




Deployment of VN-link with the Nexus1000v

BRKVIR-2006




Related Material

- **BRKVIR-2008** UCS and Nexus1000V
Virtualization for Cloud DC Services
- **BRKVIR-2011** Deploying Services in a Virtualized
Environment
- **LABDCT-1901** Introduction to Nexus 1000V
Hands-on Lab
- VMware 10 Gigabit White Paper:
http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c07-607716.html

Agenda

- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- System Architecture Overview
- Switching Overview
- Policy Management
- Connectivity and Design
- Q and A


Transparency in the Eye of the Beholder



With virtualization, VMs have a transparent view of their hardware resources...

- Hypervisor presents “dedicated” CPU and Memory resources to each VM
- Allows multiple servers running different Operating Systems to run simultaneously on the same hardware


Transparency in the Eye of the Beholder



...but, its difficult to correlate from a network point of view

- VM's have no "dedicated" network resources
- physical ports, adapters, links, wires are all shared
- **one-to-one connection between switchport and server is not there**

VN-Link Brings VM Level Granularity



Problems:

- VMotion may move VMs across physical ports—policy must follow
- Impossible to view or apply policy to locally switched traffic
- Cannot correlate traffic on physical links—from multiple VMs

VN-Link:

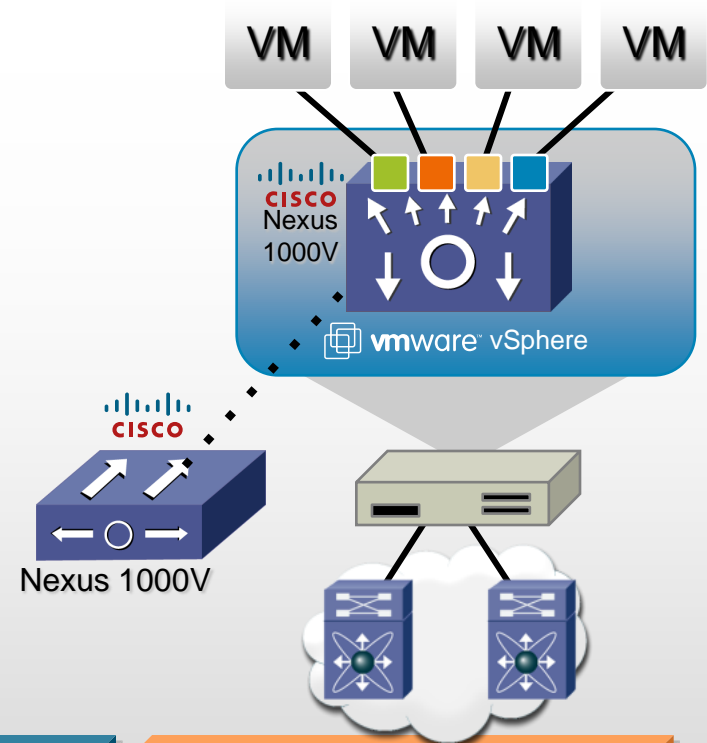
- Extends network to the VM
- Consistent services
- Coordinated, coherent management

VN-Link With the Cisco Nexus 1000V

Cisco Nexus™ 1000V Software Based

- Industry's first 3rd-party vNetwork Distributed Switch for VMware vSphere
- Built on Cisco NX-OS
- Compatible with all switching platforms
- Maintain vCenter provisioning model unmodified for server administration; allow network administration of virtual network via familiar Cisco NX-OS CLI

BEST OF
vmworld® 2008



Policy-Based
VM Connectivity



Mobility of Network and
Security Properties

Non-Disruptive
Operational Model

Agenda

- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- System Architecture Overview
- Switching Overview
- Policy Management
- Connectivity and Design

Cisco Nexus 1000V Components




Virtual Ethernet Module (VEM)


- Replaces VMware's virtual switch
- Data Plane for the N1KV
- Enables advanced switching capability on the hypervisor
- Provides each VM with dedicated "switch ports"

Virtual Supervisor Module (VSM)



- CLI interface into the Nexus 1000V
- Control Plane for the N1KV
- Leverages NX-OS
- Controls multiple VEMs as a single network device




Cisco VEM



Cisco VEM



Cisco VEM



Cisco Nexus 1000V

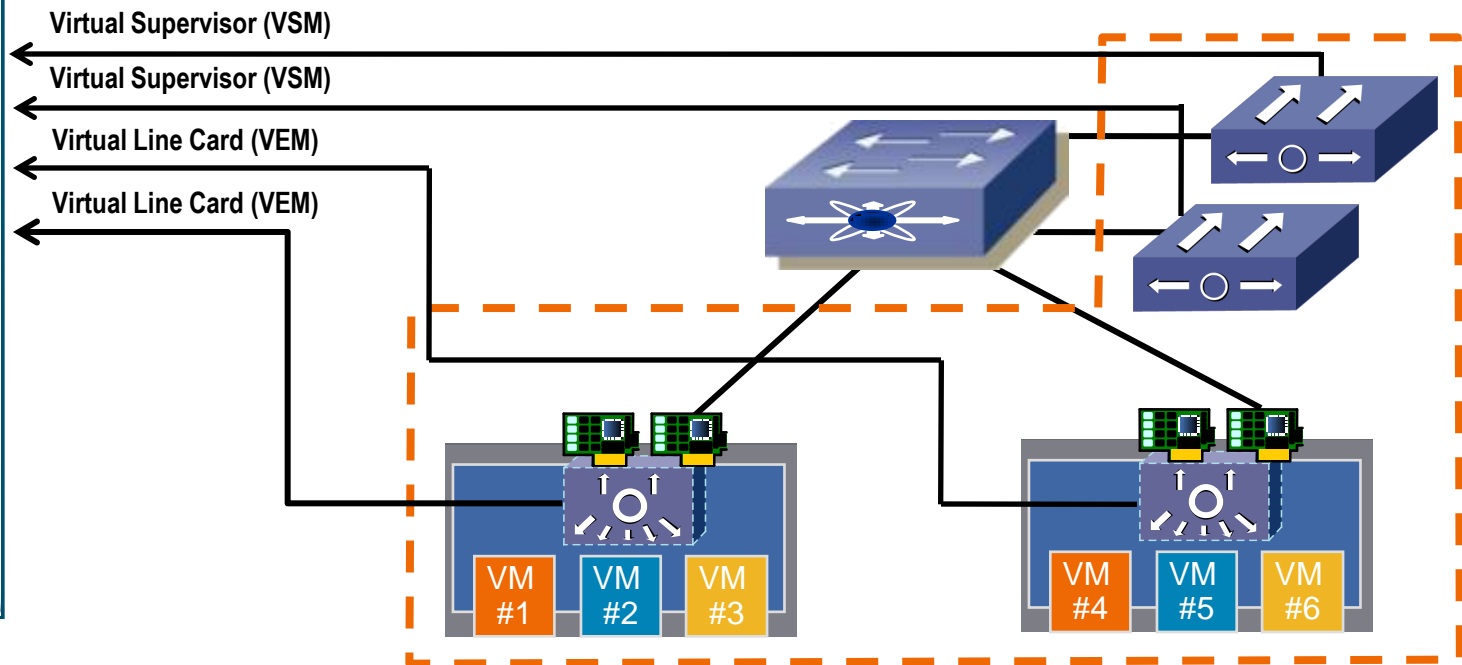
Single Chassis Management

- All components in a single vSphere cluster operate logically as a single virtualised access switch



```
pod5-vsm# show module
```

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



Cisco Nexus 1000V

Single Chassis Management


```

Upstream-Switch#show cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID           Local Intrfce      Holdtme  Capability  Platform  Port ID
N1KV-Rack10         Eth 1/8            136      S           Nexus 1000V Eth2/2
N1KV-Rack10         Eth 2/10           136      S           Nexus 1000V Eth3/2
  
```

VSMs and VEMs appear as one switch from a management and control plane perspective

- Protocols such as CDP, LACP, SNMP, IGMP operate as a single switch
- These control protocols are run on the VSM and carried over the “packet VLAN” to the VEMs



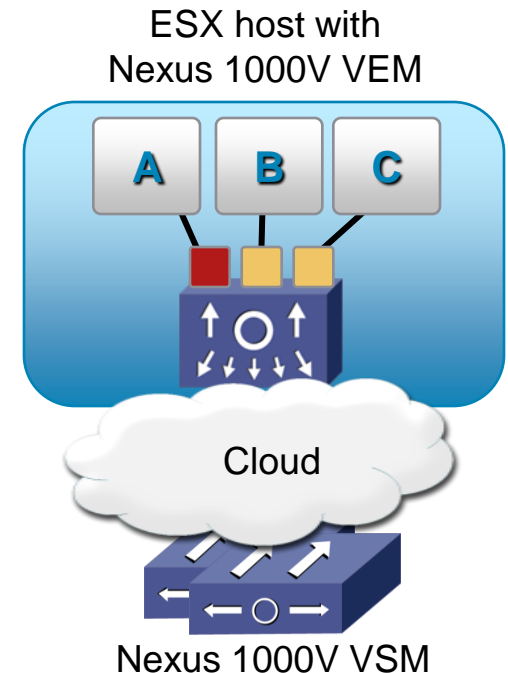
Agenda

- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- **System Architecture Overview**
- Switching Overview
- Policy Management
- Connectivity and Design

Cisco Nexus 1000V Communication

Extending the “Backplane”

- Because the VSM and VEM are not physically connected, the VSM (supervisor) must program the VEM (linecards) over a network
- The VSM to VEM communication uses the same backplane protocol found in the Nexus 7000 and MDS platforms called AIPC
- There are two ways to extend communication between VSM and VEM:
 - Over Layer 2 Cloud using Control and Packet VLANs
 - Over Layer 3 Cloud using Layer 3 Control Capability



Cisco Nexus 1000V

The “system” VLAN

- **“system”** VLANs are for the critical ports which need to be enabled and forwarding before communication with the VSM is established
- **“system”** VLANs should be used for Control, Packet, Storage and Management VLANs
- **“system”** VLANs must be declared in the virtual Ethernet port-profile as well as the uplink port-profile

```
Nexus1000(config)# port-profile VSM_Cont_Pack_Mgmt
Nexus1000(config-port-prof)# switchport mode access
Nexus1000(config-port-prof)# switchport access vlan 10
Nexus1000(config-port-prof)# system vlan 10
```

http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_2/getting_started/guide/n1000v_gsg_5vsm_behind_vem.html

Nexus 1000V VSM

Virtual Adapter Interfaces

The VSM has 3 Interfaces:

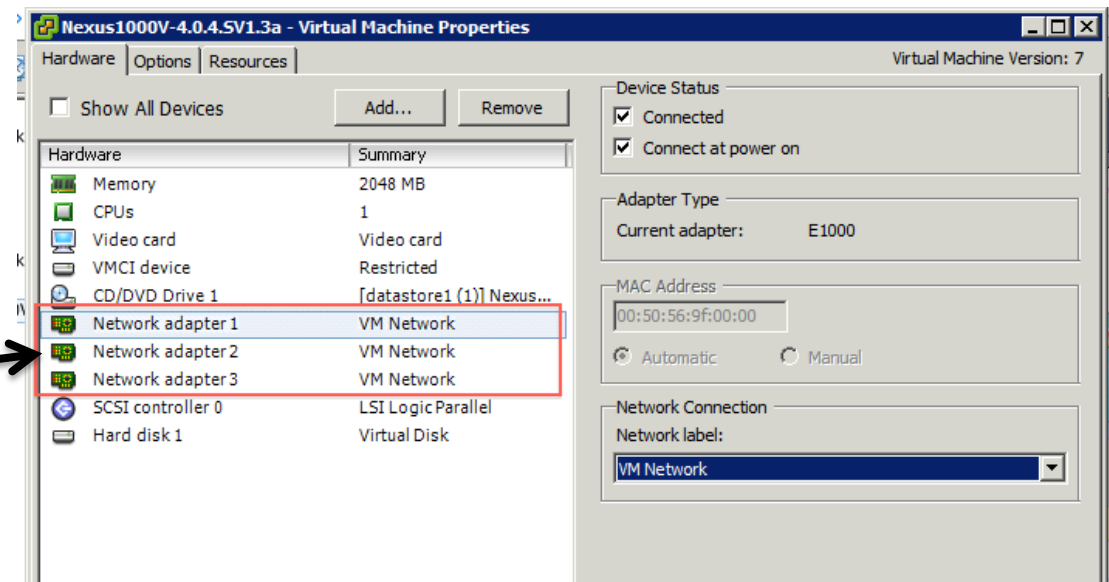
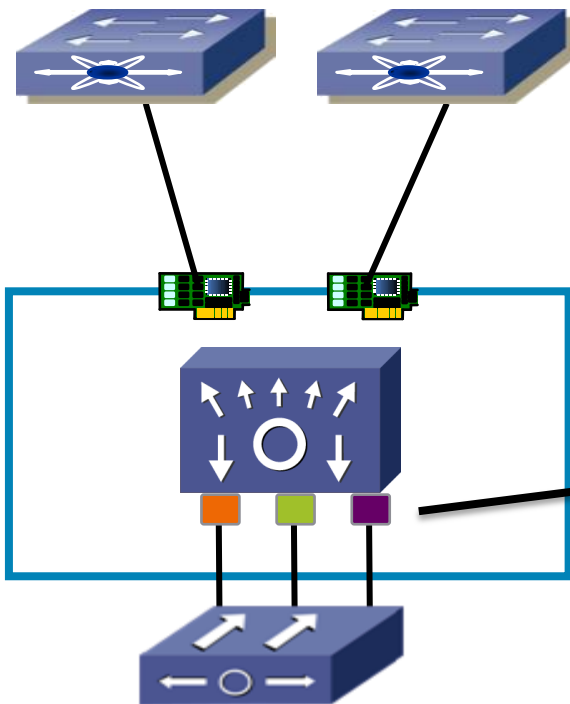
“Network adapter 1” == Control

“Network adapter 2” == Management

“Network adapter 3” == Packet

Create different port-profiles for each interface if separate VLANs are required –or– a single port-profile if all control/management/packet traffic is sharing the same VLAN

```
Nexus1000(config)# port-profile VSM_Cont_Pack_Mgmt
Nexus1000(config-port-prof)# switchport mode access
Nexus1000(config-port-prof)# switchport access vlan 10
Nexus1000(config-port-prof)# system vlan 10
```



VSM and VEM Communication

Layer 2 Connectivity

- Two distinct virtual interfaces are used to communicate between the VSM and VEM

Control: uses “Control” VLAN

Extends AIPC between “SUP” and “linecard”

Carries low level messages to ensure proper configuration of the VEM

Maintains a 1sec heartbeat with the VSM to the VEM (timeout 6 seconds)

Maintains synchronization between primary and secondary VSMs

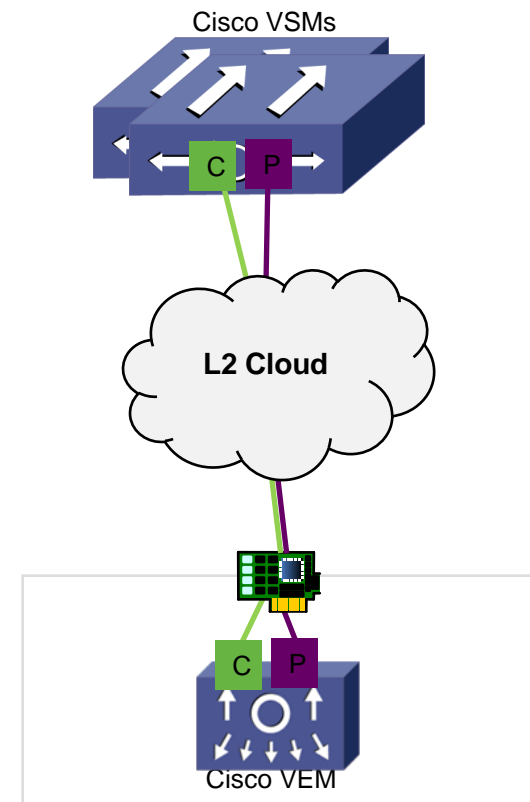
Packet: uses “Packet” VLAN

Carries any network packets from the VEM to the VSM such as CDP or IGMP control

- Control and Packet VLANs can share the same VLAN or be separate VLANs based on customer requirements


Separate offers traffic segregation

Sharing offers simplicity



VSM and VEM Communication

Layer 3 Connectivity

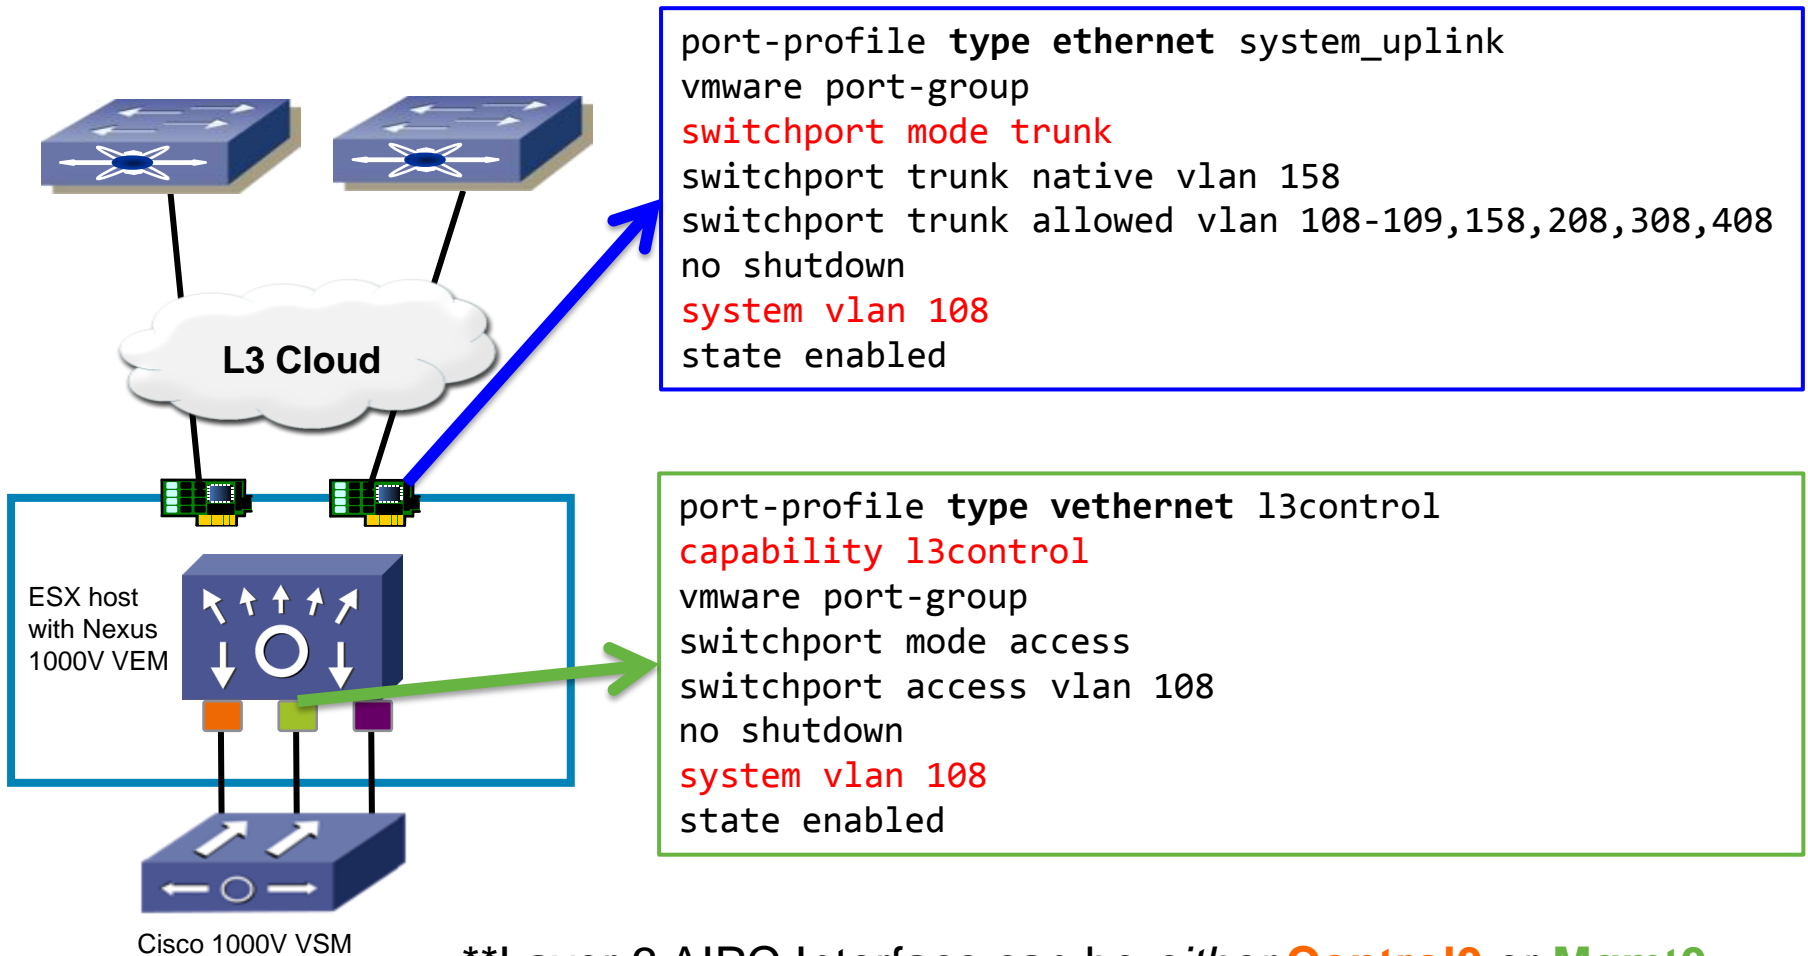


- Introduced in the 4.0(4)SV1(2) release
- User specifies an IP address belonging to a separate network for VSM to VEM communication
 - Requires enough IP addresses to span all participating ESX hosts
 - Requires IP connectivity between the ESX hosts and the VSM
- L3 AIPC functionality is accomplished by encapsulating the Control and Packet interface data in an encrypted Ethernet over IP tunnel
- Layer 2 adjacency is still required between VSMs deployed in the HA pair where both the Active and Standby VSM share the same Control and Packet VLANs

