Inside the Nexus 1000V Virtual Switch

BRKVIR-2012

Agenda

How Virtualization Affects the Network

… and How the Nexus 1000V can Help

- Two Administrators' Views of the Nexus 1000V
- Virtual Machines and Virtual Ports
- Virtual Machine Migration: Moving Targets
- The "Physical" Nexus 1000V
- Inside the Nexus 1000V
- Port Channels and Upstream Switches
- Deploying Nexus 1000V
- What's New?

What Are We Talking About?

- Virtualization is here, and becoming more dynamic
- Network management and monitoring disrupted
  - Physical → virtual
  - Static → dynamic
  - Physical boundary changes
  - Network edge moves outside the physical switch
- Network → active part of OS virtualization layer
  - **The Nexus 1000V**
- **How does the Nexus 1000V work**
  - **And how is it different from a physical switch?**

A Tale of Two Administrators…

From Physical to Virtual

When moving from physical to virtual…

- Applications are the Same
- Network Policies are the Same
- Network and Server Admins are the Same
- But
  - The network edge is now virtual, and lives in the host
  - VM lifecycles move from months/years → weeks/days
  - Ports on VMs are moving targets for the network admin
  - Port count dramatically increases

?

Good Fences

- With physical machines
- Each administrator controls his or her own side
- Physical cable is the interface

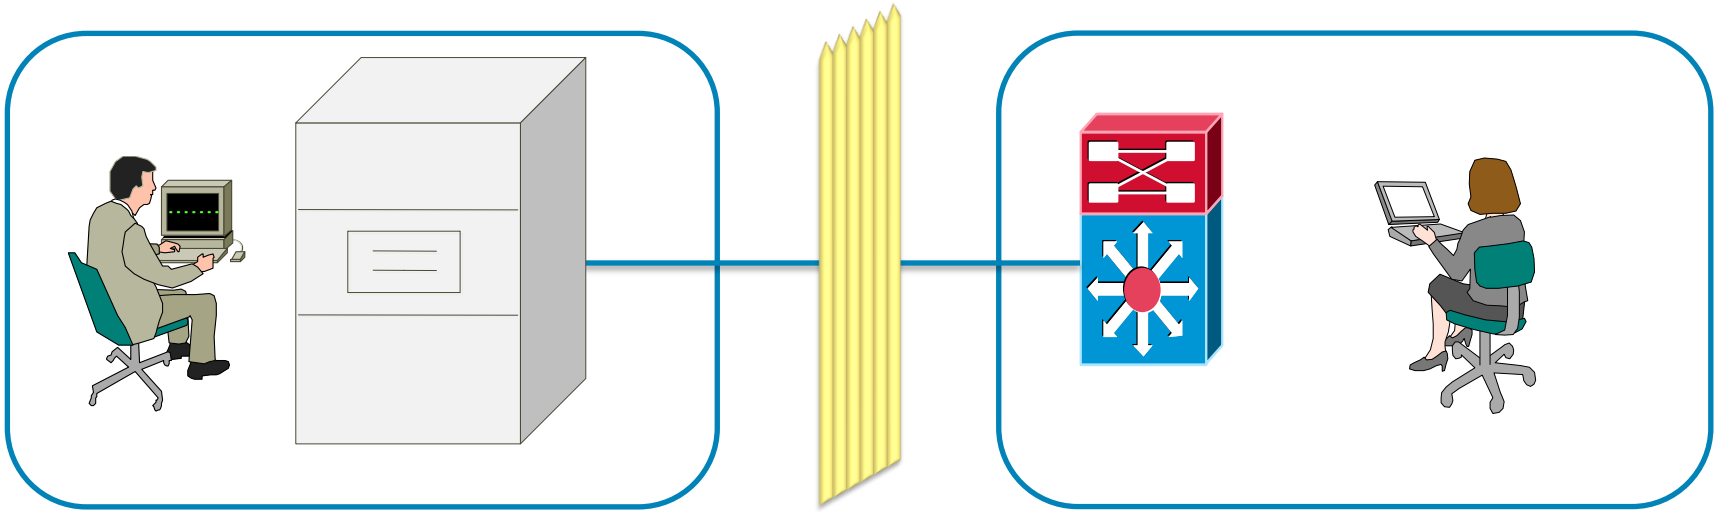

Managing the Network from vCenter

The Network is In the Computer

- The network edge has moved into the server
- Virtual network is managed by server admin
- Trust and boundary issues
- Network is a barrier to deploying virtualization

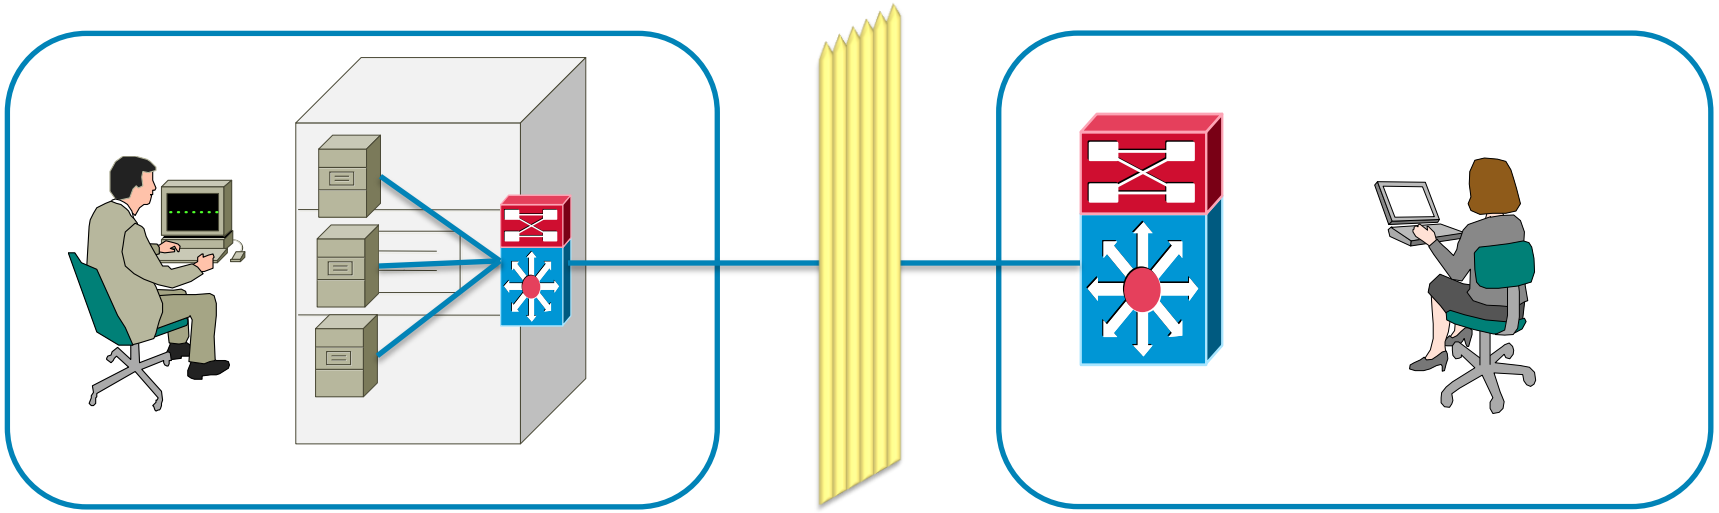

Managing the Network from vCenter

Before the Nexus 1000V

- Server admins want to deploy VMware
- Asked network admins for trunk ports

What could go wrong?

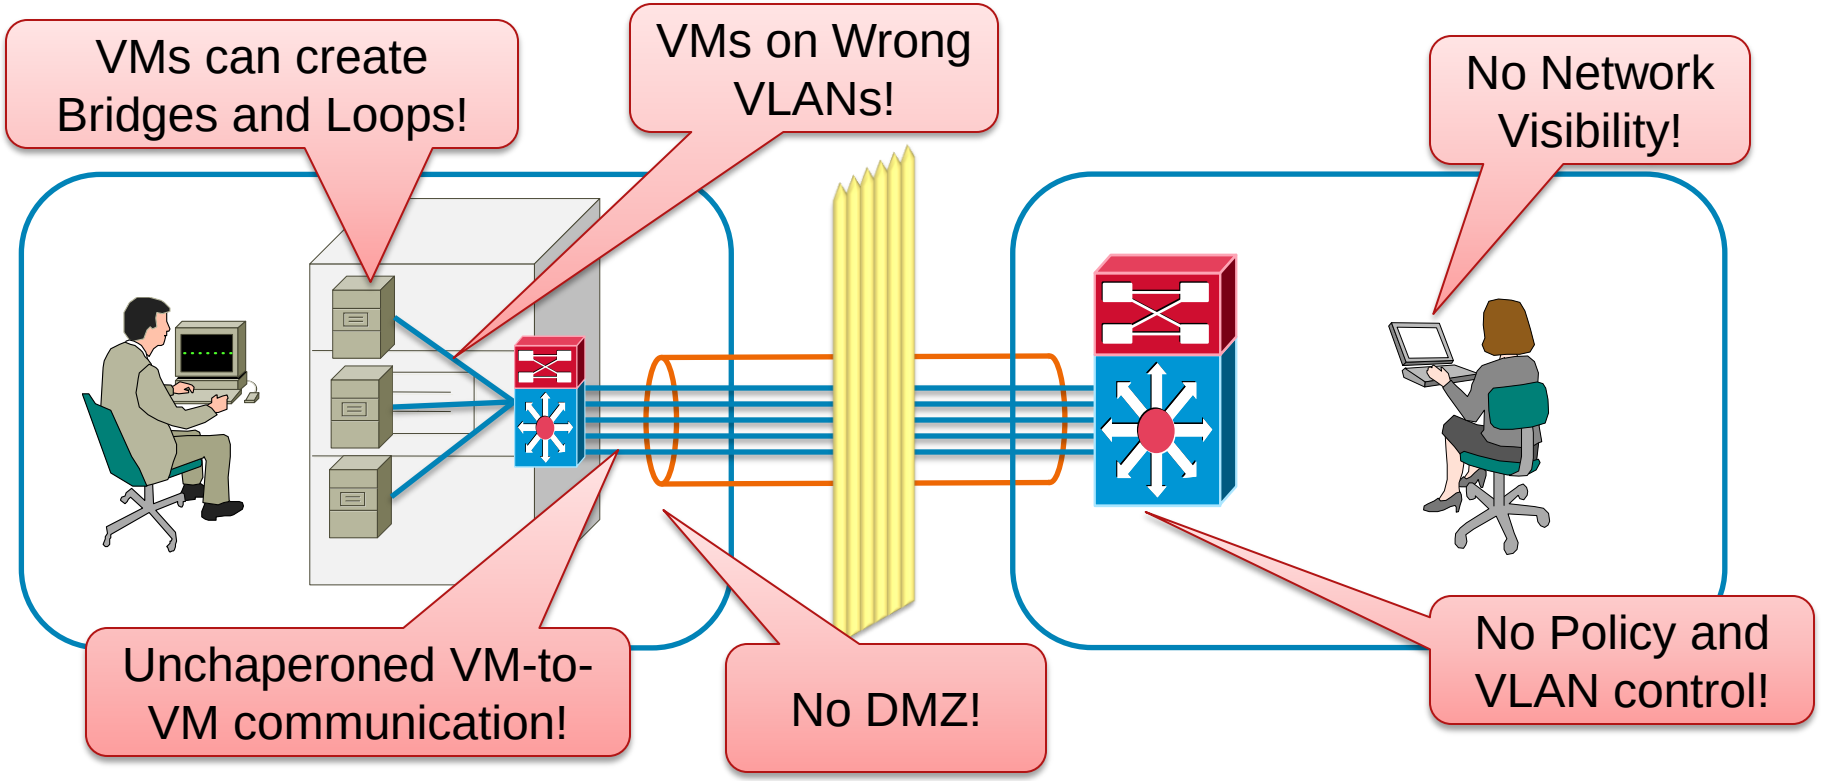

Early Customer Solution

Solved a subset of issues, but…

- Network allowed one VLAN per Physical NIC
- Many physical ports needed on switches
- Severely limited networking
- Virtualization only for development projects

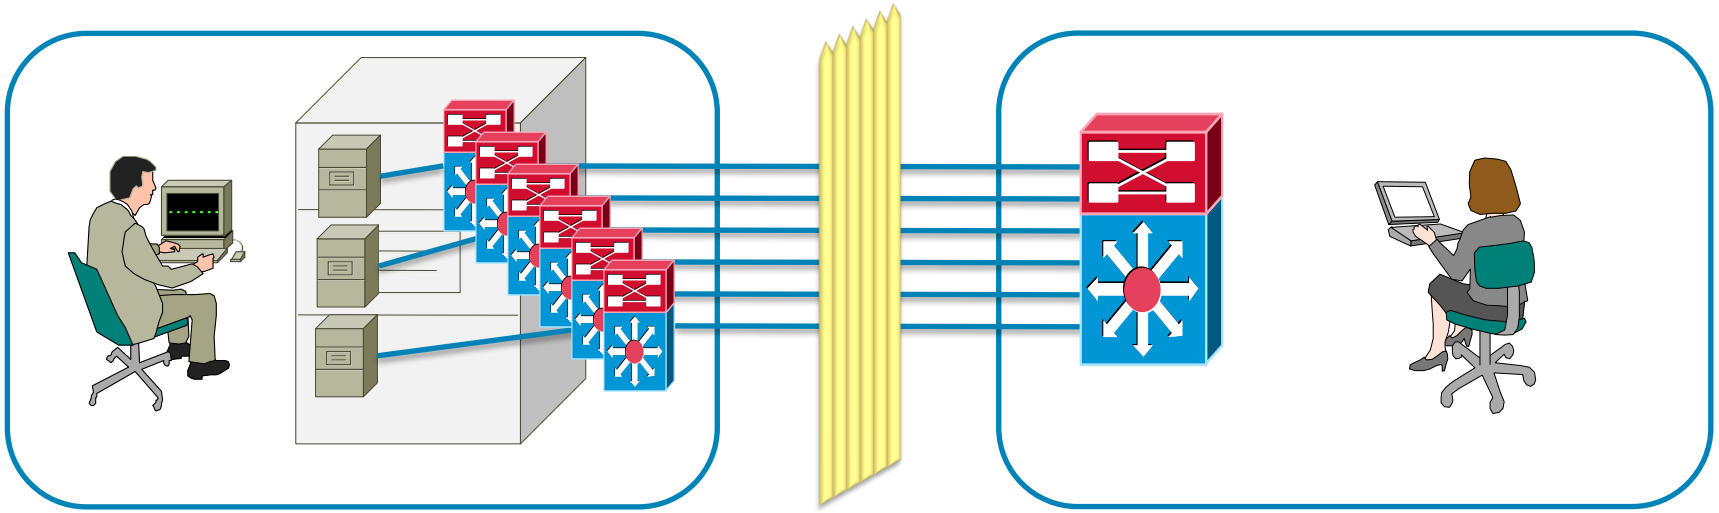

Managing the Network from Nexus 1000V

Solution!

