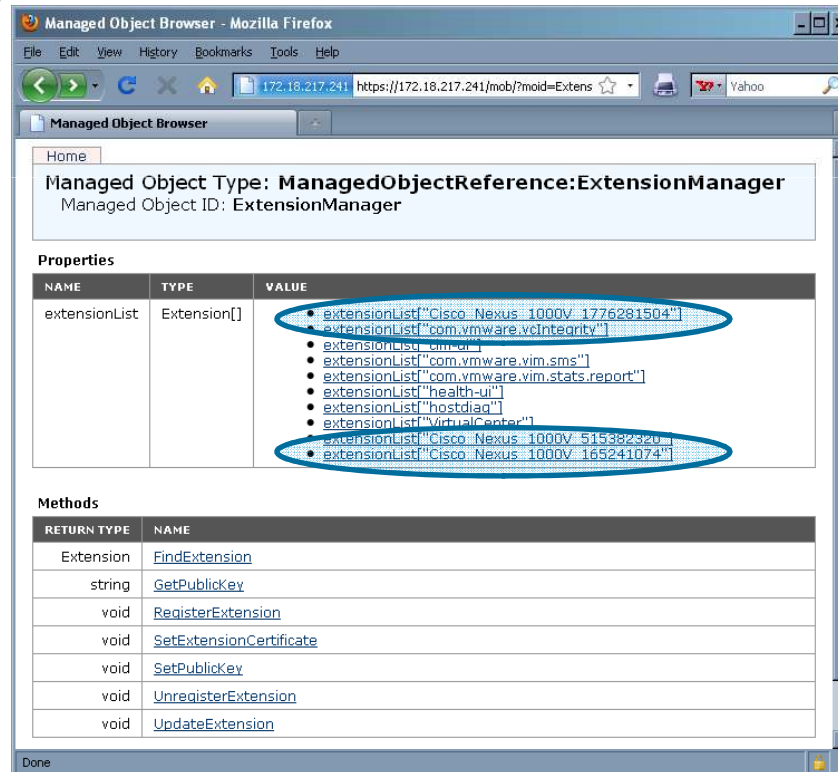
- Check the key on VSM

```
n1000v# show vmware vc extension-key
Extension ID: Cisco_Nexus_1000V_165241074
```

# Check Extension Key on VMware MOB

- Check MOB to see what is registered
- <http://VMware-vCenter-IP/mob>

Content->ExtensionManager



The screenshot shows the Managed Object Browser (MOB) interface in Mozilla Firefox. The browser window title is "Managed Object Browser - Mozilla Firefox". The address bar shows the URL "https://172.18.217.241/mob/?moid=Extens". The page content is titled "Managed Object Type: ManagedObjectReference:ExtensionManager" with the Managed Object ID: "ExtensionManager".

The "Properties" section displays a table with columns "NAME", "TYPE", and "VALUE". The "extensionList" property is of type "Extension[]" and contains a list of extension keys. Two keys are circled in blue:

- extensionList["Cisco\_Nexus\_1000V\_1776281504"]
- extensionList["Cisco\_Nexus\_1000V\_165241074"]

The "Methods" section displays a table with columns "RETURN TYPE" and "NAME". The methods listed are:

RETURN TYPE	NAME
Extension	<a href="#">FindExtension</a>
string	<a href="#">GetPublicKey</a>
void	<a href="#">RegisterExtension</a>
void	<a href="#">SetExtensionCertificate</a>
void	<a href="#">SetPublicKey</a>
void	<a href="#">UnregisterExtension</a>
void	<a href="#">UpdateExtension</a>

# 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
- All communication is https
- VMware accepts on port 80 and tunnels internally to port 8089