****NOTE:** that the L3 communication is encapsulated as a **UDP packet**, which may be important if there is a firewall device between the VSM and VEM

VSM and VEM Communication


Layer 3 Configuration



Layer 3 AIPC Interface can be *either* **Control0 or **Mgmt0**
Config shown below uses Mgmt0 for L3 interace

Cisco Nexus 1000V Component

vCenter Communication



- Communication is sent securely between the VSM and VC using the VMware VIM API
- Connection is configured manually on the VSM or establish automatically through the VSM GUI installer application
- Requires installation of vCenter plug-in (downloaded from VSM)
- Once established the Nexus 1000V is created in vCenter

```
pod5-vsm# show svcs connections

connection VC:
  hostname: phx2-dc-pod5-vc
  ip address: 10.95.5.158
  protocol: vmware-vim https
  certificate: default
  datacenter name: Phx2-Pod5
  DVS uuid: df 11 38 50 0a 95 83 4e-95 69 d6 a7 f4 76 4a 7f
  config status: Enabled
  operational status: Connected
```


Cisco Nexus 1000V

Domain ID

Each instance of the Nexus 1000V DVS must contain a **unique** Domain ID used to distinguish between multiple 1000Vs in a DC

Each packet between VSM and VEM is tagged with the appropriate Domain ID to ensure that VEMs only respond to commands from their managing VSMs

Configured during set up and can range from 1–4095

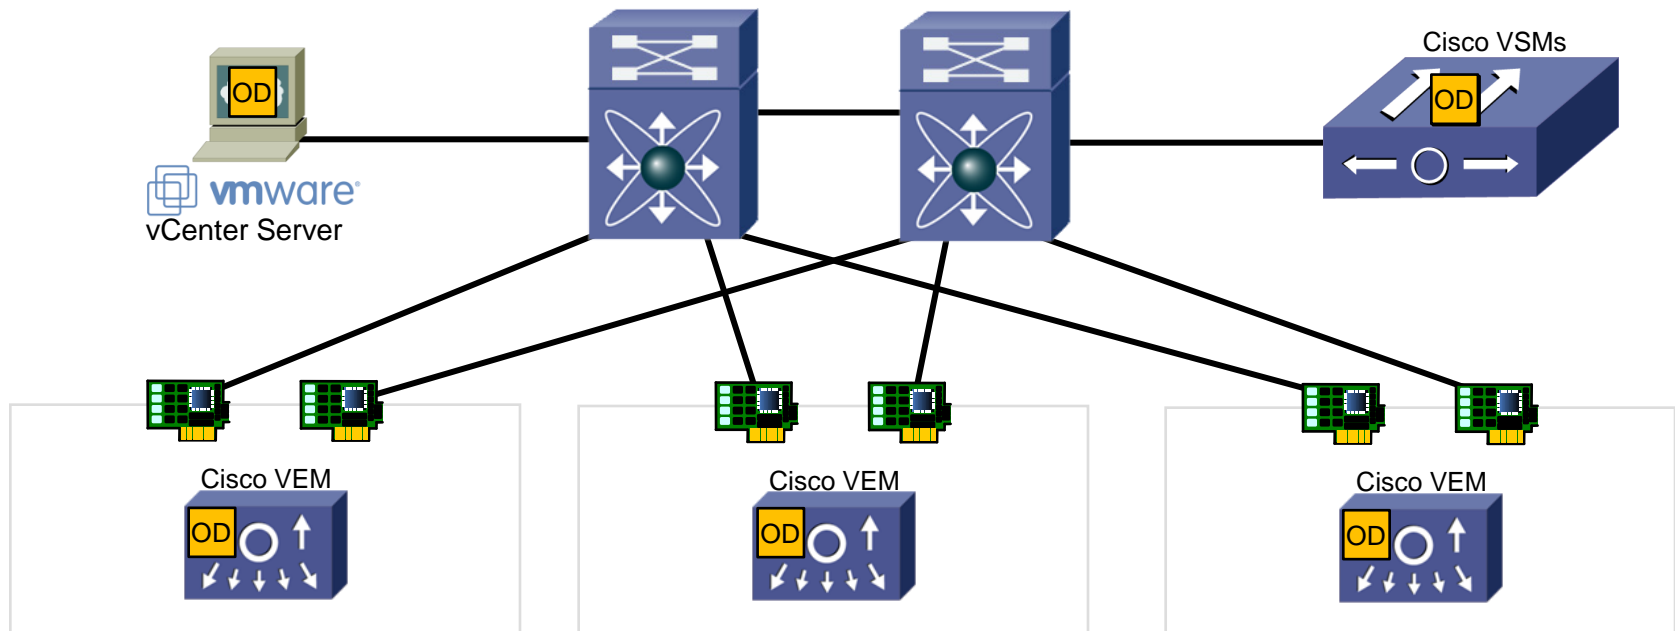


Cisco Nexus 1000V

Opaque Data

Each Nexus 1000V requires global setting on the VSMs and VEMs called ***Opaque Data***

- Contains such data as control/packet VLAN, Domain ID, System Port Profiles
- VSM pushes the opaque data to vCenter Server
- vCenter Server pushes the opaque data to each VEM as they are added



Cisco Nexus 1000V

VSM deployment options




A virtual machine in the Server Admin Domain



A physical appliance in the Network Admin Domain

Nexus 1010



Nexus 1010—Physical Appliance

- Houses up to four VSMs
- Deployed in pairs for redundancy
- Built on Cisco UCS
- Is the future platform for all virtual services (vNAM, VSG, vWAAS, etc)

VSM—Virtual Appliance



- ESX Virtual Appliance
- Supports 64 VEMs (64 hosts)
- Installable via ISO or OVA file




Cisco VEM



Cisco VEM



Cisco VEM



Cisco Nexus 1000V

VSM Deployment and High Availability Options

- As of 4.0(4)SV1(2) the VSM can live on a VEM that it is managing


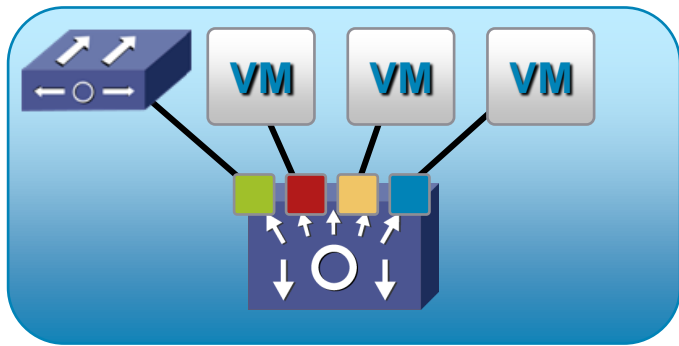
With older versions, it is recommended to deploy the VSM on a separate vSwitch

- VSM's should always be deployed in an HA Pair

VSM-VMs should utilize VMware's anti-affinity rule to ensure that both VSMs will not reside on the same ESX host

Nexus 1010 VSMs must be deployed on two Nexus 1010 boxes

VSM HA pair must reside in the same physical DC location



Agenda


- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- System Architecture Overview
- **Switching Overview**
- Policy Management
- Connectivity and Design

Cisco Nexus 1000V

Distributed Data Plane

Each Virtual Ethernet Module forwards packets independently

- No address learning/synchronization across VEMs
- No concept of Crossbar/Fabric in the data plane
Virtual Supervisor Module is **NOT** in the data path
- No concept of forwarding from an ingress linecard to an egress linecard (in one server and out another server)
- No EtherChannel across VEMs



Cisco Nexus 1000V

Switch Interface Types

- Ethernet Port (eth)

Corresponds to the **physical NIC** interfaces leaving the server: VMware's vmnics)

Numbering based on VEM Module

Up to 32 Physical NICs per ESX host

- Port Channel (po)

Aggregation of physical Ethernet ports

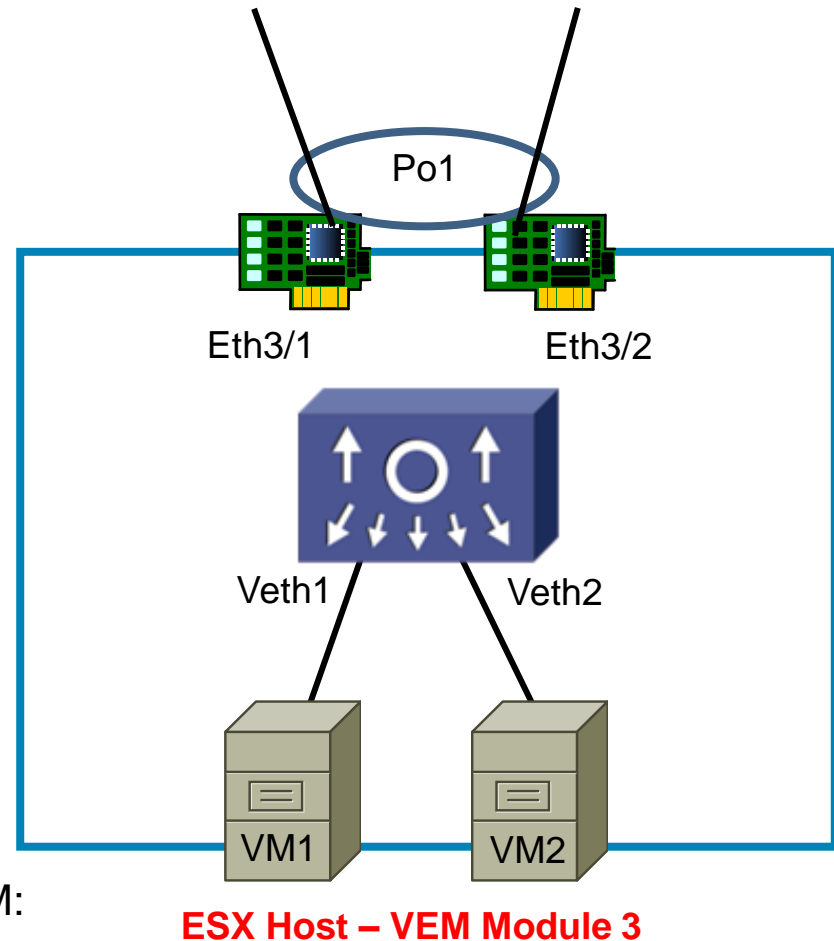
Up to 8 Port-Channels per ESX host

- Virtual Ethernet Port (veth)