- Virtual Cisco switch extends into the server
- Virtual network managed by network admin
- Boundary is restored

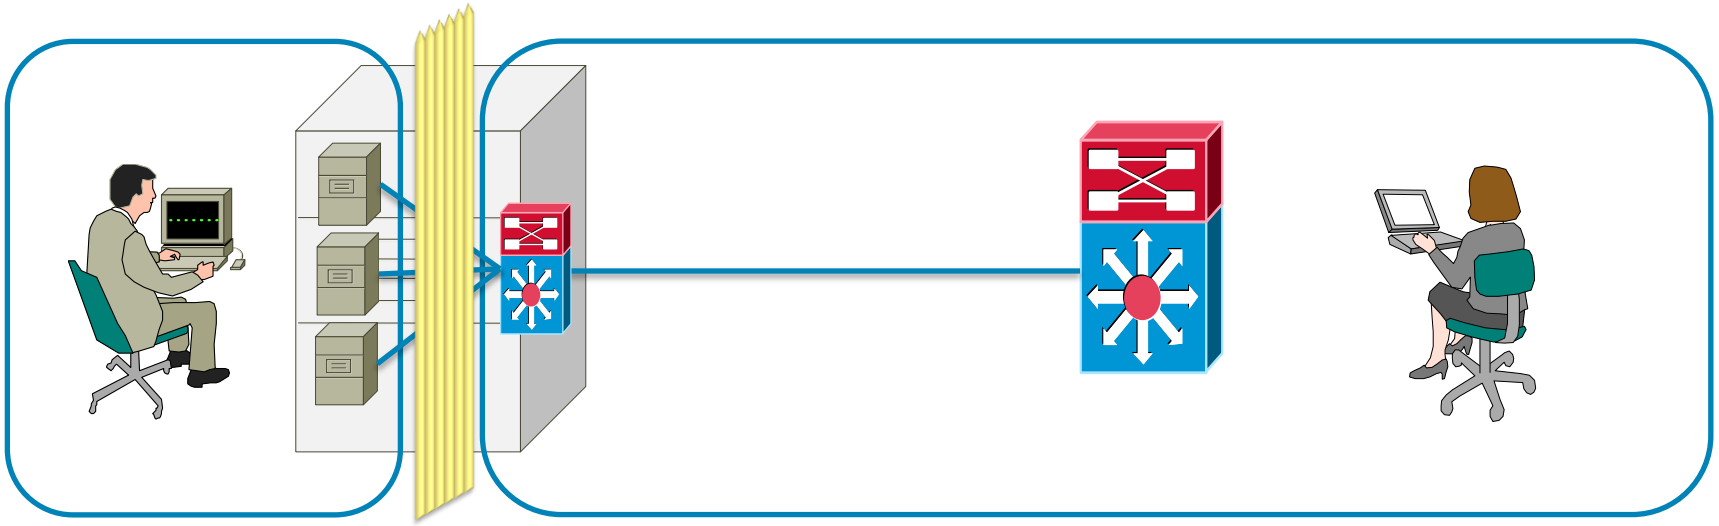

Two-Way Visibility

Opportunity!

- “Chain link” fence keeps configuration secure
- But allows visibility in both directions
- Transparency and accountability help build trust
- And can make even better neighbors

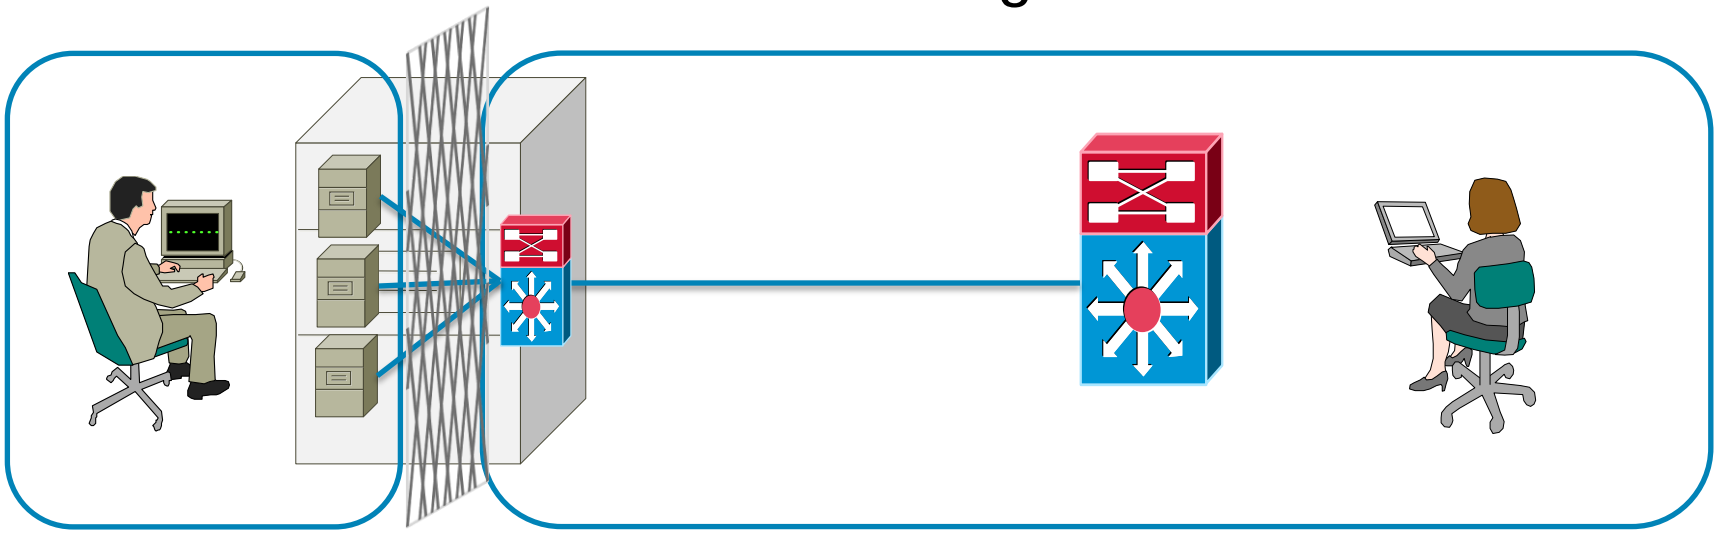

The Nexus 1000V

Virtual Supervisor Module (VSM)

Take Two; They're Small!

- Familiar CLI interface into the Nexus 1000V
- Leverages NX-OS
- Controls multiple VEMs as a single network device
- Runs as a virtual machine
- Or a physical appliance (Nexus 1010)

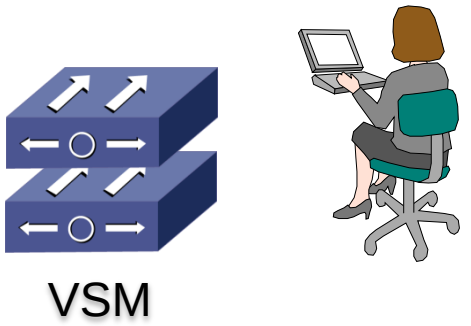

Virtual Ethernet Module (VEM)

A Switch in Every Host…

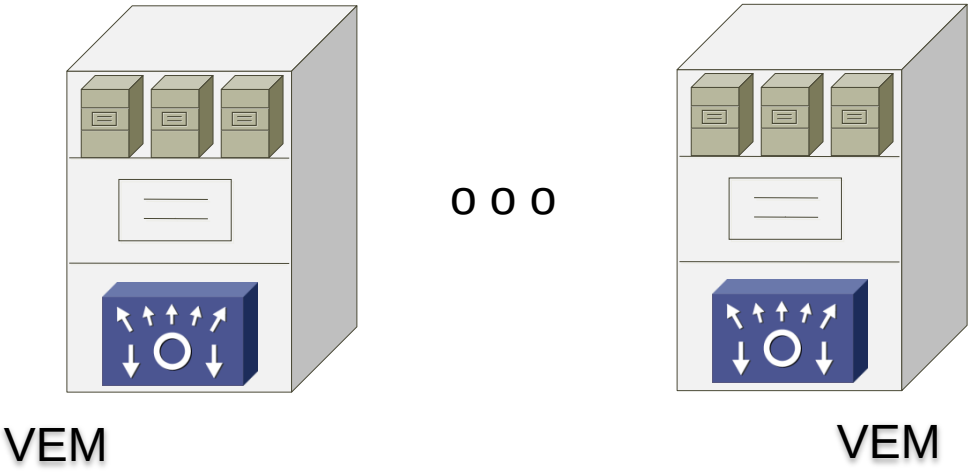

- Enhances VMware switch architecture
- Enables advanced switching capability on the hypervisor
- Provides each VM with dedicated “switch ports”

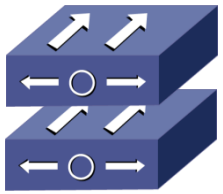

VSM

The Nexus 1000V

Nexus 1000V is a Distributed Edge Switch

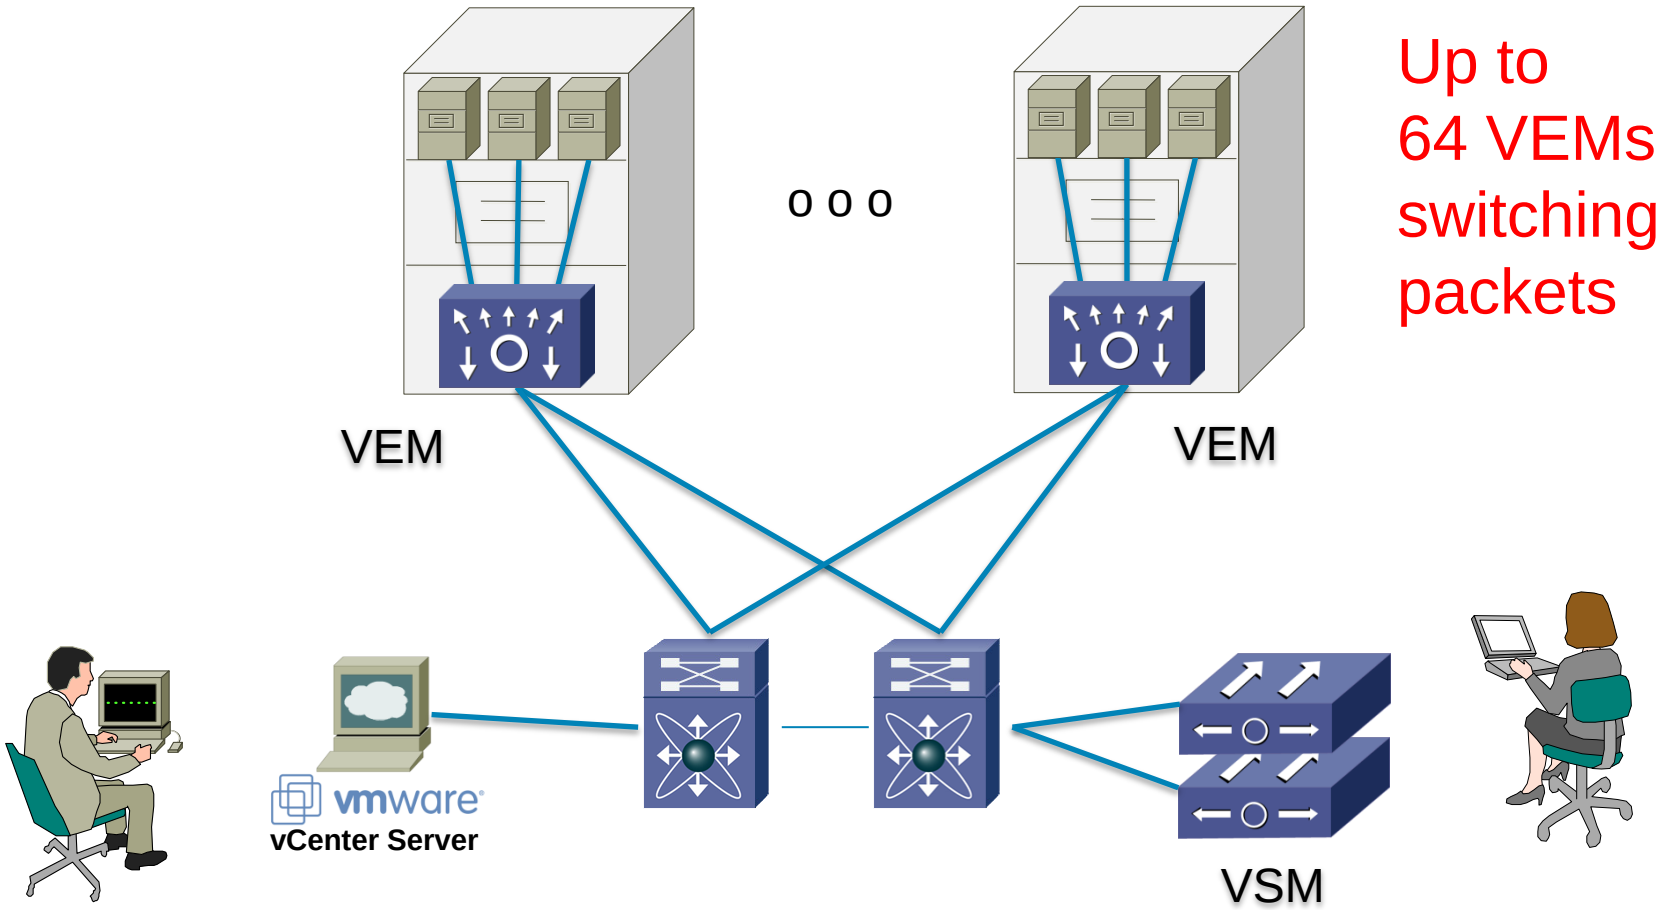

Up to
64 VEMs
switching
packets

Only one switch to manage!