# VSM to VMware vCenter Connectivity

- Make sure VSM SVS port matches vCenter http port

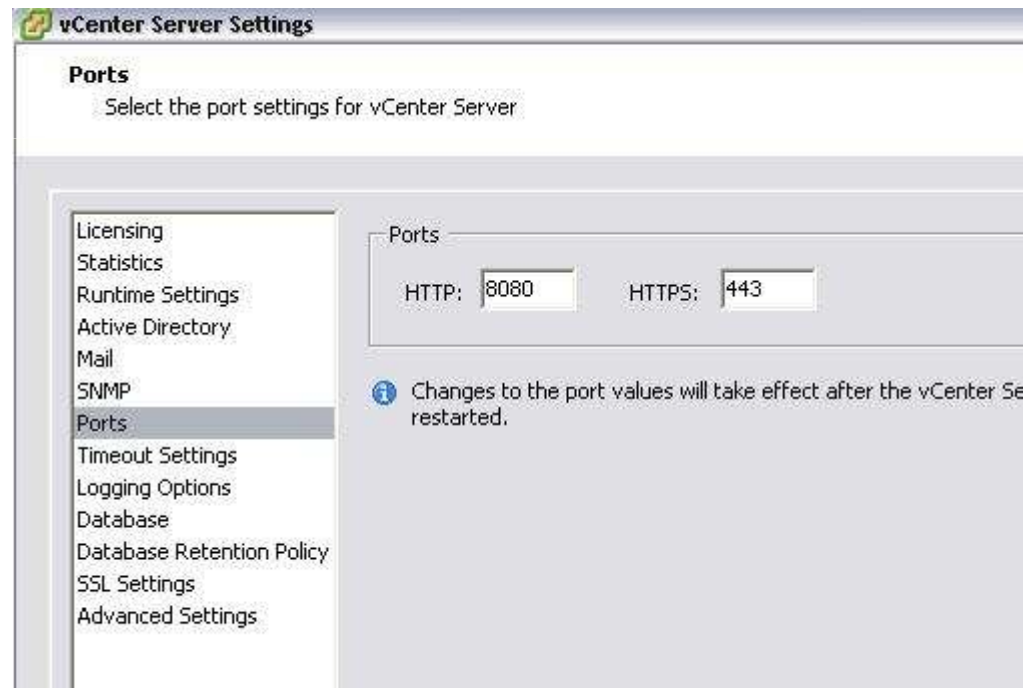
```
n1000v# show svcs connections
connection vcenter:
  ip address: 172.18.217.41
  remote port: 80
  protocol: vmware-vim https
```

- To change port

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

# VSM to VMware vCenter Connectivity

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



# 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



# VSM Best Practices

- L2 control is preferred
- Management, Control, and Packet can use same VLAN
- Do not use VLAN 1 for Control and Packet
- Primary and Standby VSM in same L2 domain!!!
- VSM on VEM is supported
- VSM primary to secondary latency  $\leq 100\text{ms}$
- VSM to VEM latency between 50ms and 100ms
- Backup your config!!!
- VMware snapshots and cloning are not supported

# Virtual Ethernet Module Troubleshooting



# Virtual Ethernet Module (VEM) Troubleshooting Agenda


- VEM installation
- VEM patches
- VEM removal
- VEM does not show up on VSM
- VEM useful commands

# VEM Installation

- Automatically with VMware Update Manager (VUM)  
VUM does all the work user just adds the host to the N1KV  
If you are having problems check the VUM logs  
On vCenter Server in  
C:/Documents and Settings/Application Data/All Users/VMware/VUM/logs  
On ESX host in  
/var/log/vmware/esxupdate.log
- Manually with esxupdate or vihostupdate

# VEM Installation DVS Error

- DVS operation failed error
  - VUM is not installed or configured
  - VUM could not find the right VEM version
  - Manual installation of VEM was not performed



# 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 using correct VEM VIB version for ESX Kernel
- Use the N1KV compatibility matrix to identify the correct VIB to load on the host.

# Matching a VEM Patch to a VIB Version

- Compatibility Matrix – from Cisco

[http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4\\_2\\_1\\_s\\_v\\_1\\_4/compatibility/information/n1000v\\_compatibility.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/compatibility/information/n1000v_compatibility.html)