One per virtual NIC (vnic) interface on a VM: includes service console and vmknics

No module number is assigned to keep naming persistent as VMs move between modules (ESX hosts/VEMs)

Up to 216 veths supported per host



Cisco Nexus 1000V

Switch Interface Types

```
Nexus1000V# show interface brief
```

Port	VRF	Status	IP Address	Speed	MTU
mgmt0	--	up	10.2.11.5	1000	1500

used for out of band management

Numbering based on Module #

Ethernet Interface	VLAN	Type	Mode	Status	Reason	Speed	Port Ch #
Eth3/4	1	eth	trunk	up	none	1000 (D)	1
Eth4/4	1	eth	trunk	up	none	1000 (D)	2

Physical NICs on ESX hosts

Port-Channeling of the Physical NICs out of the ESX hosts

Port-channel Interface	VLAN	Type	Mode	Status	Reason	Speed	Protocol
Po1	1	eth	trunk	up	none	a-1000 (D)	none
Po2	1	eth	trunk	up	none	a-1000 (D)	none

Interface	VLAN	Type	Mode	Status	Reason	MTU
Veth1	11	virt	access	up	none	1500
Veth2	11	virt	access	up	none	1500

Virtual interfaces corresponding to each VM attached

Port	VRF	Status	IP Address	Speed
ctrl0	--	up	--	1000

```
Nexus1000V#
```

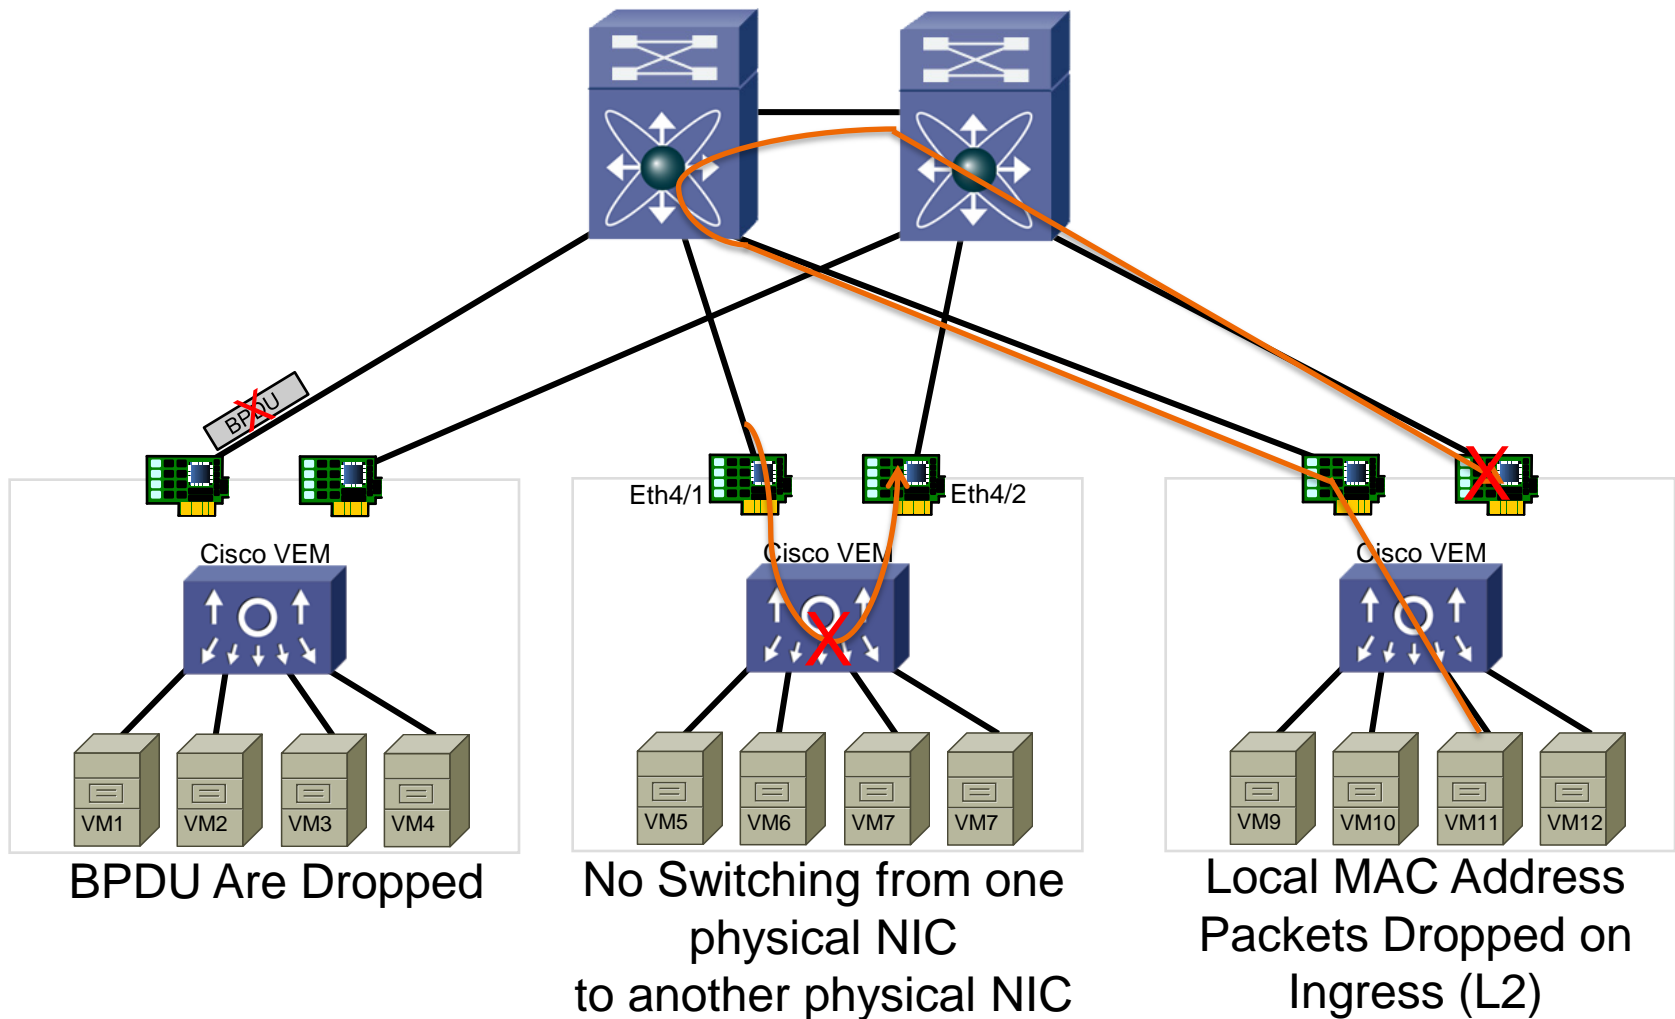
Cisco Nexus 1000V

Virtual Ethernet (vEth) Interface

- vEths are assigned sequentially (can not be manually changed or assigned)
- VM vNICs are statically bound to a vEth
 - Assignment persistent through reboots
 - May change if the vNIC is reassigned to another port profile
 - vEths will move between modules when a VM is moved (HA, Vmotion, DRS, etc.)
- Default virtual 'speed' is Gigabit as negotiated with the guest OS
 - By default performance is un-gated (i.e. 1Gb vNIC can run faster than 1Gb)
- 2048 vEths supported system-wide

Cisco Nexus 1000V Switching

Loop Prevention without STP




Upstream Switch Configuration

On interfaces upstream to the VEM uplink:

- Configure portfast
- Use BPDU guard
- Filter BPDUs
- Allow only required VLANs

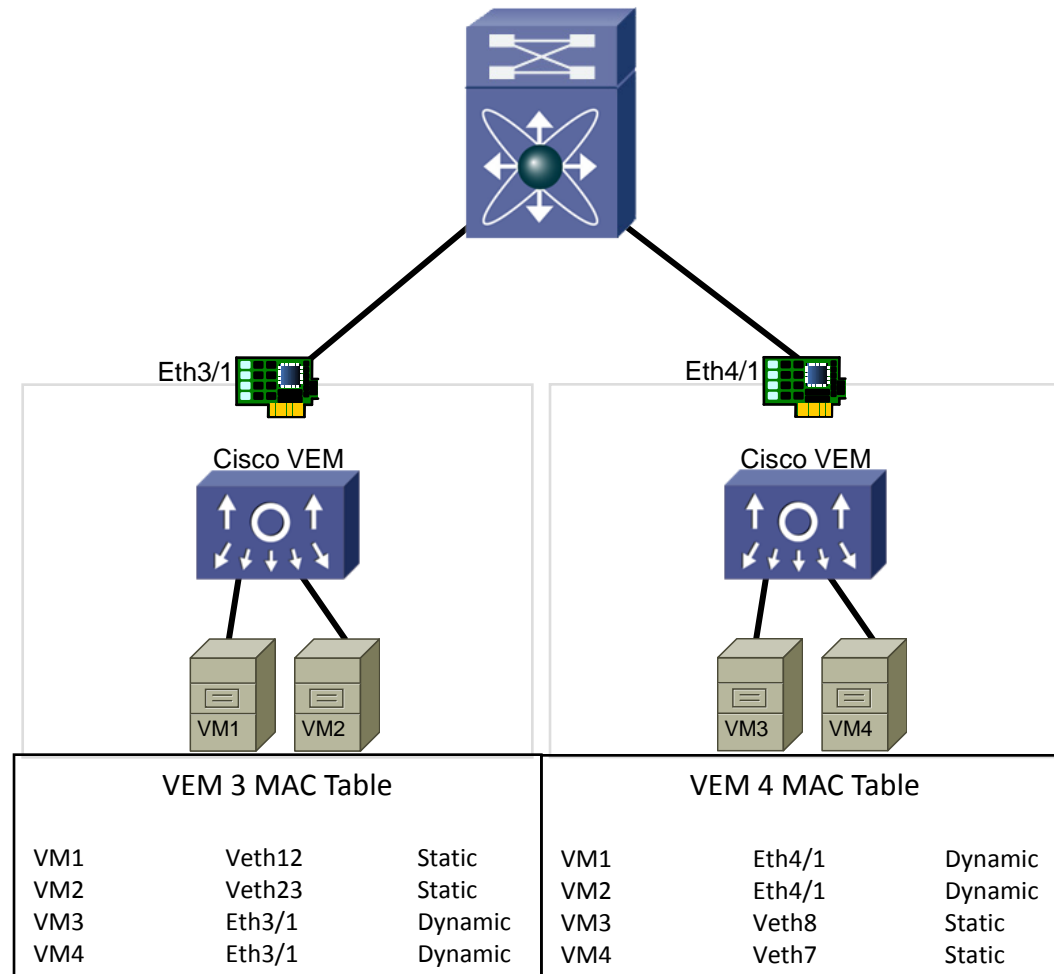
```
upstream-switch(config-if)# spanning-tree port type edge trunk
upstream-switch(config-if)# spanning-tree bpduguard enable
upstream-switch(config-if)# spanning-tree bpdufilter enable
upstream-switch(config-if)# switchport trunk allowed vlan 10-15
```



Cisco Nexus 1000V Switching

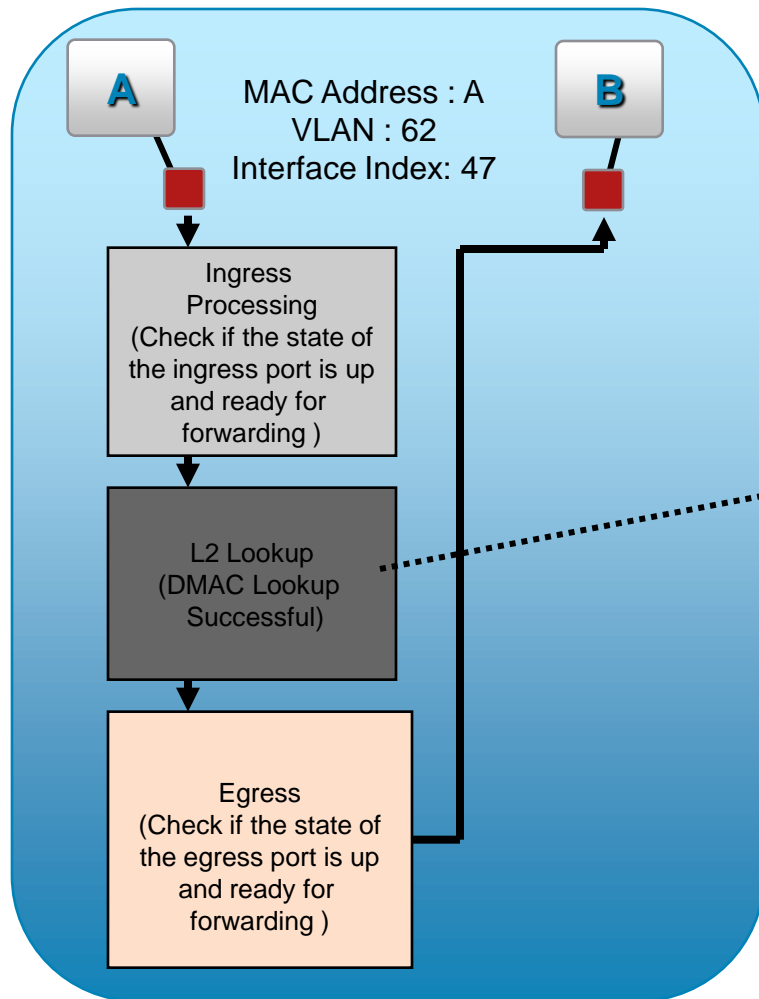
MAC Learning

- Each VEM learns independently and maintains a separate MAC table
- VM MACs are statically mapped
 - Other vEths are learned this way (vmknics and vswifs)
 - No aging while the interface is up
- Devices external to the VEM are learned dynamically
- 1000 MAC addresses per VLAN per VEM



Cisco Nexus 1000V Switching

Intra-VEM Switching

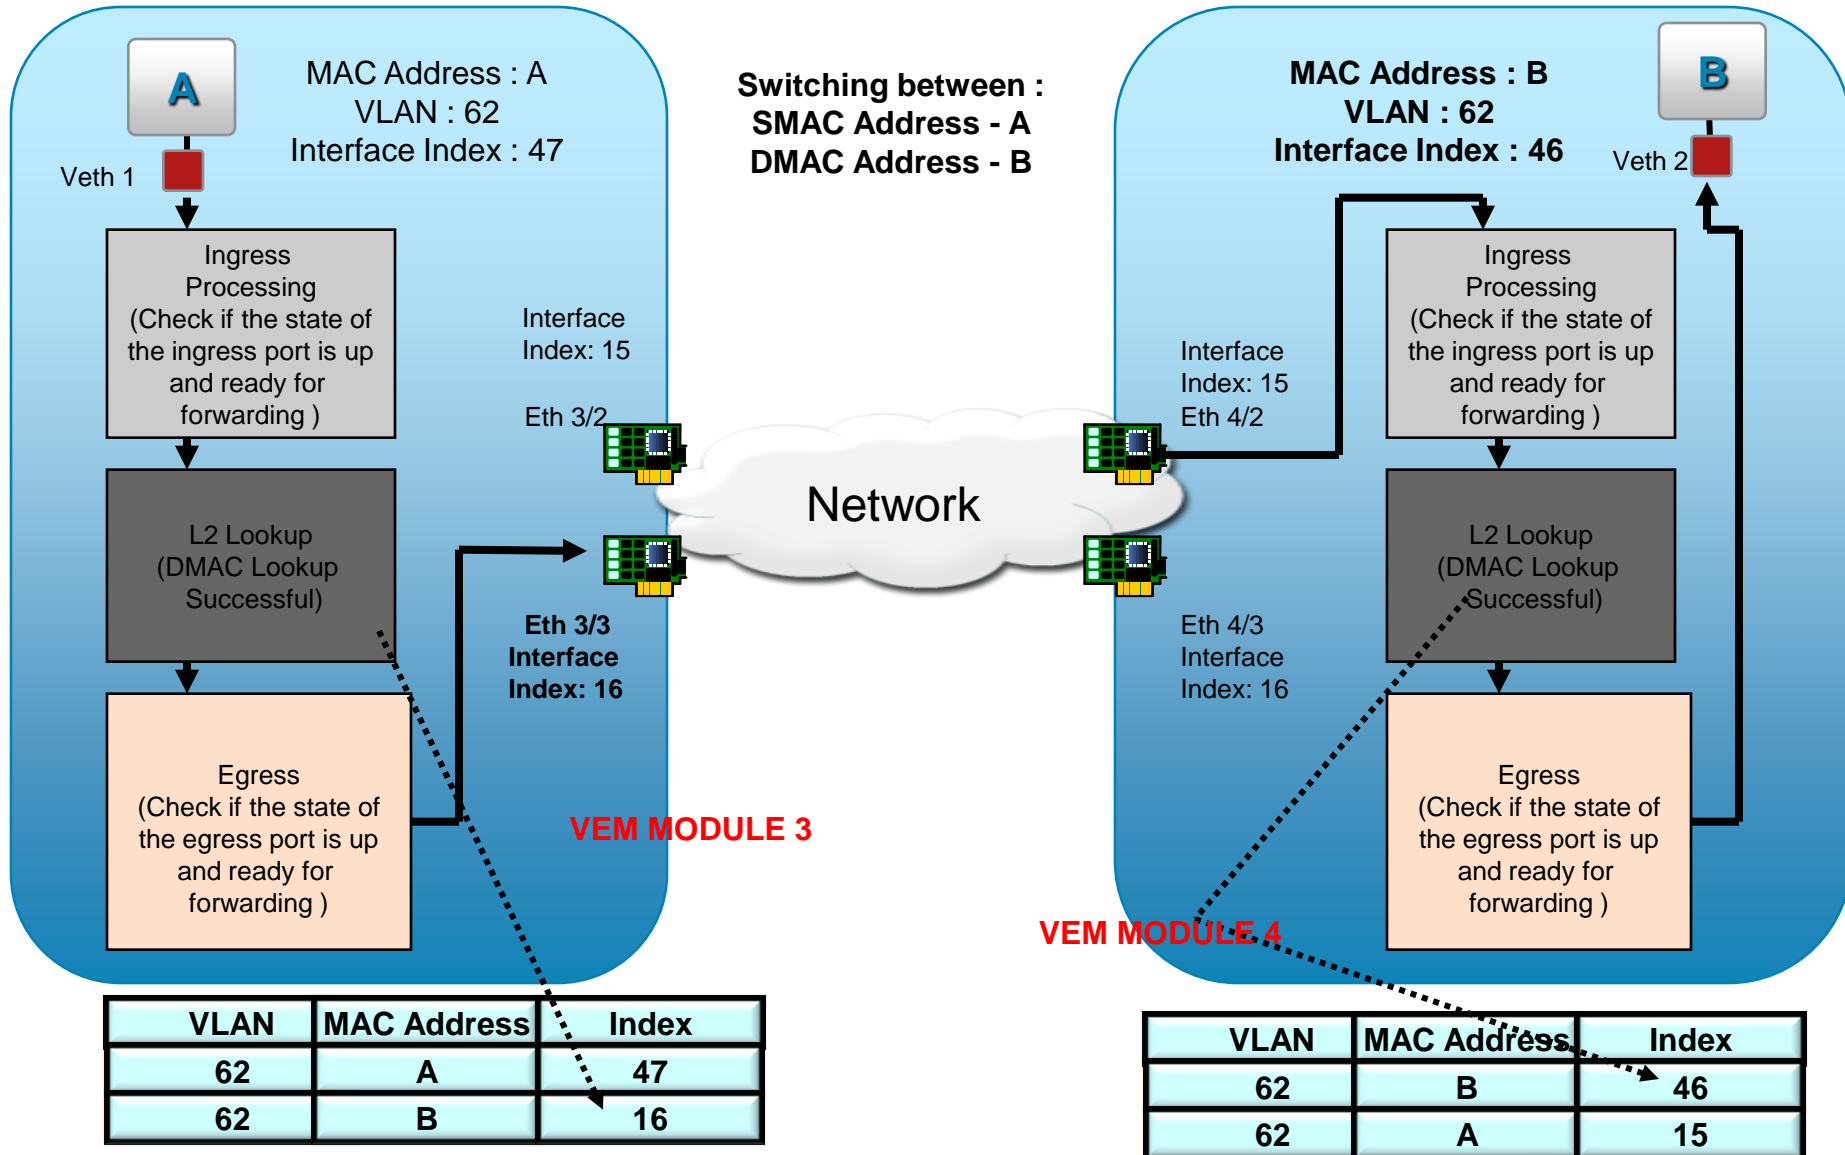


VLAN	MAC Address	Index
62	A	47
62	B	48

The MAC Address table supports 1k entries per VLAN per VEM. After 1k entries, VEM starts flooding.