Nexus 1000V on vCenter

Nexus 1000V is a Distributed Virtual Switch

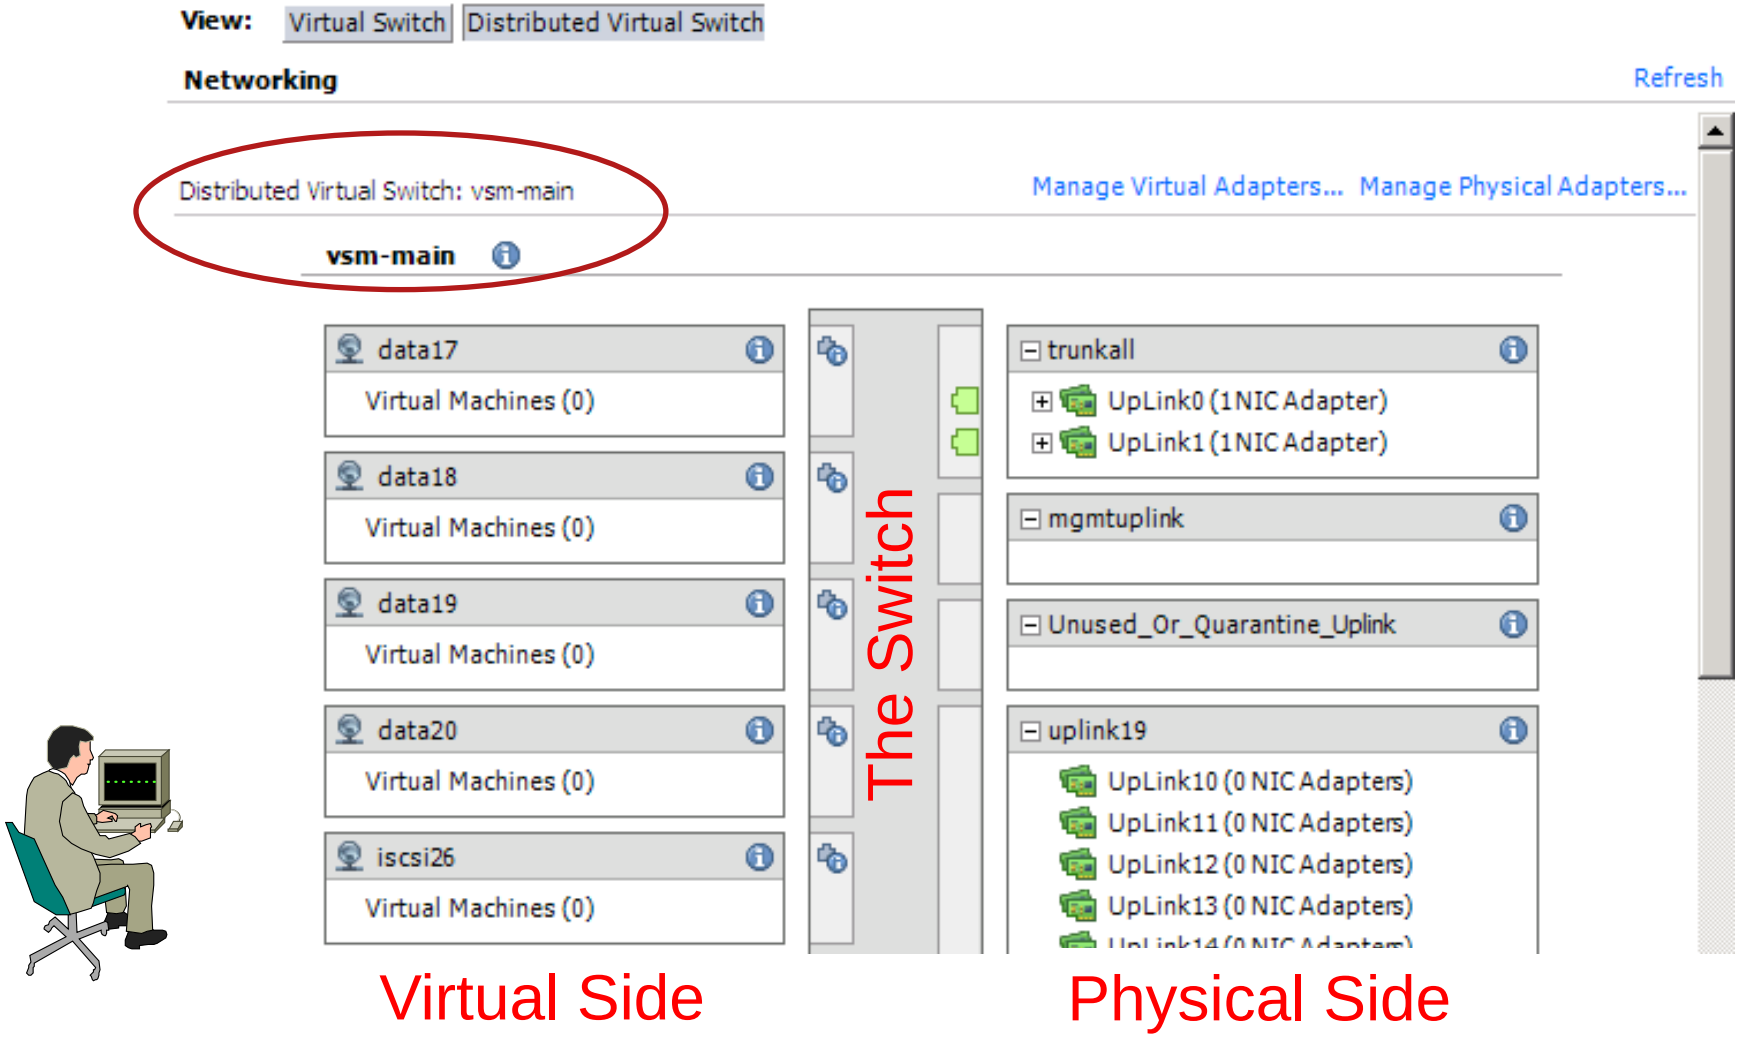

Virtual Side

Physical Side

Nexus 1000V on vCenter

Nexus 1000V is a Distributed Virtual Switch

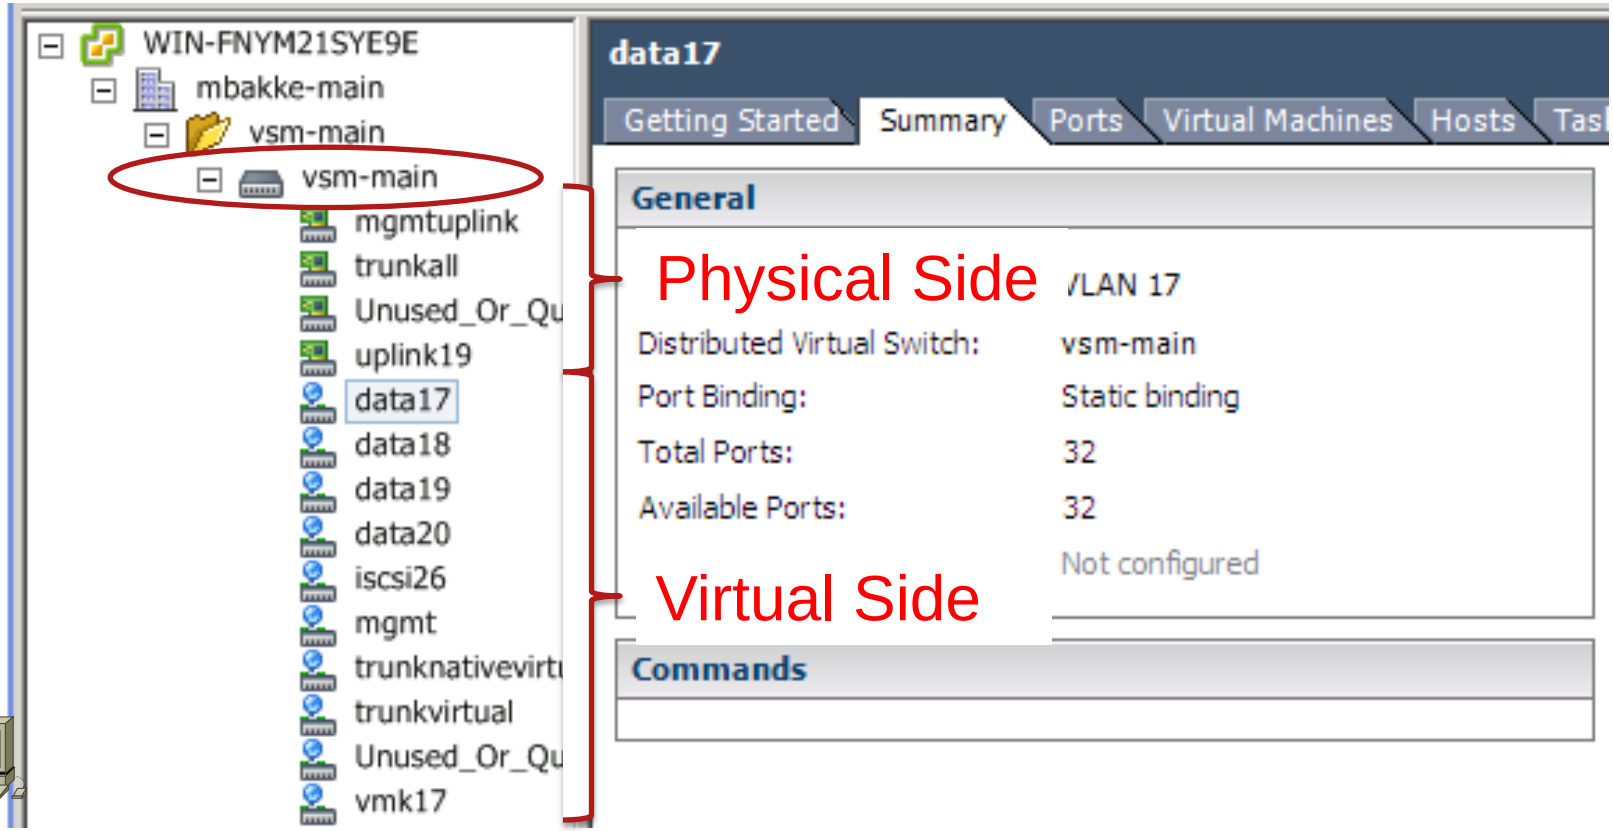

Nexus 1000V on the Host

Nexus 1000V is an ESX Host Software Package

```
esx-host # esxupdate --vib-view query
--------------------VIB ID------------------- Package State
cross_cisco-vem-v130-esx_4.2.1.1.4.0.0-2.0.1  installed
```

VEM Software Package

```
esx-host # ps | grep vemdpa
33959 33959 vemdpa
```

Data Path Agent
Communicates with VSM

Hypervisor “Drivers” – Packet Switching

```
esx-host # vmkload_mod -b
Name                 Size        Used
vmkernel             2713065     51
vem-v120-l2device    24576       5
vem-v120-n1kv        77824       3
vem-v120-vssnet      14901248    3
vem-v120-stun        90112       1
```

The Nexus 1000V VSM

Nexus 1000V is a Cisco Nexus Switch

```
vsm-main# show module
Mod  Ports  Module-Type                    Model          Status
---  -----  -----------------------------  -----------  ------
1    0      Virtual Supervisor Module      Nexus1000V   active *
3    248    Virtual Ethernet Module        NA           ok

Mod  Sw              Hw
---  --------------  ------
1    4.2(1)SV1(4)    0.0
3    4.2(1)SV1(4)    VMware ESXi 4.1.0 Releasebuild-260247 (2.0)

…

Mod  Server-IP        Server-UUID                Server-Name
---  ---------------  -------------------------  -----------
1    192.168.16.7     NA                         NA
3    192.168.16.15    bfd55a52-d564-11de-...     mbakke-ucs...
```

A single CLI for all VEMs

Virtual Machines and Virtual Ports

Network Customer Requirements

Owned by the Server / Application Administrator

Network	Security	Performance	Monitoring
DB / File Server	Back-end Network	Very High	ERSPAN for debug
DMZ	Separate Access		ERSPAN for recording
Web / Mail Server	Precise ACLs	Unpredictable	
Call Manager / UC	ACLs	Predictable	
VDI Clients	PVLANs Port Security		
Management	Separate from Others	Moderate	

Different Requirements for Different Networks

Network Implementation

Owned by the Network Admininistrator

- Network admin accepts requirements
- Chooses VLANs
- Chooses security policies—ACL, PVLAN, etc
- Chooses monitoring policies—NetFlow, ERSPAN,
- Chooses quality of service policies
- Chooses how physical NICs are used

Why Not Configure Virtual Ports?