VMware Software Version	Host Build <sup>1</sup>	VIB Version <sup>2</sup>	VEM Bundle <sup>3</sup>	Minimum Required Version	
				vCenter Server	Update Manager
ESX/ESXi 4.1.0 Update 1 ESX/ESXi410-201101201-SG	348481	cross_cisco-vem_v130-4.2.1.1.4.0.0-2.0.1.vib	VEM410-201101108-BG (Online) VEM410-201101407-BG (Offline)	345043	341095
ESX/ESXi 4.1.0 GA	260247	cross_cisco-vem_v130-4.2.1.1.4.0.0-2.0.1.vib	VEM410-201101108-BG (Online) VEM410-201101407-BG (Offline)	258902	256596
ESX/ESXi 4.0.0 P09 ESX/ESXi400-201103401	360236	cross_cisco-vem_v130-4.2.1.1.4.0.0-1.20.1.vib	VEM400-201101121-BG (Online) VEM400-201101406-BG (Offline)	258672	264019
				208111	282702
ESX/ESXi 4.0.0 Update 2 upgrade-from-esx/esxi4.0-4.0_update02	261974	cross_cisco-vem_v130-4.2.1.1.4.0.0-1.20.1.vib	VEM400-201101121-BG (Online) VEM400-201101406-BG (Offline)	258672 <sup>4</sup>	264019 <sup>4</sup>
				208111	282702
ESX/ESXi 4.0.0 P06 ESX/ESXi400-2011005001	256968	cross_cisco-vem_v130-4.2.1.1.4.0.0-1.13.1.vib	VEM400-201101115-BG (Online) VEM400-201101405-BG (Offline)	208111	282702

# 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]
  - IP address
  - 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 to VC successful.
```

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

# VEM – VSM Connectivity Troubleshooting

- 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  
Is the VMK port-profile set with system VLAN?

# Step 1: VSM MAC

- Need for L2 troubleshooting
- On VSM run show svcs neighbors
- Its the AIPC Interface MAC

```
n1000v# show svcs neighbors
```

```
Active Domain ID: 1254
```

```
AIPC Interface MAC: 0050-5681-0da4
```

```
Inband Interface MAC: 0050-5681-3595
```

## Step 2: VSM – vCenter Connectivity

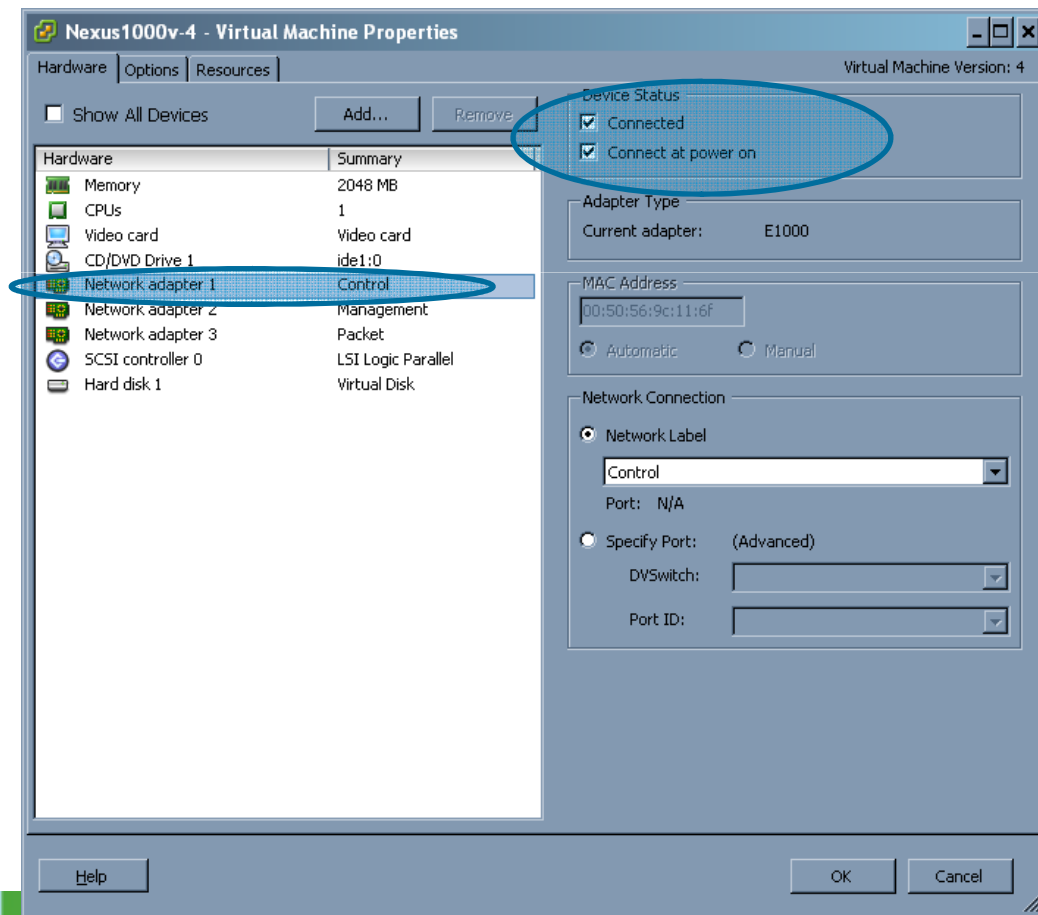
- Verify VSM is connected to vCenter