Cisco Nexus 1000V Switching

Inter-VEM Switching



Cisco Nexus 1000V Switching


Multicast Behavior



- When a switch doesn't do any MAC learning, all broadcast, multicast and unknown unicast packets are flooded
- With VMs, this can present a challenge. In order to avoid loops from the end host, only one uplink can be designated to carry multicast traffic
- All virtual switches must be configured in a way that will drop duplicate multicast packets

Cisco Nexus 1000V Switching

Multicast Behavior



— Control VLAN
— Packet VLAN

Multicast Packet

In order to improve this behavior, the Nexus 1000V supports IGMP Snooping:

1. When a VM wants to participate to a multicast group it sends an IGMP join.
1. IGMP join message is sent to the VSM over the “packet” VLAN
1. VSM programs the VEM multicast tables and sends back the IGMP join message
2. IGMP join message is then forwarded to upstream device

Cisco Nexus 1000V Scalability

- A single Nexus 1000V supports:


- 2 Virtual Supervisor modules (HA)

- 64 Virtual Ethernet modules

- 512 Active VLANs

- 2048 Ports (Eth + Veth)

- 256 Port Channels




- A single Virtual Ethernet module supports:

- 216 Ports Veths

- 32 Physical NICs

- 8 Port Channels



https://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0_4_sv1_2/release/notes/n1000v_rn.html#wp42527

Agenda

- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- System Architecture Overview
- Switching Overview
- **Policy Management**
- Connectivity and Design

Cisco Nexus 1000V

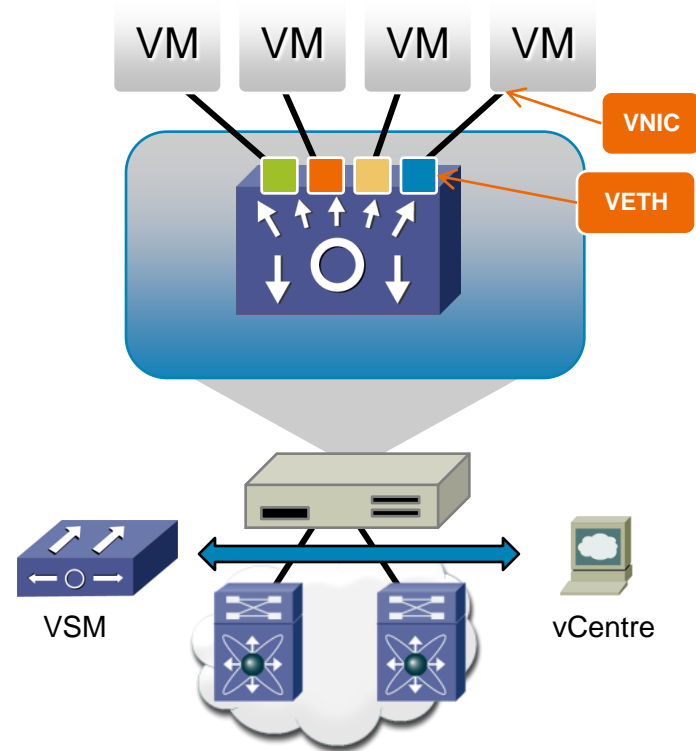
Port Profiles

- A port-profile is a container used to define a common set of network configuration commands for multiple interfaces
- Define once and apply many times to Physical and virtual NICs

****As of the 4.0(4)SV1(2) release, defined by “type Ethernet” or “type vEthernet” when creating – no longer need “capability uplink” under the port-profile config**

```
n1000v(config)# port-profile type vethernet WebServers
n1000v(config-port-prof)# switchport mode access
```

- Simplifies management by storing interface configuration
- Key to collaborative management of virtual networking resources



Cisco Nexus 1000V

Port Profile Configuration

```
n1000v# show port-profile name WebServers
port-profile WebServers
  description:
  type: vethernet
  status: enabled
  capability l3control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: none
  port-group: WebServers
  max ports: 32
  inherit:
```

config attributes:

```
switchport mode access
switchport access vlan 100
no shutdown
```

evaluated config attributes:

```
switchport mode access
switchport access vlan 100
no shutdown
```

assigned interfaces:

```
Vethernet9
```


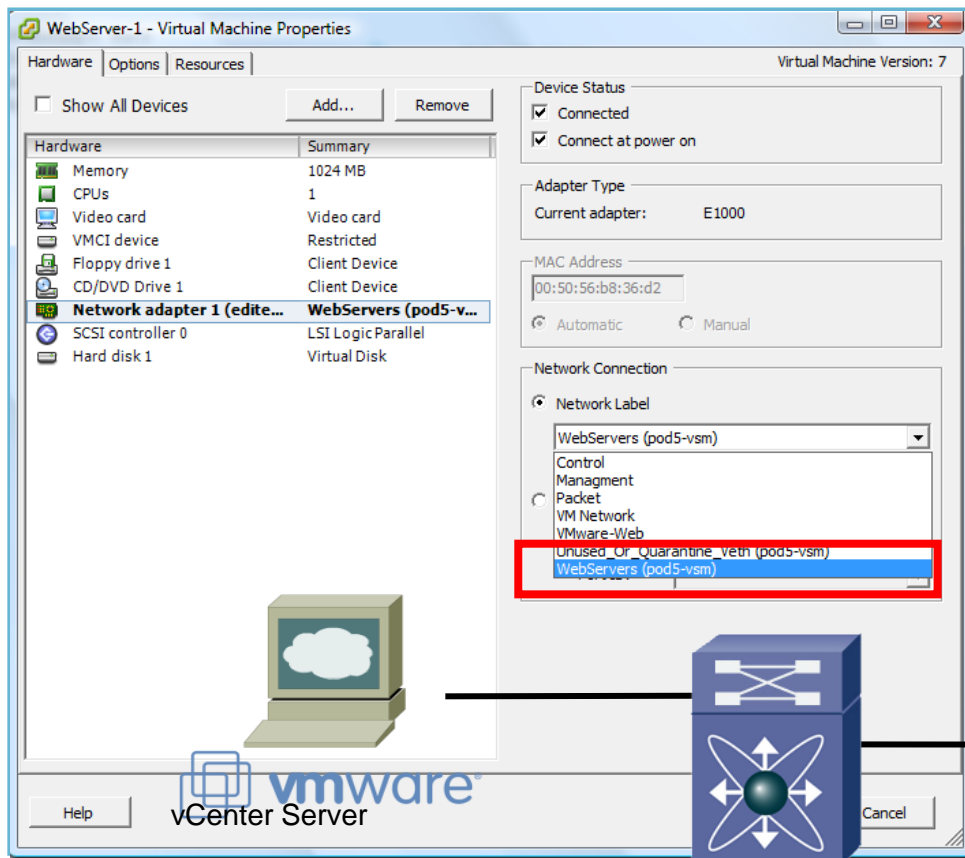
Support Commands include:

- ✓ Port management
- ✓ VLAN
- ✓ PVLAN
- ✓ Port-channel
- ✓ ACL
- ✓ NetFlow
- ✓ Port Security
- ✓ QoS
- ✓ Port-Mirroring (ERSPAN)

Cisco Nexus 1000V

Port Profile Distribution

```
n1000v(config)# port-profile WebServers
n1000v(config-port-prof)# vmware port-group
n1000v(config-port-prof)# switchport mode access
n1000v(config-port-prof)# switchport access vlan 100
n1000v(config-port-prof)# no shut
n1000v(config-port-prof)# state enable
```




Cisco Nexus 1000V

Uplink Port Profile Configuration

- Special profiles that define physical NIC properties
- Usually configured as a trunk
- Defined by “type Ethernet” when creating the port-profile (default it “type vEthernet”)
- Uplink profiles cannot be applied to vEths
- Non-uplink profiles cannot be applied to physical NIC interfaces
- Only selectable in vCenter when adding a host or additional NICs

```
n1000v(config)# port-profile type Ethernet DataUplink
n1000v(config-port-prof)# vmware port-group
n1000v(config-port-prof)# switchport mode trunk
n1000v(config-port-prof)# switchport trunk allowed vlan 10-15,51,52
n1000v(config-port-prof)# system vlan 51, 52
n1000v(config-port-prof)# channel-group auto mode on mac-pinning
n1000v(config-port-prof)# no shut
```

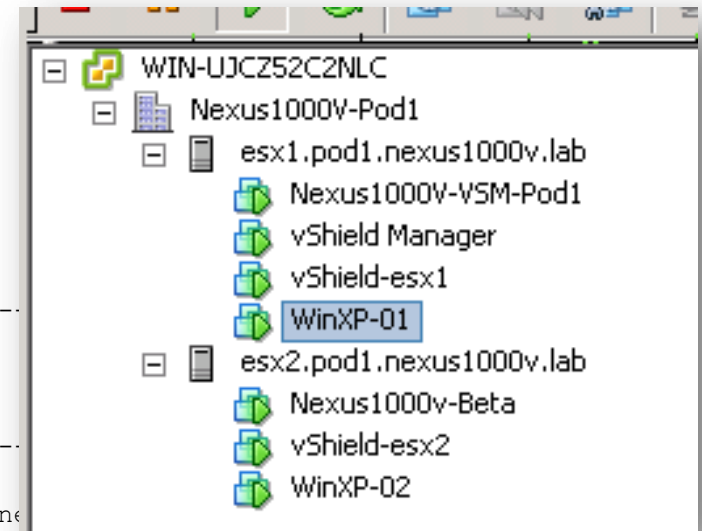


Cisco Nexus 1000V

VM Visibility

Pod1-VSM# **show interface virtual**

Port	Adapter	Owner	Mod	Host
Veth1	vmk1	VMware VMkernel	3	esx1.pod1.nexus1000v.la
Veth2	vmk1	VMware VMkernel	4	esx2.pod1.nexus1000v.la
Veth3	Net Adapter 1	Nexus1000V-VSM-Pod1	3	esx1.pod1.nexus1000v.la
Veth4	Net Adapter 1	Nexus1000v-Beta	4	esx2.pod1.nexus1000v.la
Veth5	Net Adapter 1	vShield-esx1	3	esx1.pod1.nexus1000v.la
Veth6	Net Adapter 1	vShield Manager	3	esx1.pod1.nexus1000v.la
Veth7	Net Adapter 1	vShield-esx2	4	esx2.pod1.nexus1000v.la
Veth8	Net Adapter 1	WinXP-01	3	esx1.pod1.nexus1000v.la
Veth9	Net Adapter 1	WinXP-02	4	esx2.pod1.nexus1000v.la