- Too many ports, and they move too fast
- Configuring expanding, moving targets
- Network admin needs **sanity**
- Server admin needs **freedom**
  - To deploy and move virtual machines
  - To deploy and move physical hosts

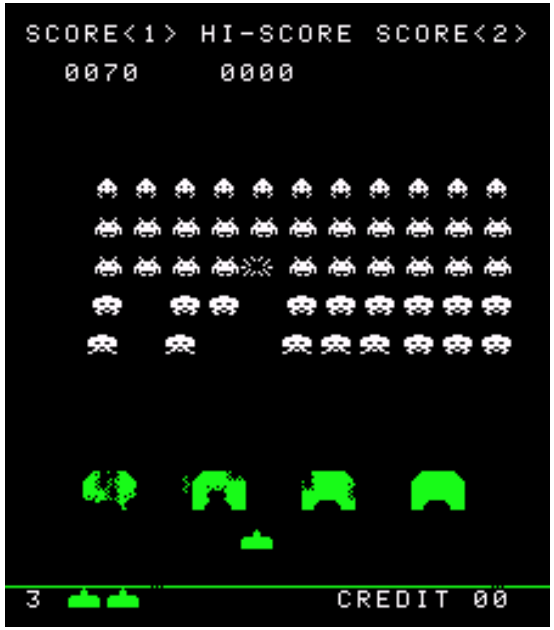

Source: http://images.webmagic.com/klov.com/screens/S/wSpace_Invaders.png

Port Profiles

Doing Your Homework…

- Instead of individual Ports
- Configure a Port Profile
- Set up ahead of time:
  - VLANs
  - ACLs
  - NetFlow
  - QoS
  - Private VLANs
  - and all other port config!
- Use this multiple times!

```
# port-profile database
switchport mode access
switchport access vlan 10
ip port access-group myacl in
no shut
state enabled
```

Port Profiles—Network View

Setting Port Policies Ahead of Time

```
# port-profile database
switchport mode access
switchport access vlan 10
no shut
```

```
# port-profile webserver
switchport mode access
switchport access vlan 752
access list, etc. commands
no shut
```

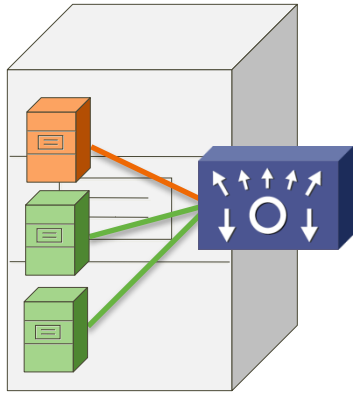

Port Profiles are "Live":

Network Admin can change them any time!

Port Profiles—Server View

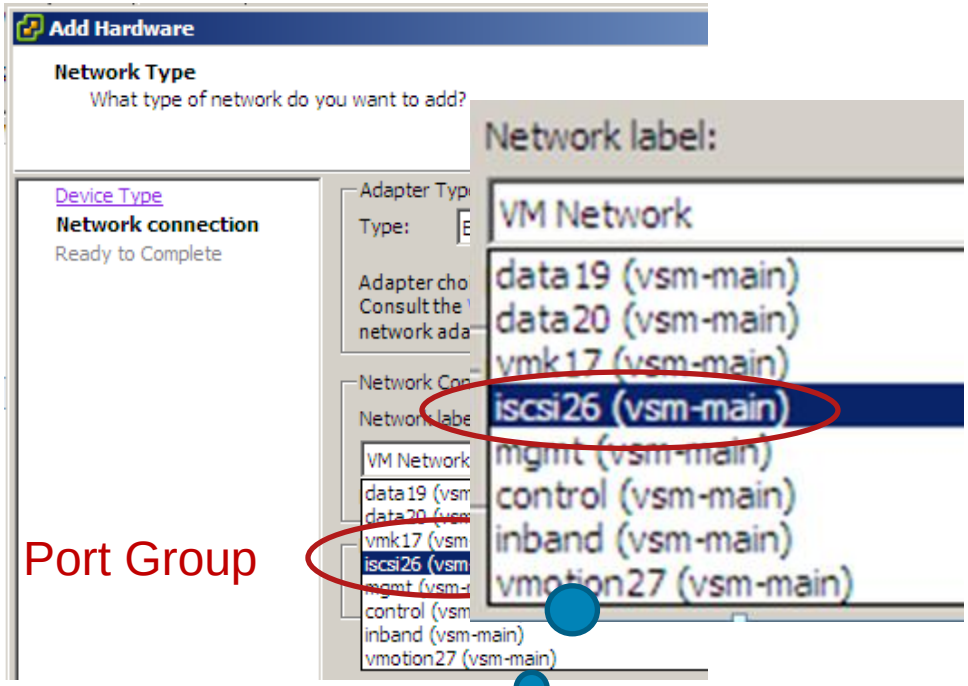

Port Group

Port Profile

```
# port-profile iscsi26
switchport mode access
switchport access vlan 26
Access lists, etc here
state enable
```

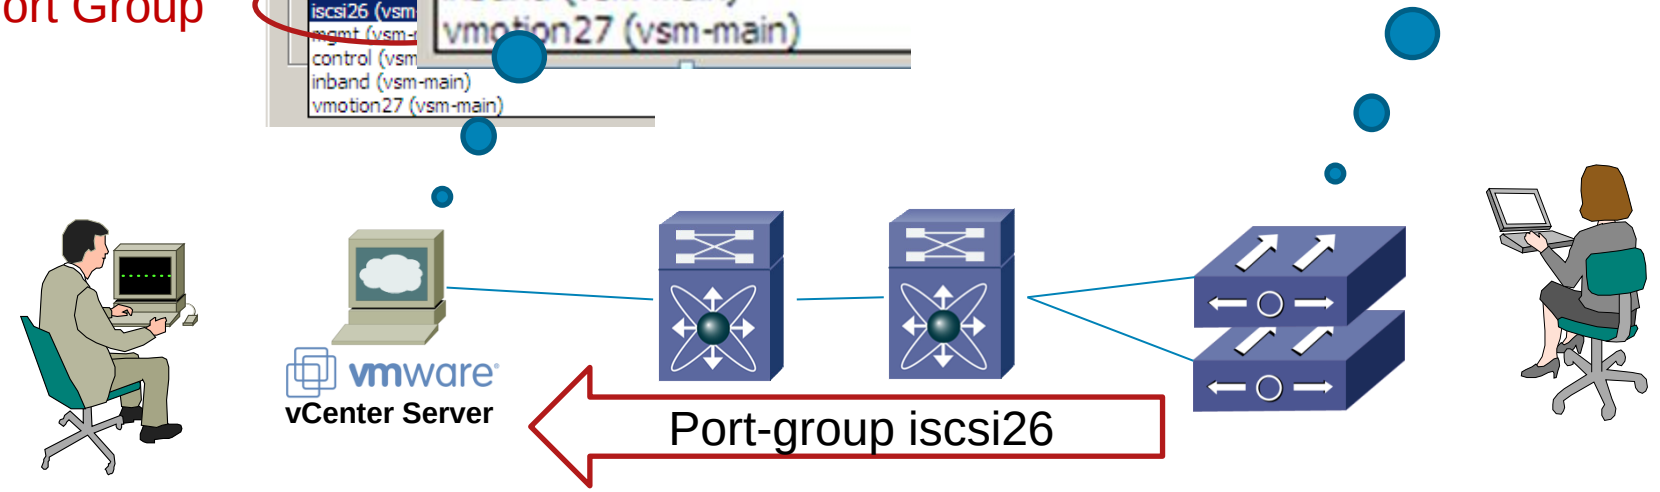

A Physical Analogy for Port Profiles

A Named Bundle of Wires

General

Description:	VLAN 26
Distributed Virtual Switch:	vsm-main
Port Binding:	Static binding
Total Ports:	8
Available Ports:	8
IP Pool:	Not configured

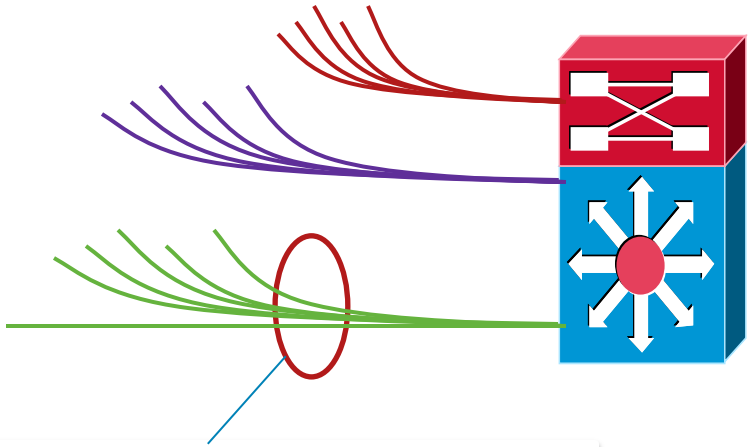

```
# port-profile iscsi26
switchport mode access
switchport access vlan 26
vmware max-ports 8
…
no shut
State enable
```

Plugging in the Virtual Wire

Inside the Host…

- Server admin chooses the port group
  - Like plugging the virtual wire from the switch into the NIC

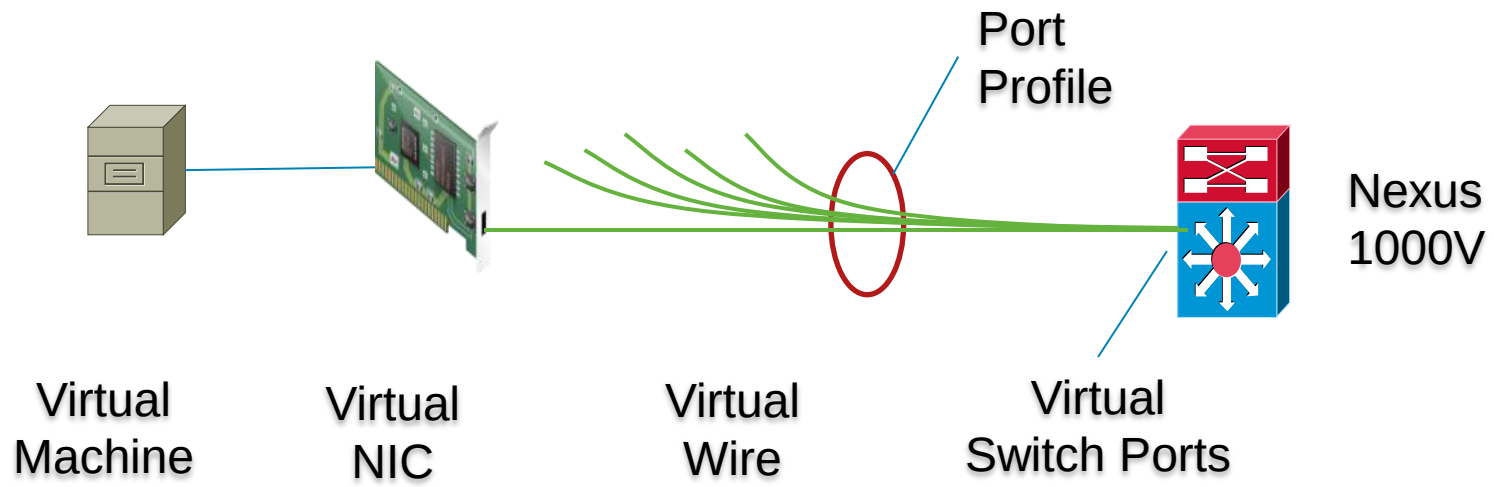

Server Admin's Virtual Port

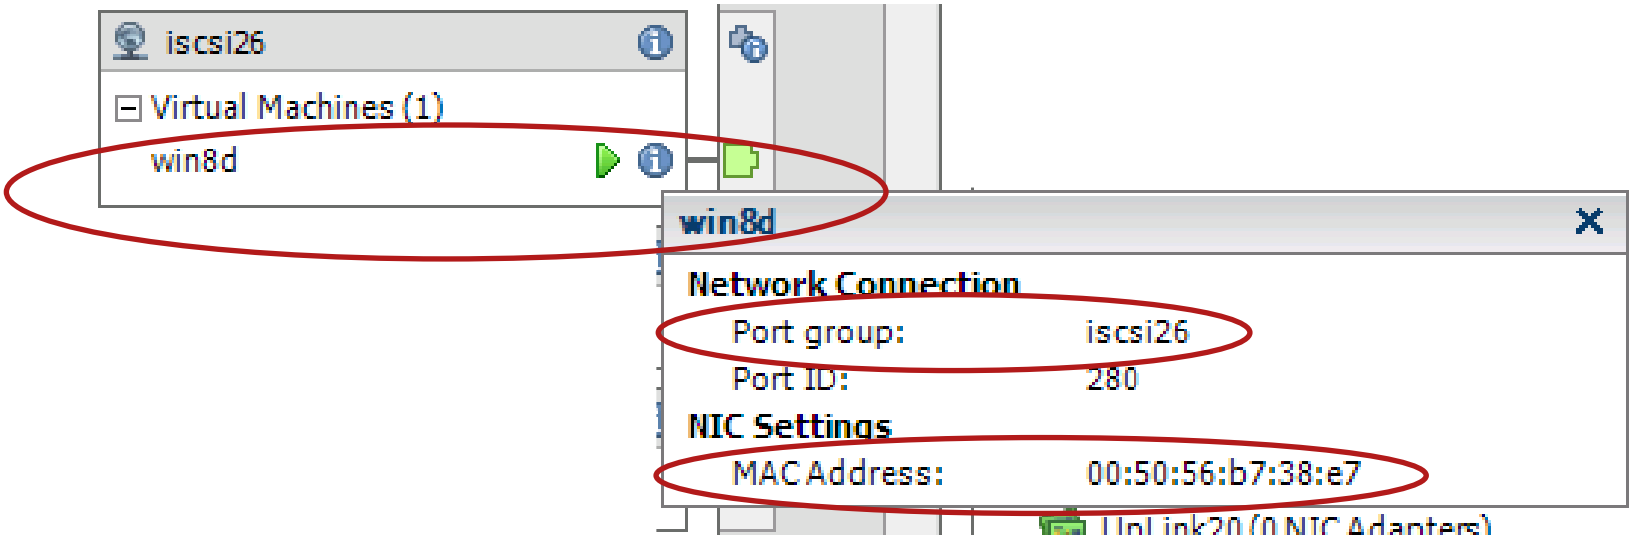

- Virtual machine name
- Adapter MAC address
- Port profile (port group)