```
n1000v# show svcs connections

connection MV-testing:
  ip address: 172.18.217.241
  protocol: vmware-vim https
  certificate: default
  datacenter name: Harrison
  DVS uuid: 34 21 01 50 8e 07 cf 03-77 80 29 55 4f 77 4e 89
  config status: Enabled
  operational status: Connected
```

# Step 3: Verify VSM VM Control interface

- 1<sup>st</sup> interface listed is Control Interface
- Is Interface connected?



# Step 4: Verify Uplink Port-Profile

- The first ESX interface added to the N1KV must have Control VLAN
- Verify uplink port-profile has Control VLAN defined and system VLAN

```
n1000v# show port-profile name uplink
port-profile uplink
  description:
  status: enabled
  capability uplink: yes
  capability l3control: no
  system vlans:2,10,150-152
  port-group: uplink
  max-ports: -
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan all
```

# Step 5: Verify VEM Sees Control VLAN

- Verify VEM sees control VLAN with commands
  - vemcmd show card
  - vemcmd show port-old
  - vemcmd show trunk

# Vemcmd show card

- Control, packet vlans and domain-ID match with VSM

```
[root@cae-cali-host6 ~]# vemcmd show card
Card UUID type 2: 8c1d5178-8c02-11d9-0000-00000000000d
Card name: cae-cali-host6
Switch name: n1000v-MV
Switch alias: DvsPortset-0
Switch uuid: 72 f7 01 50 b2 01 7b 8b-55 68 cf df 10 5a db 55
Card domain: 234
Card slot: 4
VEM Control (AIPC) MAC: 00:02:3d:10:ea:03
VEM Packet (Inband) MAC: 00:02:3d:20:ea:03
VEM Control Agent (DPA) MAC: 00:02:3d:40:ea:03
VEM SPAN MAC: 00:02:3d:30:ea:03
Management IP address: 172.18.217.177
Max physical ports: 32
Max virtual ports: 216
Card control VLAN: 150
Card packet VLAN: 150
Processors: 16
```

MAC the VSM  
should learn for  
VEM



# Vemcmd show port-old

- Ports with LTLs 8(I20), 9(I21), 10(I22) are UP and CBL states are 4.
- ESX Physical ports are UP and CBL states 4.

```
root@cae-cali-host6 ~]# vemcmd show port-old
```

LTL	IfIndex	Vlan	Bndl	SG_ID	Pinned_SGID	Type	Admin	State	CBL	Mode	Name
8	0	3969	0	2	2	VIR	UP	UP	4	Access	l20
9	0	3969	0	2	2	VIR	UP	UP	4	Access	l21
10	0	150	0	2	0	VT	UP	UP	4	Access	l22
11	0	3968	0	2	2	VIRT	UP	UP	4	Access	l23
12	0	151	0	2	1	VIRT	UP	UP	4	Access	l24
13	0	1	0	2	2	VIRT	UP	UP	0	Access	l25
14	0	3967	0	2	2	VIRT	UP	UP	4	Access	l26
15	1a030000	1 T	303	0	2	ES	UP	UP	4	Trunk	vmnic0
16	1a030100	1 T	303	1	2	ES	UP	UP	4	Trunk	vmnic1
47	1b030000	10	0	2	0	VIRT	UP	UP	4	Access	vmk0
48	1b030010	10	0	2	0	VIRT	UP	UP	4	Access	vmk1
49	1b030020	2	0	2	1	VIRT	UP	UP	4	Access	vswif0
50	1b030030	152	0	2	1	VIRT	UP	UP	4	Access	linux1.eth0
303	16000002	1 T	0	2	2	VIRT	UP	UP	4	Trunk	

# Vemcmd show trunk

- Control and packet are CBL states 4 on the physical ports.