Cisco Nexus 1000V

Tracking VM Statistics

```
Pod1-VSM# show interface veth8
```

```
Vethernet8 is up
```

```
< ---- SNIP --- >
```

```
Port mode is trunk
```

```
5 minute input rate 0 bits/second, 0 packets/second
```

```
5 minute output rate 40 bits/second, 0 packets/second
```

```
Rx
```

```
426 Input Packets 125 Unicast Packets
```

```
15 Multicast Packets 286 Broadcast Packets
```

```
50941 Bytes
```

```
Tx
```

```
81182 Output Packets 136 Unicast Packets
```

```
18 Multicast Packets 81028 Broadcast Packets 81046 Flood Packets
```

```
8387936 Bytes
```

```
1 Input Packet Drops 0 Output Packet Drops
```



Agenda

- Introduction to Virtual Networking
- Cisco Nexus 1000V Introduction
- System Architecture Overview
- Switching Overview
- Policy Management
- **Connectivity and Design**

Cisco Nexus 1000V

Connectivity Options

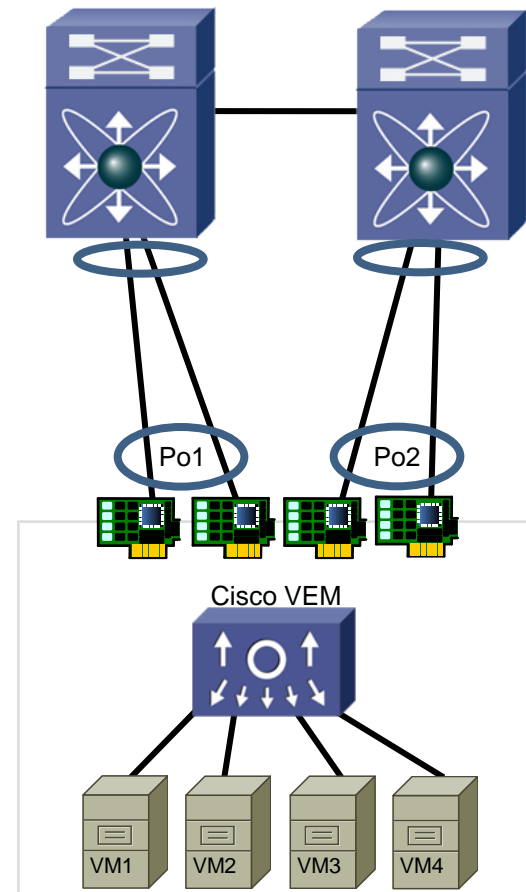
The Nexus 1000V has different ways of connecting the VEM to the upstream switch:

- Regular port-channel (recommendation LACP port-channel)
- vPC Host Mode (recommendation “mac-pinning”)
- Static pinning a port-profile to a sub-group

Cisco Nexus 1000V


Uplink Port Channel Configuration

- Standard Cisco Port Channels
 - Behaves like EtherChannel
- Link Aggregation Control Protocol (LACP) Support
- 17 hashing algorithms available
 - Selected either system-wide or per module
 - Default is source MAC
- Host side configured within the port-profile Port Profiles (don't forget to configure the upstream!)



Cisco Nexus 1000V

PC and vPC




- For redundancy, hosts are often connected using multiple links (and often to multiple devices)
- Without port-channels, the same MAC address will be seen on two different ports
- This is avoided when the links connecting the host to a switch are bundled together so that the MAC address is now seen on the same logical port

An arrow points from the physical switch in the diagram to a table. The table has two columns: 'MAC Address' and 'Port'. The first row is highlighted in yellow and has a red line extending from both sides. The second, third, and fourth rows are light blue. The fifth row is highlighted in yellow and has a red line extending from both sides.

MAC Address	Port
A	Po 1
B	Gig 1/1
C	Gig 1/2
D	Gig 1/2
A	Po 1

Cisco Nexus 1000V

PC and vPC




- Virtual Port-Channels takes host redundancy up one level by allowing a single device to be dual attached and forward out both paths without a STP loop
- vPCs allow two upstream switches to act as one
- The port-channel is configured across the 2 physical switches and the MAC address is seen on both links on both switches

MAC Address	Port
A	Po 1
B	Po 1
C	Po 1
D	Po 1
A	Po 1

Cisco Nexus 1000V

vPC-Host Mode


- Allows a single PC to span multiple upstream **NON-CLUSTERED** upstream switches using “subgroups”
- 1st release “**subgroup cdp**”: formed subgroups based on CDP information from upstream
 - Limitation of 2 subgroups (made deployments with FEX and UCS difficult)
 - Subgroups could be manually defined outside of a port-profile
- 2nd release “**mac-pinning**”: forms “subgroups” based on uplinks
 - Each uplink (or bundled PC uplinks) is assigned a subgroup – up to 32
 - VMs pinned to an uplink upon boot up



Cisco Nexus 1000V

MAC-Pinning

MAC Pinning provides the dynamism of VPC Host-Mode without requiring CDP to be configured on the upstream switch




The MAC address of the VM will be used to select which link to use

This is the same hashing used by vmware today

```
Nexus1000(config)#port-profile type ethernet sys-uplink
Nexus1000(config-port-prof)#no shut
Nexus1000(config-port-prof)#channel-group auto mode on mac-pinning
Nexus1000(config-port-prof)#switchport mode trunk
Nexus1000(config-port-prof)#switchport trunk allowed vlan 10-25
Nexus1000(config-port-prof)#state enabled
Nexus1000(config-port-prof)#vmware port-group
```

Cisco Nexus 1000V

MAC-Pinning Failover




If a failure occurs within the network, all the traffic pinned to the failed interface will automatically failover and be re-pinned to the other available interfaces

```
Nexus1000(config)#port-profile type ethernet sys-uplink
Nexus1000(config-port-prof)#no shut
Nexus1000(config-port-prof)#channel-group auto mode on mac-pinning
Nexus1000(config-port-prof)#switchport mode trunk
Nexus1000(config-port-prof)#switchport trunk allowed vlan 10-25
Nexus1000(config-port-prof)#state enabled
Nexus1000(config-port-prof)#vmware port-group
```

Cisco Nexus 1000V

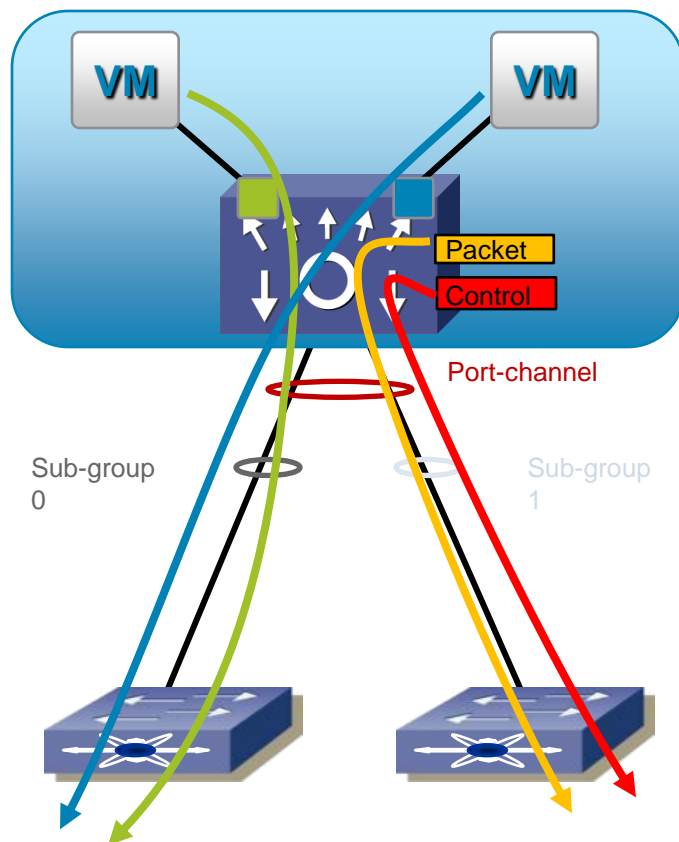
Static Pinning



- By default, VM traffic is pinned in a round-robin fashion to the available uplinks within a sub group
- ***static pinning*** allows customers to select specific uplinks for specific traffic types based on need and availability requirements
- For example, with static pinning, customers can use one link for **SERVICE CONSOLE** and another for **VMOTION** yet still retain active/standby failover for both traffic types

Cisco Nexus 1000V

Static Pinning Example



Static Pinning has options for port-profiles as well as **control** and **packet traffic** to be pinned to a specific sub-group

```
Nexus1000(config)#port-profile green  
Nexus1000(config-port-prof)#pinning sub-group id 0
```

```
Nexus1000(config)#port-profile blue  
Nexus1000(config-port-prof)#pinning sub-group id 0
```

```
Nexus1000(config)#port-profile type ethernet uplink  
Nexus1000(config-port-prof)# pinning control-vlan 1  
Nexus1000(config-port-prof)# pinning packet-vlan 1
```