Network Admin's Virtual Port

The Other End of the Wire

```
vsm-main# show int veth2
Vethernet2 is up
    Port description is win8d, Network Adapter 4
    Hardware is Virtual, address is 0050.56b7.38e7
    Owner is VM "win8d", adapter is Network Adapter 4
    Active on module 3
    VMware DVS port 280
    Port-Profile is iscsi26
    Port mode is access
 …
```

- Virtual machine name
- Adapter number and MAC address
- Host (module number) VM is running on
- Port profile

Deploying Large Numbers of VMs

- Virtual machine templates
  - Many duplicates from one master VM
  - Templates have port groups!
- Cloning virtual machines
  - Clones also bring port groups along!
- New VMs have own identity, MACs, etc.

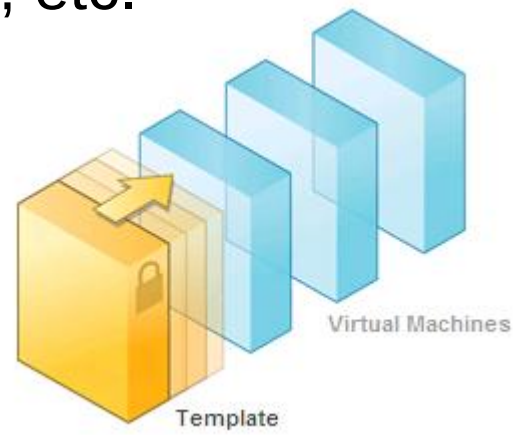

Review—Deploying Virtual Machines

- Network admin sets up port profiles in advance
  - All features are specified that will be needed
  - Goes to get coffee or on vacation
- Server admin creates VM templates
  - Template virtual NICs use port profiles
- Server admin clones templates
- Server admin starts up VMs
- Nexus 1000V sets up ports from port profiles
- No action from the Network Admin!

Possibly Thousands of VMs!

Back to the Fence

- Clear configuration boundaries
  - Server admin requests policies
  - Network admin creates port profiles
  - Server admin connects port profiles
- Transparent monitoring boundaries

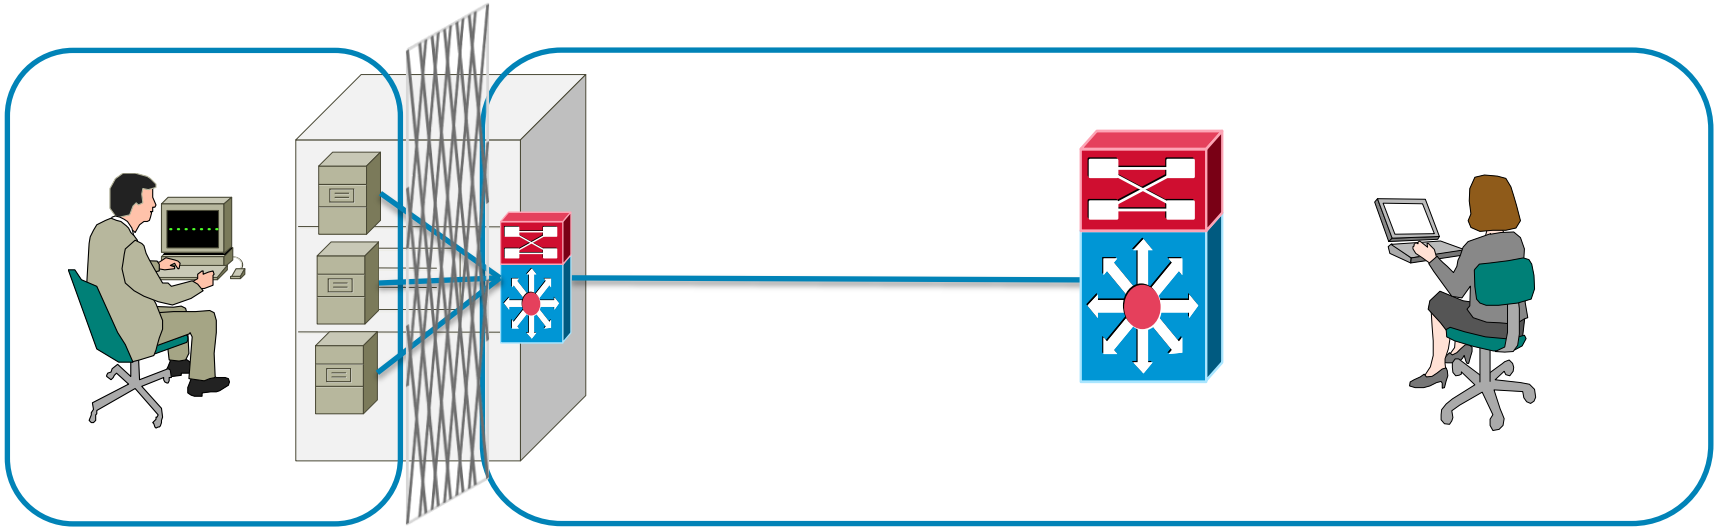

Moving Targets…

“It [The Nexus 1000V] also enables vMotion to work properly by giving server specialists the freedom to move hosts around without concerning themselves with Quality of Service and security settings.”

Rory Regan, Telecom Manager

Before the vMotion

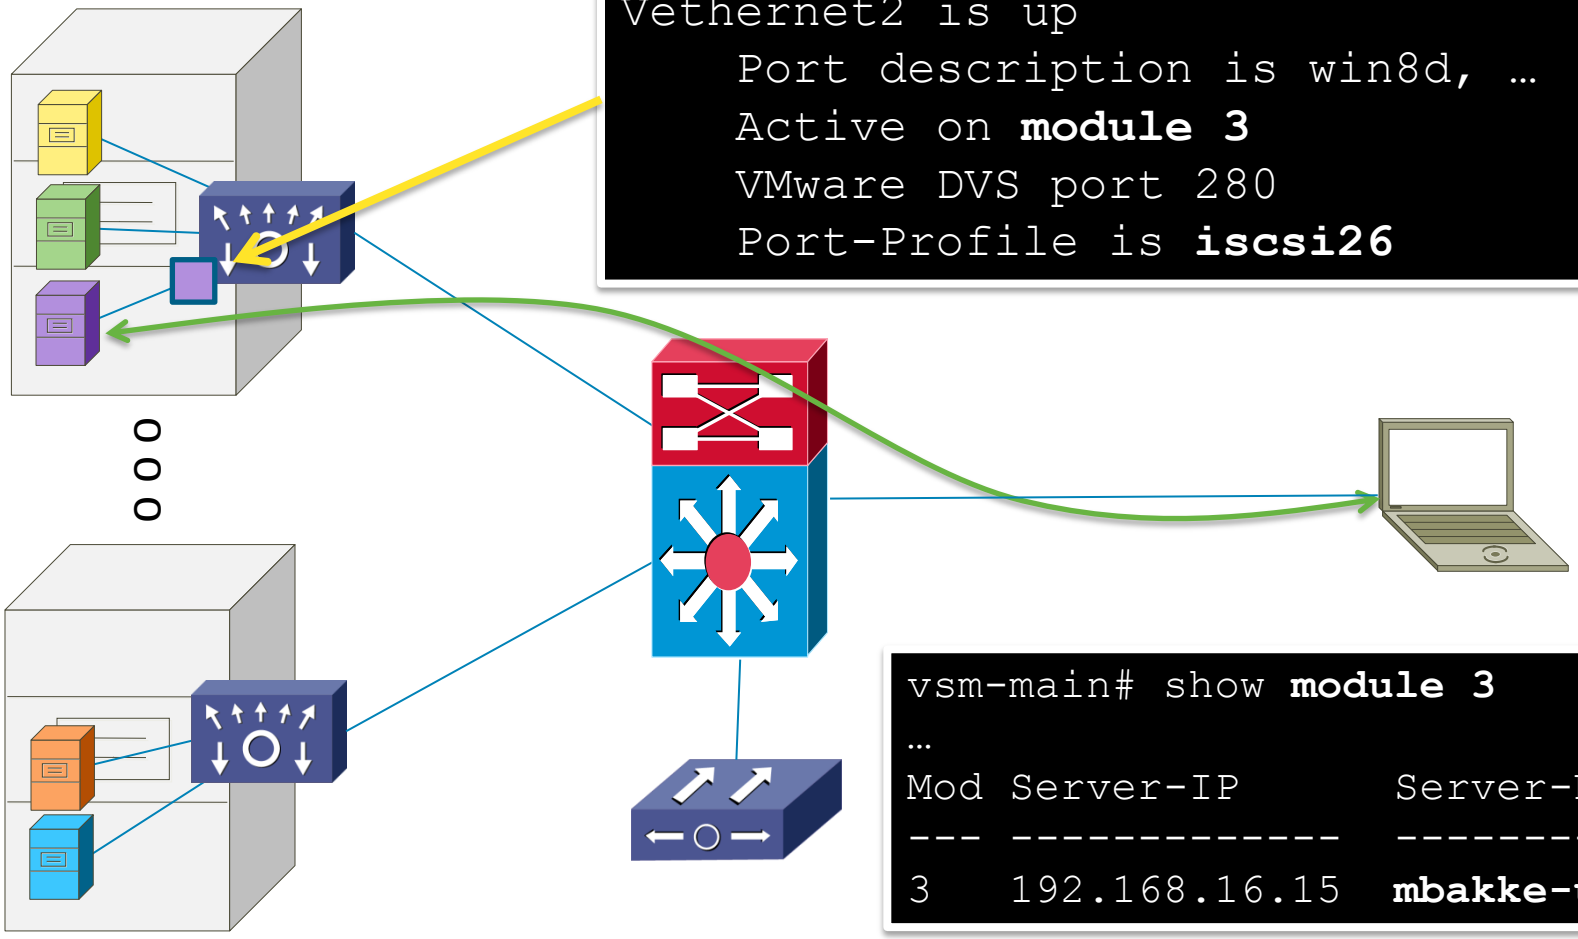

```
vsm-main# show int veth2
Vethernet2 is up
    Port description is win8d, …
    Active on module 3
    VMware DVS port 280
    Port-Profile is iscsi26
```

```
vsm-main# show module 3
…
Mod Server-IP       Server-Name
--- --------------  -----------
3   192.168.16.15   mbakke-ucs
```