```
[root@cae-cali-host6 ~]# vemcmd show trunk
Trunk port 15 native_vlan 1 CBL 4
vlan(1) cbl 4, vlan(2) cbl 4, vlan(10) cbl 4, vlan(150) cbl 4, vlan(151)
cbl 4, vlan(152) cbl 4, vlan(153) cbl 4, vlan(154) cbl 4, vlan(155) cbl 4,

Trunk port 16 native_vlan 1 CBL 4
vlan(1) cbl 4, vlan(2) cbl 4, vlan(10) cbl 4, vlan(150) cbl 4, vlan(151)
cbl 4, vlan(152) cbl 4, vlan(153) cbl 4, vlan(154) cbl 4, vlan(155) cbl 4,
```

## ■ Vemcmd show port vlans

```
[root@cae-esx-184 ~]# vemcmd show port vlans
```

LTL	VSM Port	Mode	Native VLAN	VLAN State	Allowed Vlans
17	Eth7/1	T	1	FWD	2,10,150-155
18	Eth7/2	T	1	FWD	2,10,150-155
49	Veth8	A	10	FWD	10
50	Veth9	A	2	FWD	2
305	Po2	T	1	FWD	2,10,150-155

# Step 6: VEM and VSM See Each Other's MAC

- Is the VEM learning the MAC of the VSM?
- On VEM “**vemcmd show l2 <control-vlan>**” do you see the mac of the VSM?

```
[root@cae-esx-180 ~]# vemcmd show l2 150
Bridge domain 150 brtmax 1024, brtcnt 3, timeout 300
Flags: P - PVLAN S - Secure
      Type          MAC Address    LTL    timeout    Flags
PVLAN
Dynamic 00:50:56:81:0d:a4 304      1
Static 00:02:3d:40:ea:03 10       0
```

# VEM and VSM See Each Other's MAC

- Is the VSM learning the MAC of the VEM?

```
n1000v-AV# show mac address-table vlan 150
```

VLAN	MAC Address	Type	Age	Port	Mod
150	<b>0002.3d40.ea03</b>	static	0	N1KV Internal Port	6
150	0002.3d47.db05	static	0	N1KV Internal Port	6
150	0017.a4a8.340a	dynamic	16	Eth6/2	6

# Step 7: Physical Switch Mac Table

- Check the physical switch MAC address table
- Are the MACs of the VEM and VSM getting learned by the physical switches in the right VLANs?

```
cae-cat6k-1#show mac-address-table vlan 150
```

```
Legend: * - primary entry
```

```
age - seconds since last seen
```

```
n/a - not available
```

vlan	mac address	type	learn	age	ports
* 150	0050.5677.7770	dynamic	Yes	360	Gi3/48
* 150	0002.3d40.ea03	dynamic	Yes	330	Gi3/48
* 150	3333.0000.0016	static	Yes	-	Switch,Stby-Switch
* 150	0050.5681.0da4	dynamic	Yes	0	Gi4/19

# Step 8: VEM – VSM Heartbeat

- One Heartbeat per second per VEM from VSM
- Timeout for VEM from VSM is 6 seconds of missed heartbeats
- After 6 seconds VSM will drop VEM

# Check VSM Counters for Heartbeat

- On VSM

```
N1K-VSM# show module vem counters
```

```
-----  
Mod   InNR  OutMI  InMI  OutHBeats  InHBeats  InAipcMsgs  OutTO  OutTOC  InsCnt  RemCnt  
-----  
    3     1     1     1     82243     246554     24752         0         0         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

**InAipcMsgs** - Number of AIPC msgs received from this VEM

OutTO - Number of aipc transmit timeout errors recorded for this VEM

# View Heartbeat Messages on VEM

- On the ESX host
- Use vempkt command to view Heartbeat messages
  - vempkt capture ingress vlan 150
  - vempkt display brief all
  - vempkt display detail all
  - vempkt stop
- Look for heartbeat messages from VSM – from detail all

```
***** Entry 484 *****
-----Packet Entry Information-----
      Timestamp   : Aug 10 21:42:18.388822
-----SF Packet Information-----
      Capture Stage : Ingress
      Source LTL    : 303
      Vlan          : 150
-----Packet L2 Header Information-----
      Source MAC Address : 00:50:56:9c:11:6f
      Destination MAC Address : 00:02:3d:40:ea:03
      Length          : 98
```