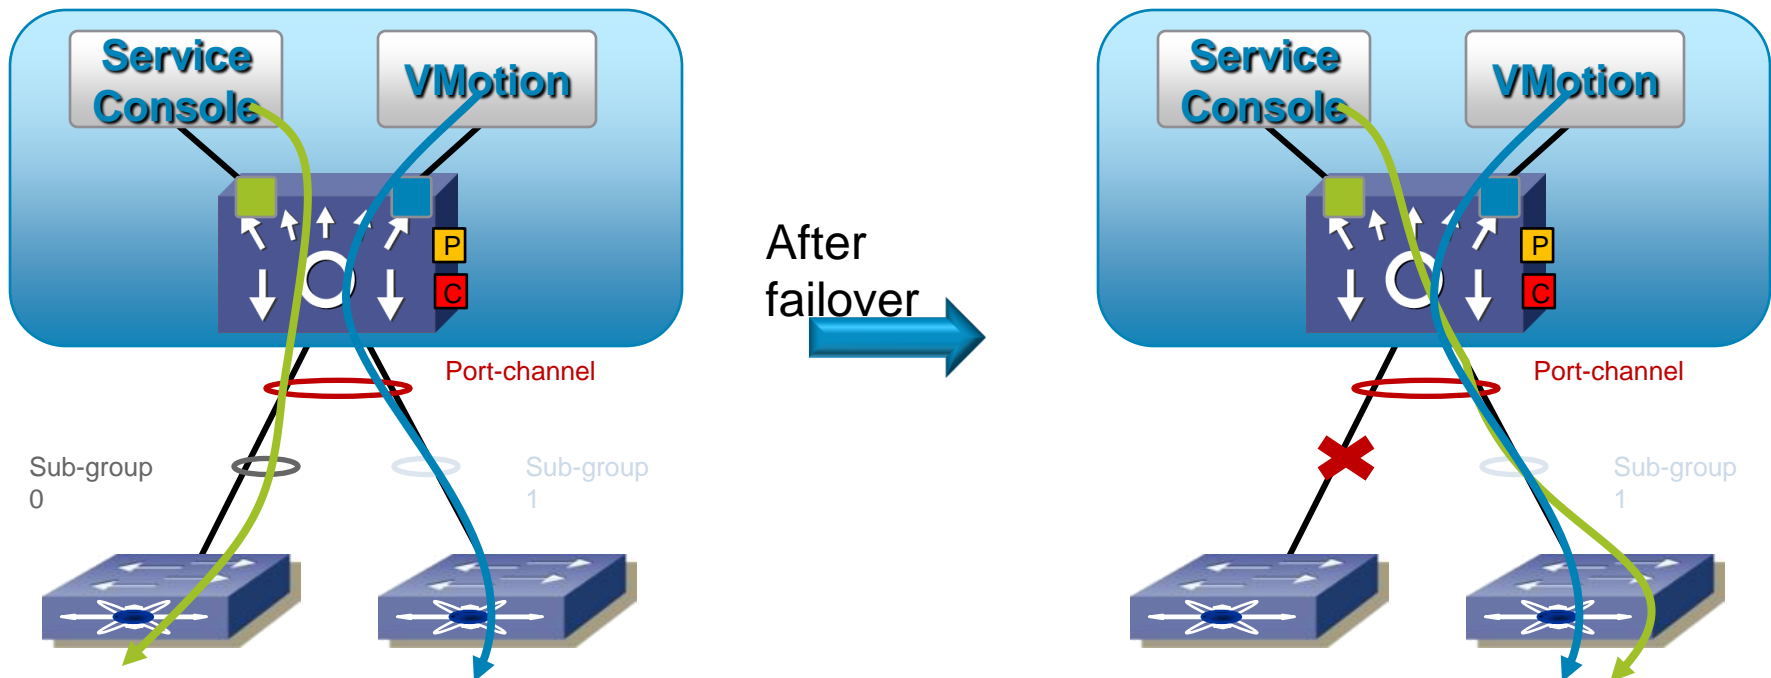
Cisco Nexus 1000V

Static Pinning Failover

One of the big advantages of **static pinning** is to migrate the **active/standby** design that customers have been deploying with VMware vSwitch.


```
Nexus1000(config)#port-profile service-console  
Nexus1000(config-port-prof)#pinning sub-group id 0
```

```
Nexus1000(config)#port-profile vmotion  
Nexus1000(config-port-prof)#pinning sub-group id 1
```



Cisco Nexus 1000V

LACP Connectivity



upstream switches are “clustered”
(vPC, VSS, VBS, Stack...)

- LACP is a port-channeling control mechanism used to ensure proper configuration of a port-channel on both ends

Allows for the VMs and VMKernel interfaces to utilize more than one link for traffic

Allows for fast vMotion and faster VM connectivity by using flow based hashing

- *****BEST PRACTICE** when connecting to upstream switches which are “clustered”

LACP is also a differentiator when comparing to VMware DVS/vSwitch

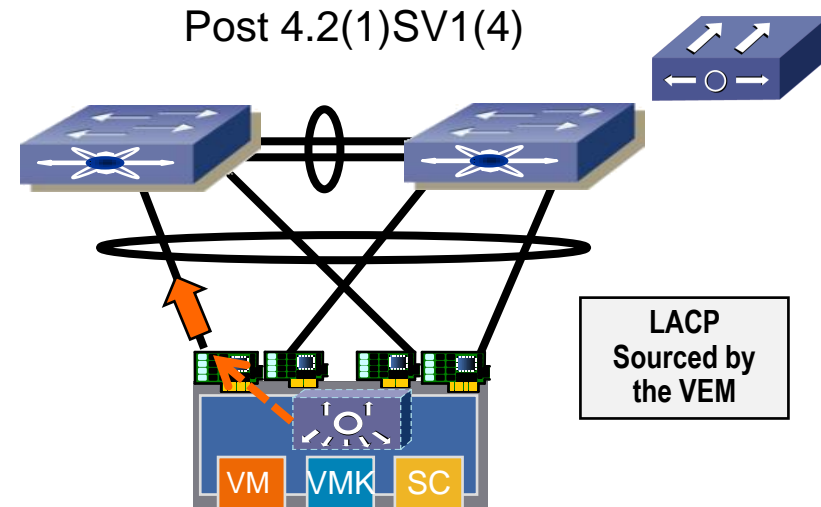

- Is NOT used with vPC-Host Mode
- LACP is configured within the port-profile for the physical interfaces

```
channel-group auto mode [active/passive]
```

Data Centre Access Architecture

Connecting Nexus 1000V to MCEC capable upstream devices

- Prior to the 4.2(1)SV1(4) release it was best practice to leverage *static configuration of 802.3ad port channels when running 1kv with FCoE*
- Prior to this release LACP bootstrap problem could impact certain topologies
 - VEM brings up system VLAN and uses one uplink to communicate with VSM
 - VSM sends LACP PDU to the VEM inside packet channel
 - VEM sends and receives LACP PDU with upstream switch
 - Port Channel negotiated with upstream switch
- If VEM to VSM communication had issues uplink port channel did not come up
- Post 4.2(1)SV1(4) use of **active** mode for port channels is preferred



Cisco Nexus 1000V

Port-Channel Hashing Options

```
pod5-vsm(config) # port-channel load-balance ethernet ?
dest-ip-port                Destination IP address and L4 port
dest-ip-port-vlan          Destination IP address, L4 port and VLAN
destination-ip-vlan        Destination IP address and VLAN
destination-mac             Destination MAC address
destination-port           Destination L4 port
source-dest-ip-port        Source & Destination IP address and L4 port
source-dest-ip-port-vlan   Source & Destination IP address,L4 port and VLAN
source-dest-ip-vlan        Source & Destination IP address and VLAN
source-dest-mac            Source & Destination MAC address
source-dest-port           Source & Destination L4 port
source-ip-port             Source IP address and L4 port
source-ip-port-vlan        Source IP address, L4 port and VLAN
source-ip-vlan             Source IP address and VLAN
source-mac                 Source MAC address
source-port                Source L4 port
source-virtual-port-id     Source Virtual Port Id
vlan-only                  VLAN
```

Cisco Nexus 1000V

Port-Channel Hashing Differences

- Source Based Hashing

Hashes all traffic from a single source down the same link
vPC-HM requires no upstream special configuration
(EtherChannel)

Examples are source MAC, VLAN, Virtual Port

- Flow Based Hashing


Each flow may take a different path

vPC may require EtherChannel upstream

Examples include any hash using dst, L4 port,
or combinations of src/dst/port

Cisco Nexus 1000V

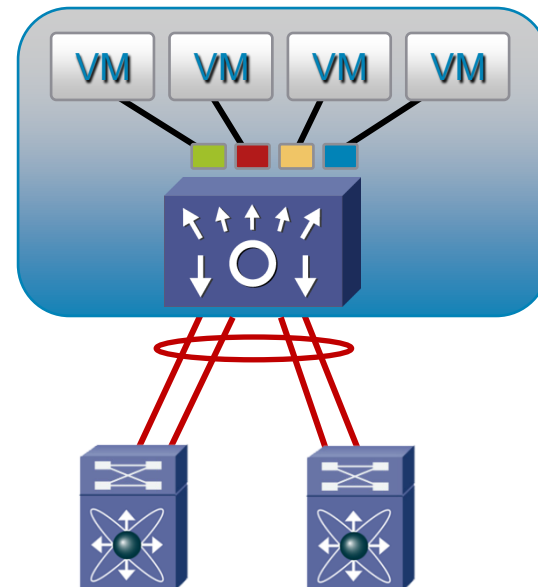
Best Practice Configurations



If the upstream switch can be clustered (VPC, VBS Stack, VSS) **use LACP**

****LACP offload, available in the 1.4 release, should be used in any FCoE environments or boot from SAN environments**

If the upstream switch can NOT be clustered **use MAC-PINNING**



1.4 Coming Features

- ISSU
- Class Based WFQ
- ERSPAN
- Network State Tracking
- LACP Offload
- HTTP Server
- Virtual Security Gateway
- LISP **

Summary

- Nexus1000v provides insight into vnic's for added security, troubleshooting, and in general better visibility in the network operations within ESX servers
- Nexus 1000v provides a better operational interface between network admin and server admin
- Nexus1000v provides a comprehensive feature set for securing datacenter virtualized environments
- Nexus 1000v offers a flexible solution for virtual environments living on any hardware in any data center

BRKVIR-2006

Recommended Reading

Please browse on-site Cisco Store for suitable reading.

Please complete your Session Survey

- We value your feedback - don't forget to complete your online session evaluations after each session. Complete 4 session evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Networkers 20th Anniversary t-shirt.
- All surveys can be found on our onsite portal and mobile website: www.ciscoliveeurope.com/connect/mobi/login.ww
- You can also access our mobile site and complete your evaluation from your mobile phone:
 1. Scan the Access Code
(See <http://tinyurl.com/qrmelist> for software, alternatively type in the access URL)
 2. Login
 3. Complete and Submit the evaluation





CISCO

For More information

- Nexus1000v Best Practices with UCS:
- http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-558242.html