After the vMotion

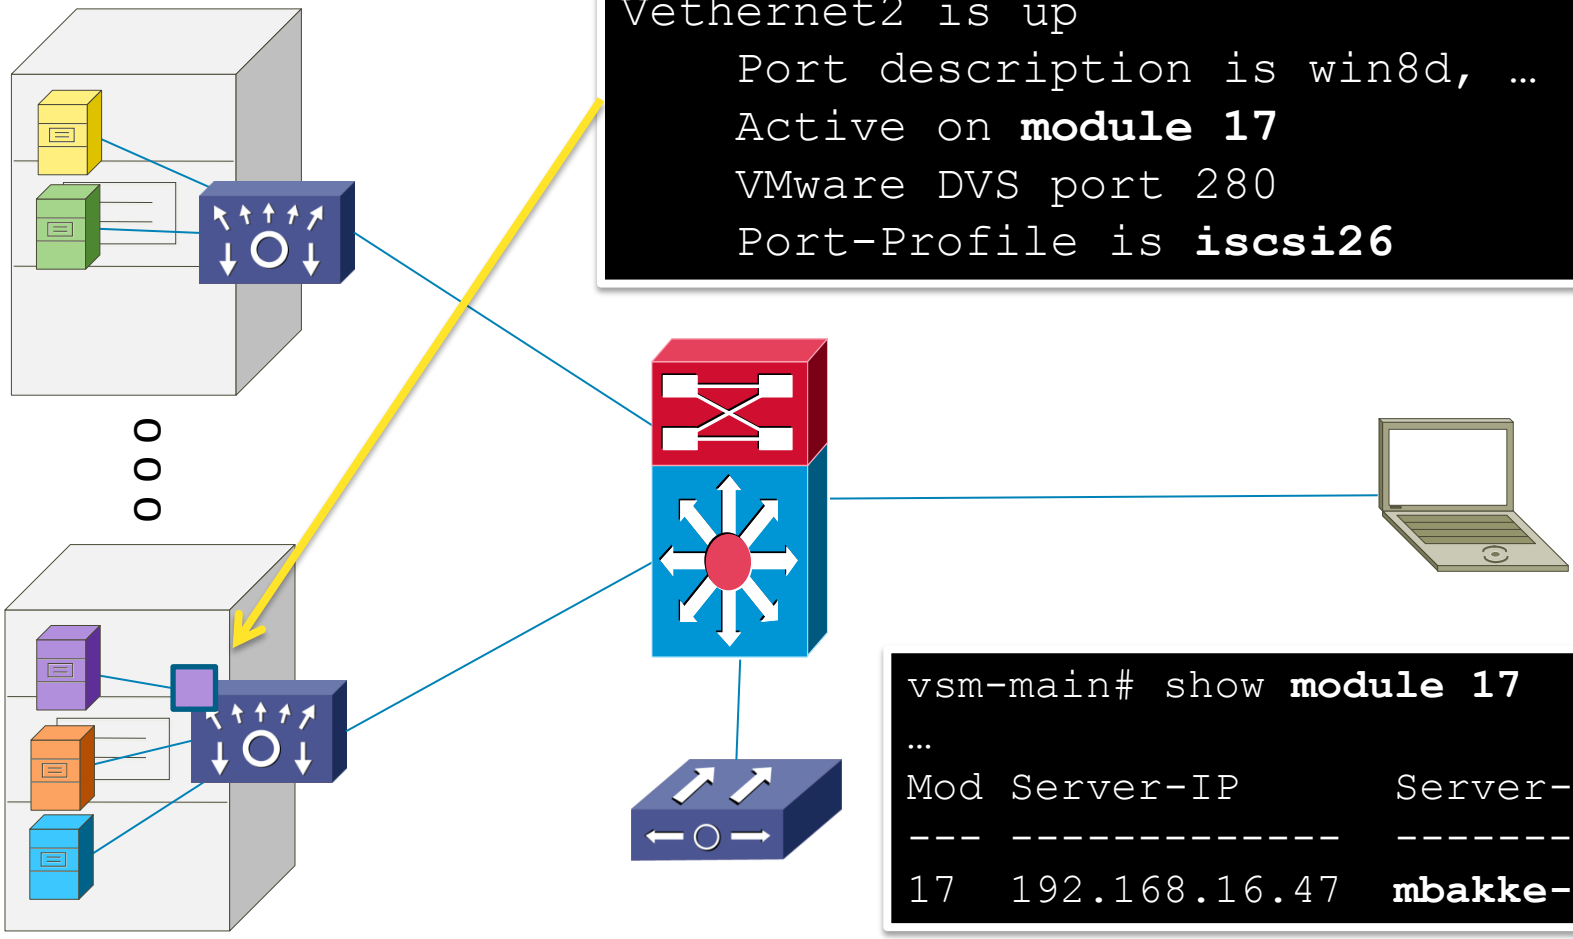

```
vsm-main# show int veth2
Vethernet2 is up
    Port description is win8d, …
    Active on module 17
    VMware DVS port 280
    Port-Profile is iscsi26
```

```
vsm-main# show module 17
…
Mod Server-IP       Server-Name
--- --------------  -----------
17  192.168.16.47   mbakke-svr2
```

Notifying the L2 Network

Learning the New Location of the VM

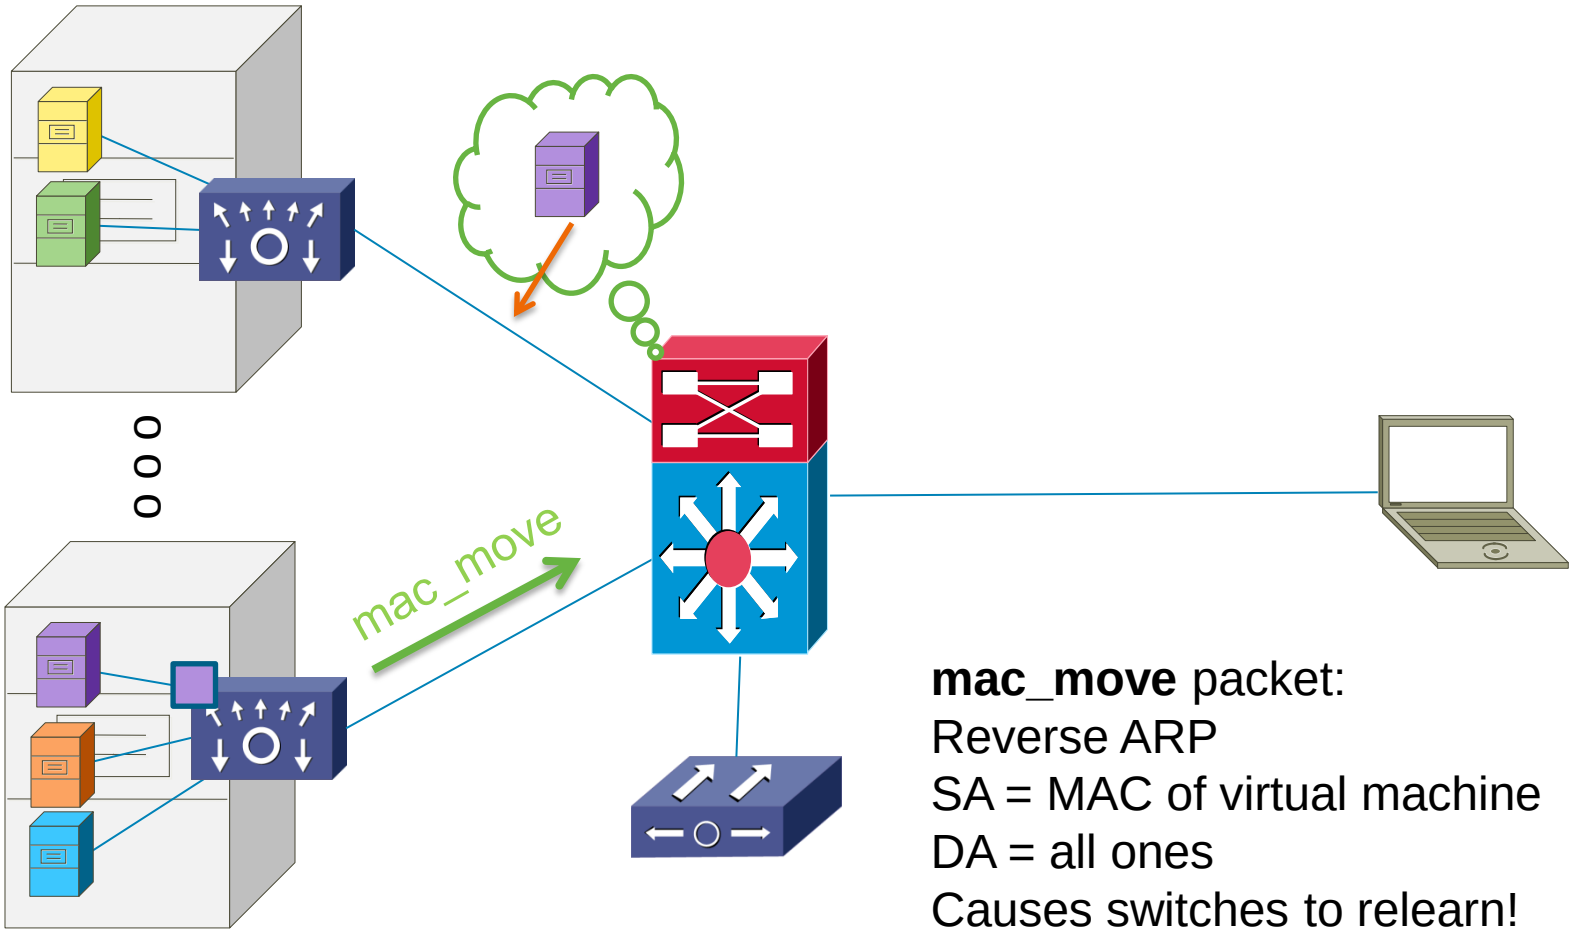

mac_move packet:
Reverse ARP
SA = MAC of virtual machine
DA = all ones
Causes switches to relearn!

Notifying the L2 Network

Learning the New Location of the VM

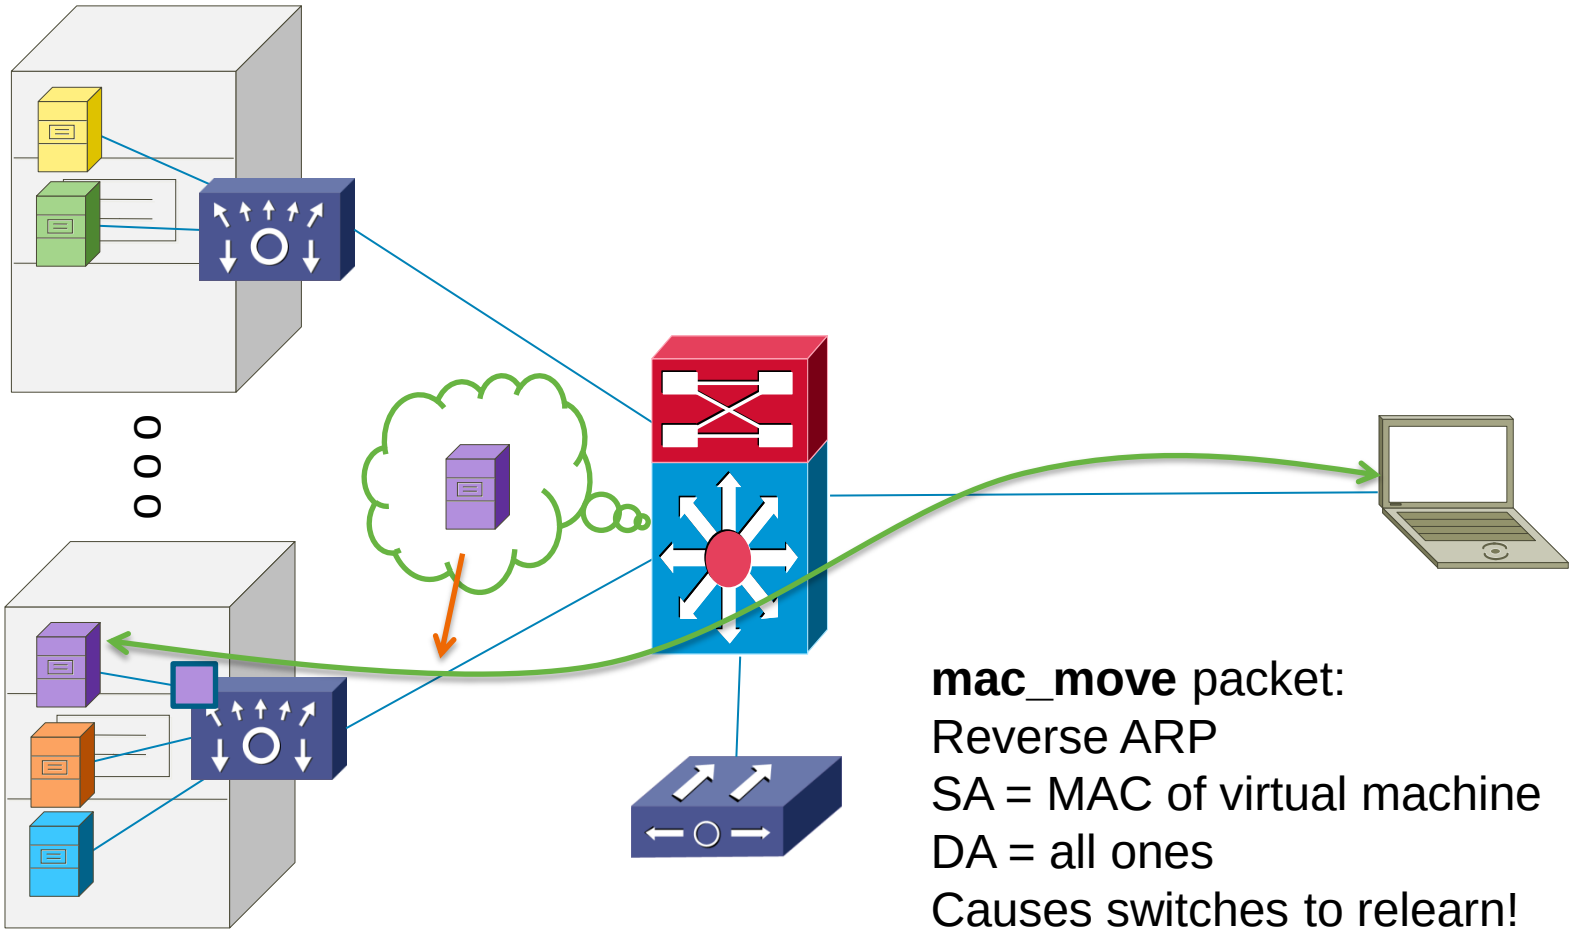

mac_move packet:
Reverse ARP
SA = MAC of virtual machine
DA = all ones
Causes switches to relearn!