# VEM Best Practices

- Control network should have low latency (50-100ms) and available bandwidth
- Match VEM version to VSM
- Using one uplink to carry all traffic is ok
  - Segregate VM traffic with port-profiles and MAC pinning
- Upstream switch ports configured identically
- 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-Profiles



# Uplink(eth) Port-Profile Troubleshooting

- Do not add multiple pnic's from same ESX host to same uplink port-profile if no port-channeling is configured
  - While it may work you will end up with the ESX host receiving duplicate packets and require extra processing from CPU to deal with this improper configuration
- Do not configure multiple uplink port-profiles to an ESX host carrying the same vlan
  - Uplink1 and uplink2 to same ESX host both carrying vlan 100
  - vPC-Mac pinning
- If you want NIC teaming use one of the approved port-channel mechanisms

# VM(veth) Port-Profile Troubleshooting

- VM port-profiles common issues
  - Port-profiles for Service Console and VMK should have system vlan set
- Be careful modifying veths directly
  - If at all possible 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

# System VLAN

- System VLANs enable interface connectivity before the VEM 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'
  - 32 in from 1.3b onward
  - 16 in all other versions
- 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

# System VLAN Guidelines

- The system VLAN list must be a subset of the allowed VLAN list on trunk ports
- There must be only one system VLAN on an access port (the access VLAN)
- The 'no system vlan' command can be given only when no interface is using the profile.
- Once a system profile is in use by at least one interface, you can only add to the list of system VLANs, but not delete any VLANs from the list.
- Required System VLANs  
Control/Packet, IP Storage, Service Console, VMKernel, Management Networks

# System VLAN Example

- Migrate VMware Service Console to VEM
- SC interface uses VLAN 2
- Uplink port-profile must define VLAN 2 as system

```
n1000v# show run port-profile uplink-pinning
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
```

- Service Console Port-profile must also define system vlan

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

# Jumbo Frames Support

- MTU setting for “eth” type port-profile – 1.4
  - Simply use “mtu size” in port-profile and nothing else
  - Add system vlan directive to Port-profile if needed
- System MTU setting for “eth” type port-profile – 1.3
  - “system mtu 9000” use with “system vlan” directive
  - No need to set MTU directly on eth interfaces or port-channels
  - Will consume a system port-profile
- Modify interfaces directly - 1.3
  - If not you do no need “system vlan” or do not want to consume “system port-profile”
  - Modify eth or port-channel interfaces directly
- “System jumbo mtu” global setting – all versions
  - Sets the system wide jumbo mtu size
  - Generally do not need to change



# Port-profiles max-ports

- We default to 32 max-ports per port-profile
  - The default number of veths that can be assigned to a port-profile
- Max-ports number counts toward the maximum number of DVS ports that VMware can support
  - Even if veths are not yet assigned the ports are pre-provisioned
- 8192 ports per DVS in ESX 4.0U1/U2
- 20000 ports per DVS in ESX 4.1
- Note if you upgrade from 4.0 to 4.1 the number stays at 8192
  - VMware has a workaround to change the setting to 20000
- Example with vCenter 4.0U1
  - Create 10 port-profiles with max-ports set to 1024
  - Limit is 8192 ports per DVS
  - Will get error from vCenter about Max Total Ports Exceeded

# Port Channels



# Port Channels

- 3 load balancing modes

- LACP Port-channels – require support on physical switch

- vPC – HM MAC Pinning works with any switch

- Preferred Channeling method over vPC-HM CDP/Manual

- Allows for pinning of veths (VM) to specific links.