VMware Clusters

A Group of Physical Hosts and Virtual Machines

Small clusters for failover (HA)

Larger clusters for load sharing (DRS)

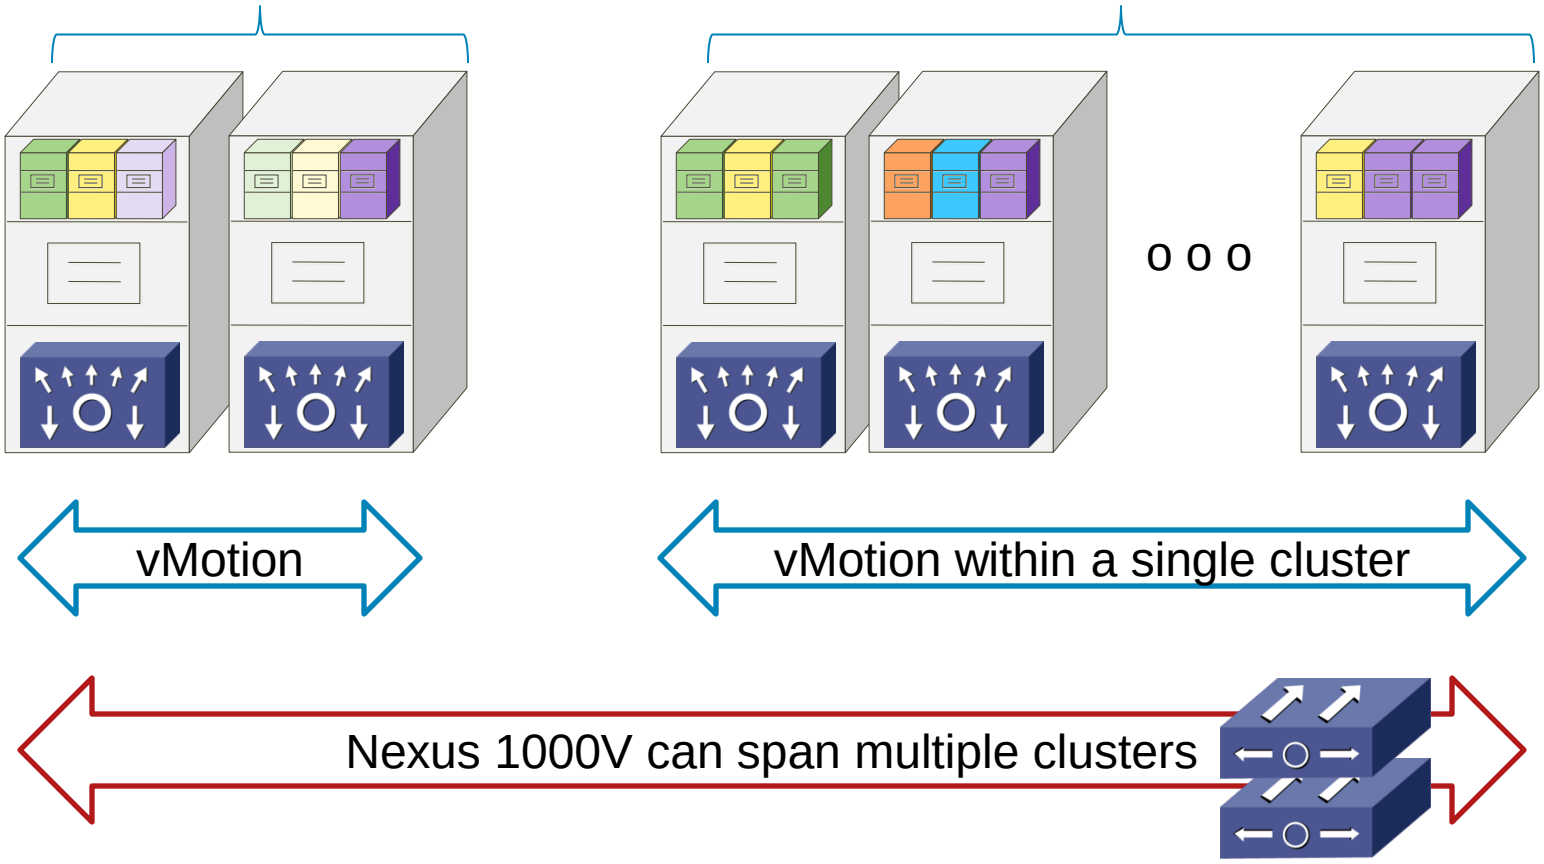

High Availability and Fault Tolerance

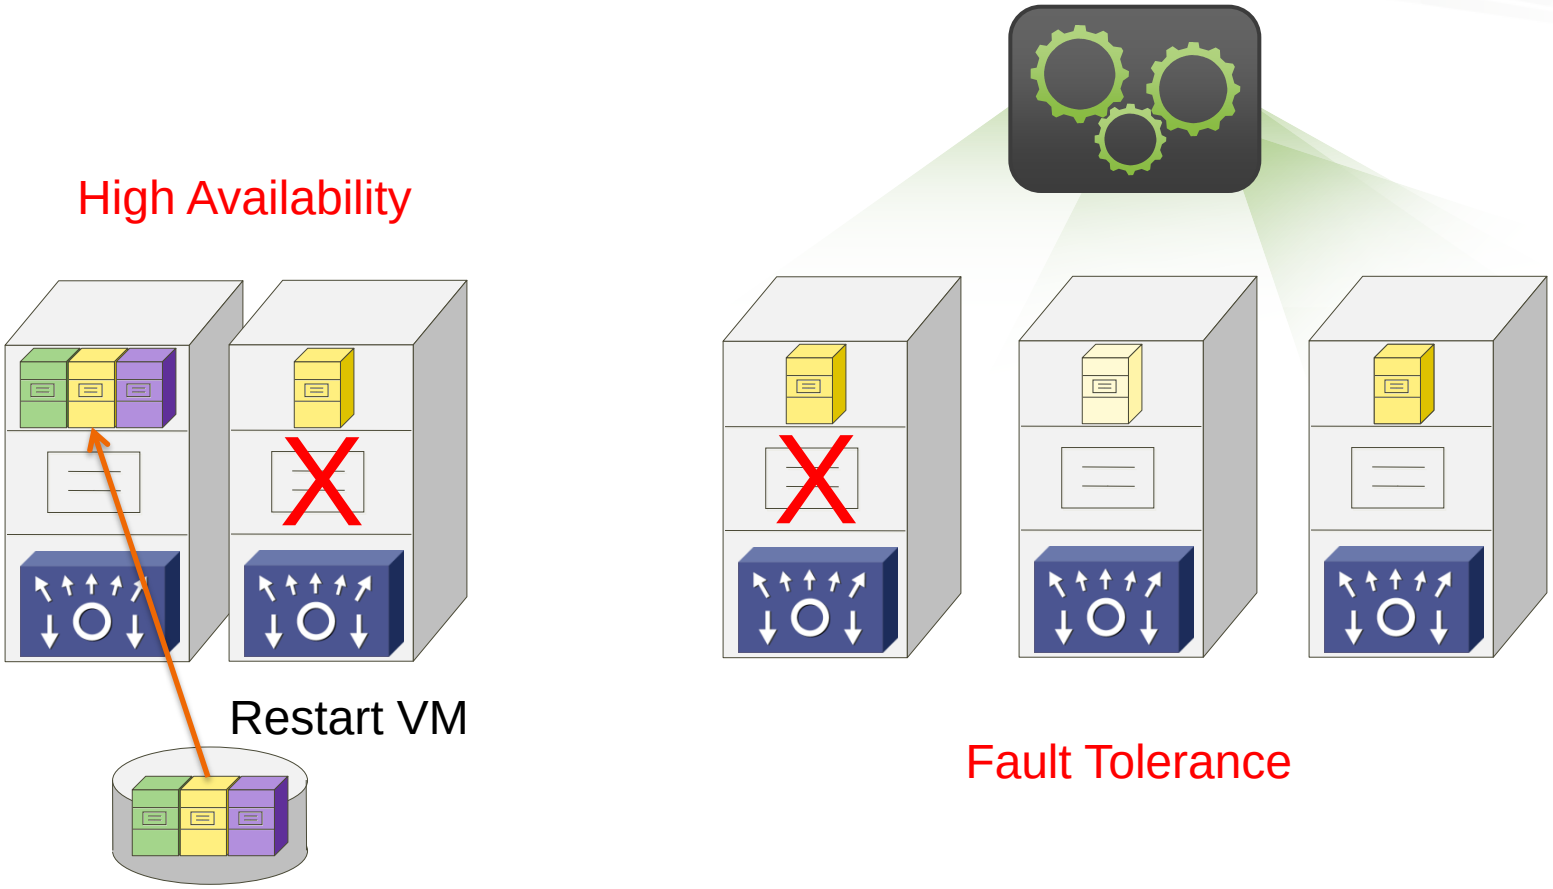

- Both are supported by Nexus 1000V

  Virtual ethernet moves with the virtual machine

The “Physical” Nexus 1000V

Physical and Virtual Switches

Supervisors and Line Cards

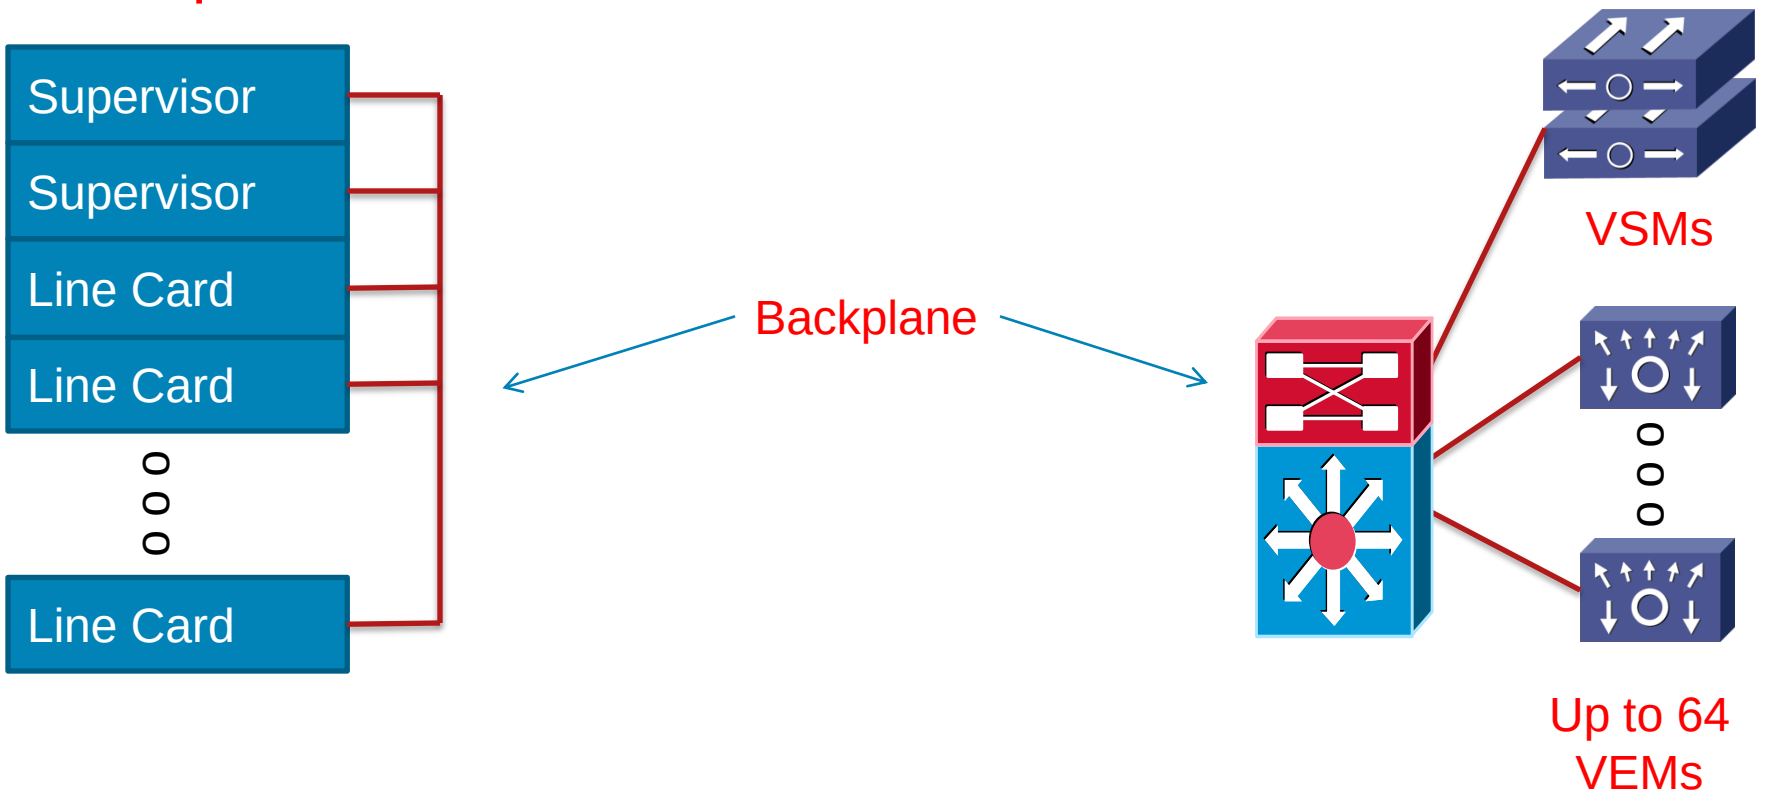

?