- vPC – Host Mode allows balancing to multiple physical switches

- NIC association is either Manual or via CDP

- If more than one connection per physical switch then port-channel is required

- Channeling mode must be configured on uplink port-profile

- Not recommended to create a channel by modifying individual Ethernet interfaces

# LACP Port Channels

- Use when single upstream or clustered (vPC,VSS) switch
- Use “**channel-group auto mode active**” on N1KV
- Use “**channel-group # mode active**” on upstream switch
- Switchports must be configured with
  - spanning-tree portfast trunk
  - spanning-tree bpdudfilter enable
- LACP PC will cycle between w and h states when uplink carries control and packet
  - One uplink is held back from LACP negotiation to guarantee control and packet have a valid link path while other links are aggregated
- VSM negotiates LACP for every VEM.
  - If VSM is down VEM cannot negotiate LACP and link will not come up

# LACP Offload

- New in 1.4
- Turned on with  
n1000v-AV(config)# lacp offload  
Requires a reboot of VSM to become active
- Feature moves LACP negotiation from VSM to VEM module
- Allows a VEM to negotiate LACP if VSM is not up
- Improves LACP stability

# Port Channels – vPC HM

- vPC-HM uses Service Group (SG)
  - Service Group is a collection of Ethernet interfaces from ESX host
  - One Service Group per physical path
- CDP is used to determine SG membership
  - Can be a 60 second delay while VSM determines NIC membership because of CDP
- Can configure SG membership manually for switches without CDP support
- Multiple links per physical path must be configured as a port-channel upstream

# Port Channels – vPC HM MAC Pinning

- Each Eth interface added is a unique Service Group
- Use “pinning id” command under vethernet port-profile
- veths will failover to another interface if pinned Ethernet interface fails
- If port-profile is not “pinned” vethernet interfaces are assigned Round Robin to an SG

# Port Channels – How to Tell Pinning

- Use “vemcmd show port” on ESX host

```
[root@cae-esx-184 ~]# vemcmd show port
```

LTL	VSM Port	Admin	Link	State	PC-LTL	SGID	Vem Port
17	Eth7/1	UP	UP	FWD	305	0	vmnic0
18	Eth7/2	UP	UP	FWD	305	1	vmnic1
49	Veth1	UP	UP	FWD	0	0	VSM-1.3a-2.eth0
50	Veth2	UP	UP	FWD	0	0	VSM-1.3a-2.eth1
51	Veth3	UP	UP	FWD	0	1	VSM-1.3a-2.eth2
52	Veth5	UP	UP	FWD	0	1	vswif0
53	Veth4	UP	UP	FWD	0	1	vmk0
305	Po1	UP	UP	FWD	0		



# 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  
Remember requires a VSM reboot to turn on
- 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 bpduguard  
    cat65k(config)# spanning-tree portfast bpdufilter  
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 Filter/Guard when Global is not possible

# Summary



# Summary

- Cisco Nexus 1000V troubleshooting document  
[http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4\\_0\\_4\\_s\\_v\\_1\\_2/troubleshooting/configuration/guide/trouble\\_n1000v.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_2/troubleshooting/configuration/guide/trouble_n1000v.html)
- Cisco MyCommunity Nexus 1000V Space  
<https://www.myciscocommunity.com/community/products/nexus1000v>
- Cisco Nexus 1000V FAQ  
<https://www.myciscocommunity.com/docs/DOC-14464>
- VMware ESX Networking Community  
<http://communities.vmware.com/community/vmtn/vsphere/networking?view=discussions&start=0>



Sign up at: <http://tinyurl.com/1000v-webinar>

Date	Business Sessions
22-Mar	Nexus 1000V Family Overview and Update
5-Apr	Virtual Network Services (vPath, vWAAS, NAM)
19-Apr	Virtual Security Gateway Introduction
3-May	Journey to the Cloud w/ N1KV: vCloud Director & Long Distance vMotion
17-May	Secure VDI with Nexus1000V & VSG

Date	Technical Sessions
29-Mar	Nexus 1000V New Features and Installation Overview
12-Apr	Nexus1010 Installation & Upgrade
26-Apr	Virtual Security Gateway Installation and Basic Configuration Overview
10-May	Nexus 1000V Advanced Configuration
24-May	Nexus 1000V Troubleshooting

### Web Sites

[www.cisco.com/go/1000v](http://www.cisco.com/go/1000v)

[www.cisco.com/go/nexus1010](http://www.cisco.com/go/nexus1010)

[www.cisco.com/go/vsg](http://www.cisco.com/go/vsg)

[www.cisco.com/go/vnmc](http://www.cisco.com/go/vnmc)

[www.cisco.com/go/1000vcommunity](http://www.cisco.com/go/1000vcommunity)  
(Preso and Q&A posted here)

Thank you.