Physical and Virtual Line Cards

From Hardware to Software

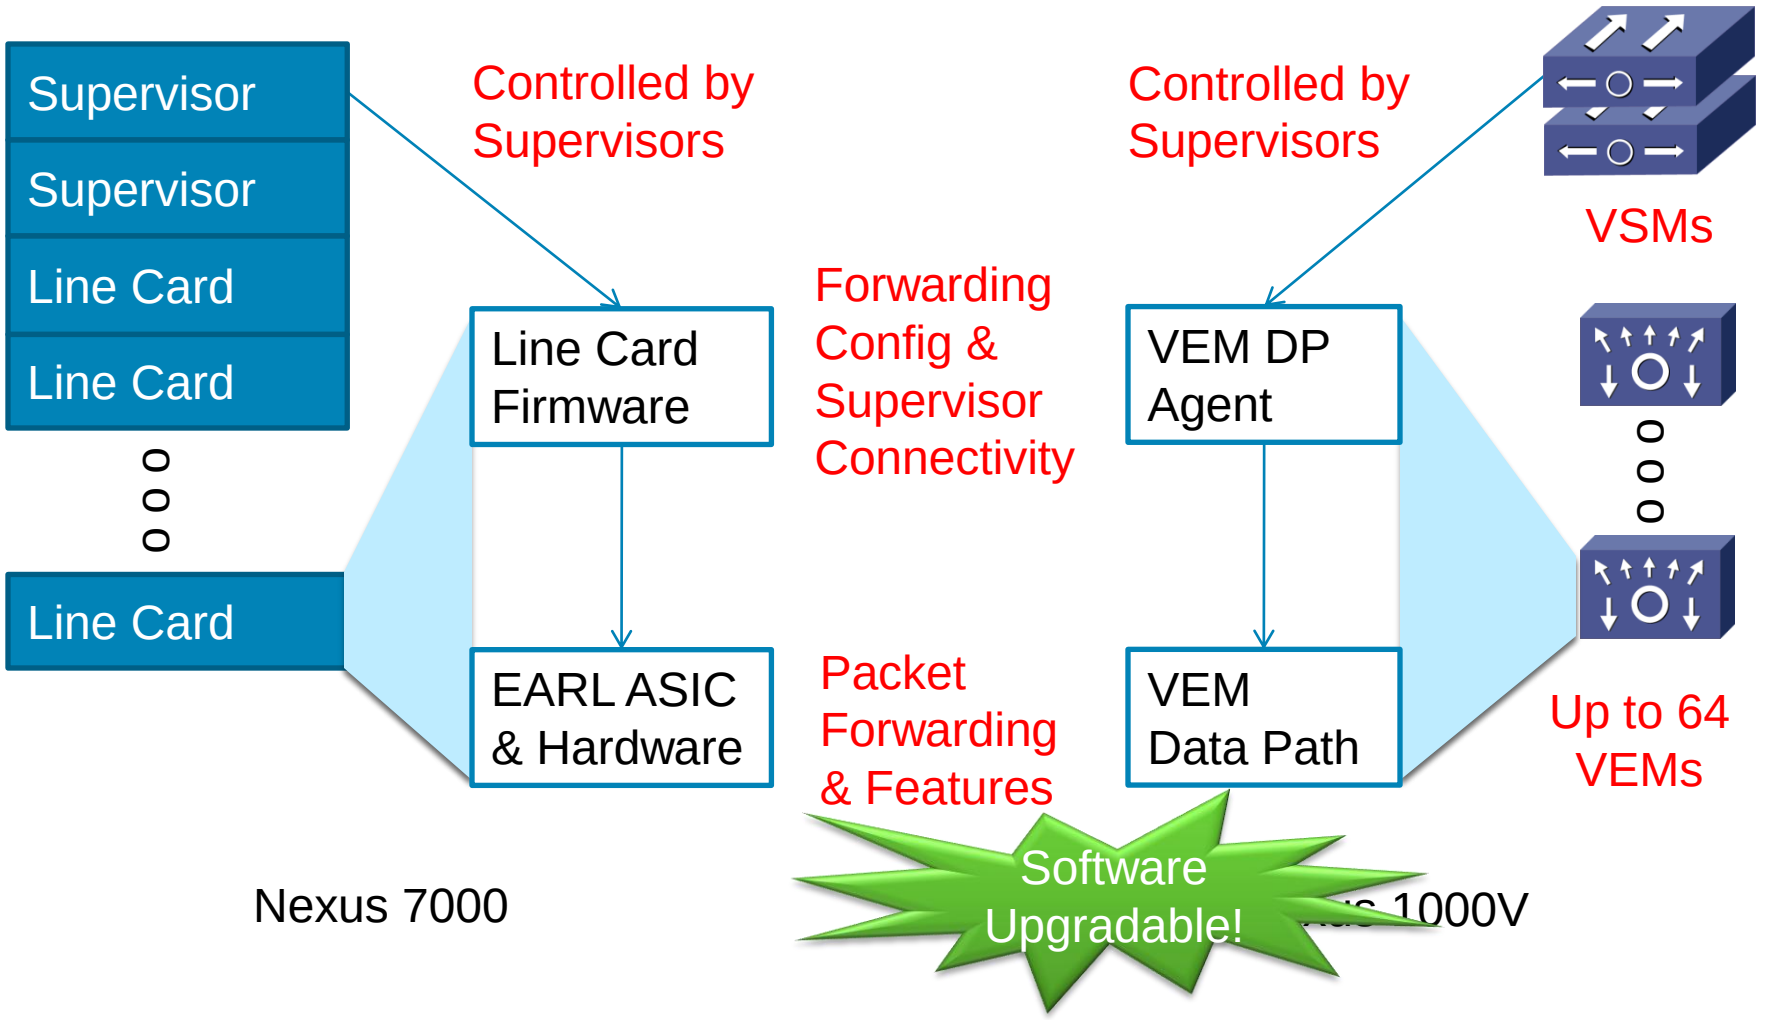

Switch in vCenter

VSM and VEMs Appear as One Switch

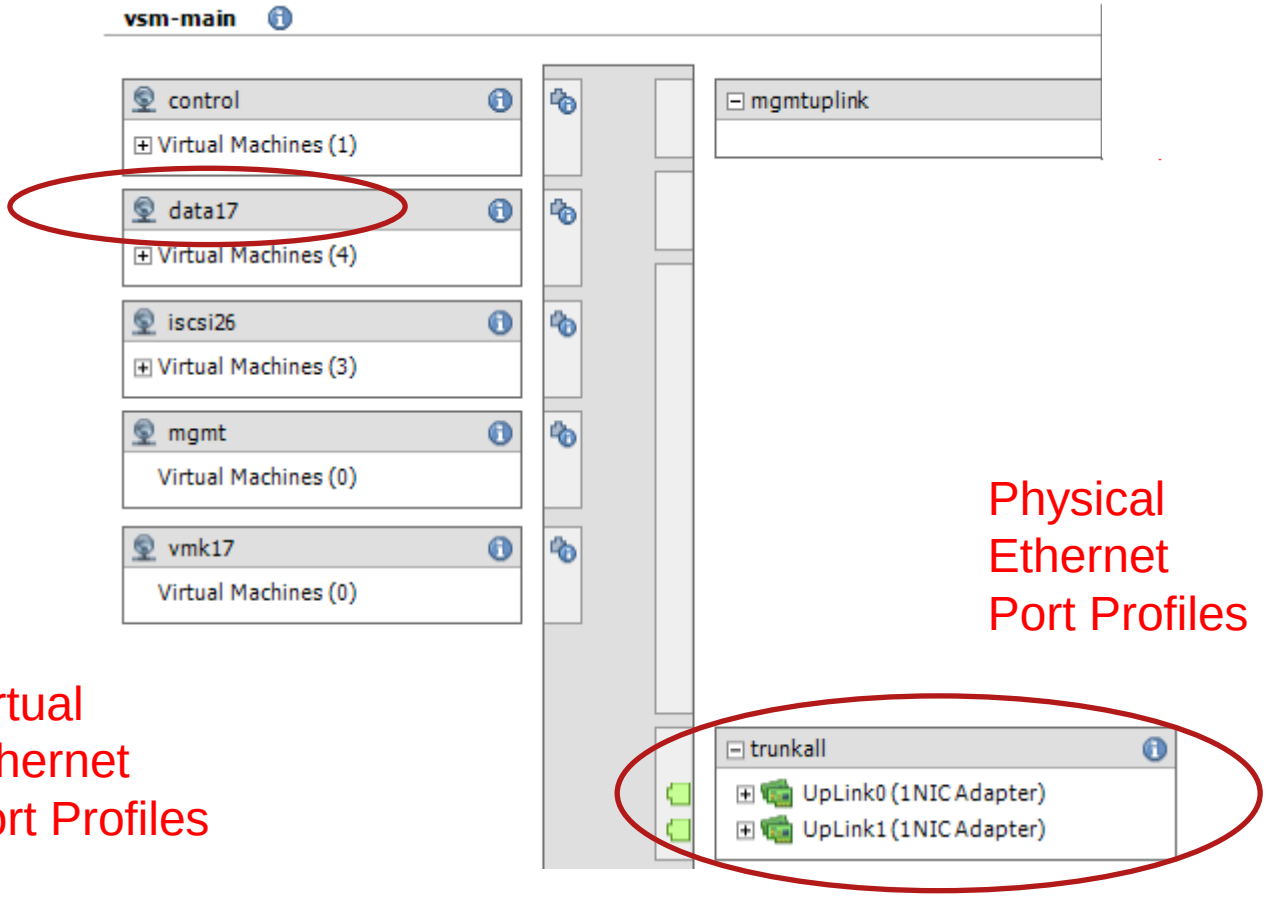

Physical Ethernet Port Profiles

Virtual Ethernet Port Profiles

Physical Hosts and the Nexus 1000V

Host Connectivity Requirements

Each Physical Host Is Typically on Several Networks

- Management—to talk to vCenter
- Storage—iSCSI and NFS
- VMotion—for moving VMs
- VSM to VEM communication—the “backplane”
- Virtual machine networks—(why we are all here)

- Port channels for physical NICs
  - Many configurations possible
  - From dual 10G to many 1G

?

Port Profiles for Physical Ports

"Uplink" or "Ethernet"

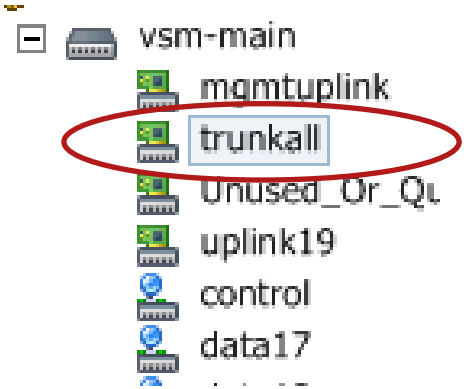


```
# port-profile type ethernet trunkall
switchport mode trunk
switchport trunk allowed vlan 16-30
channel-group auto mode on
no shutdown
```

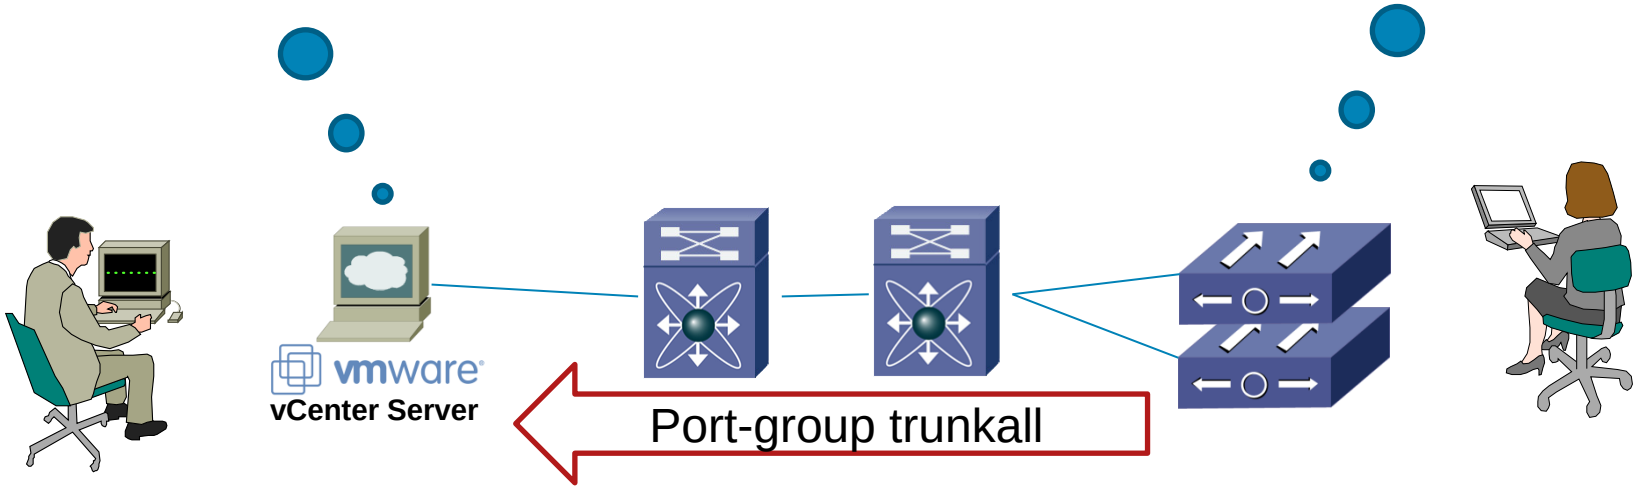

Port Profiles for Physical Ports

Adding a Host to the Nexus 1000V

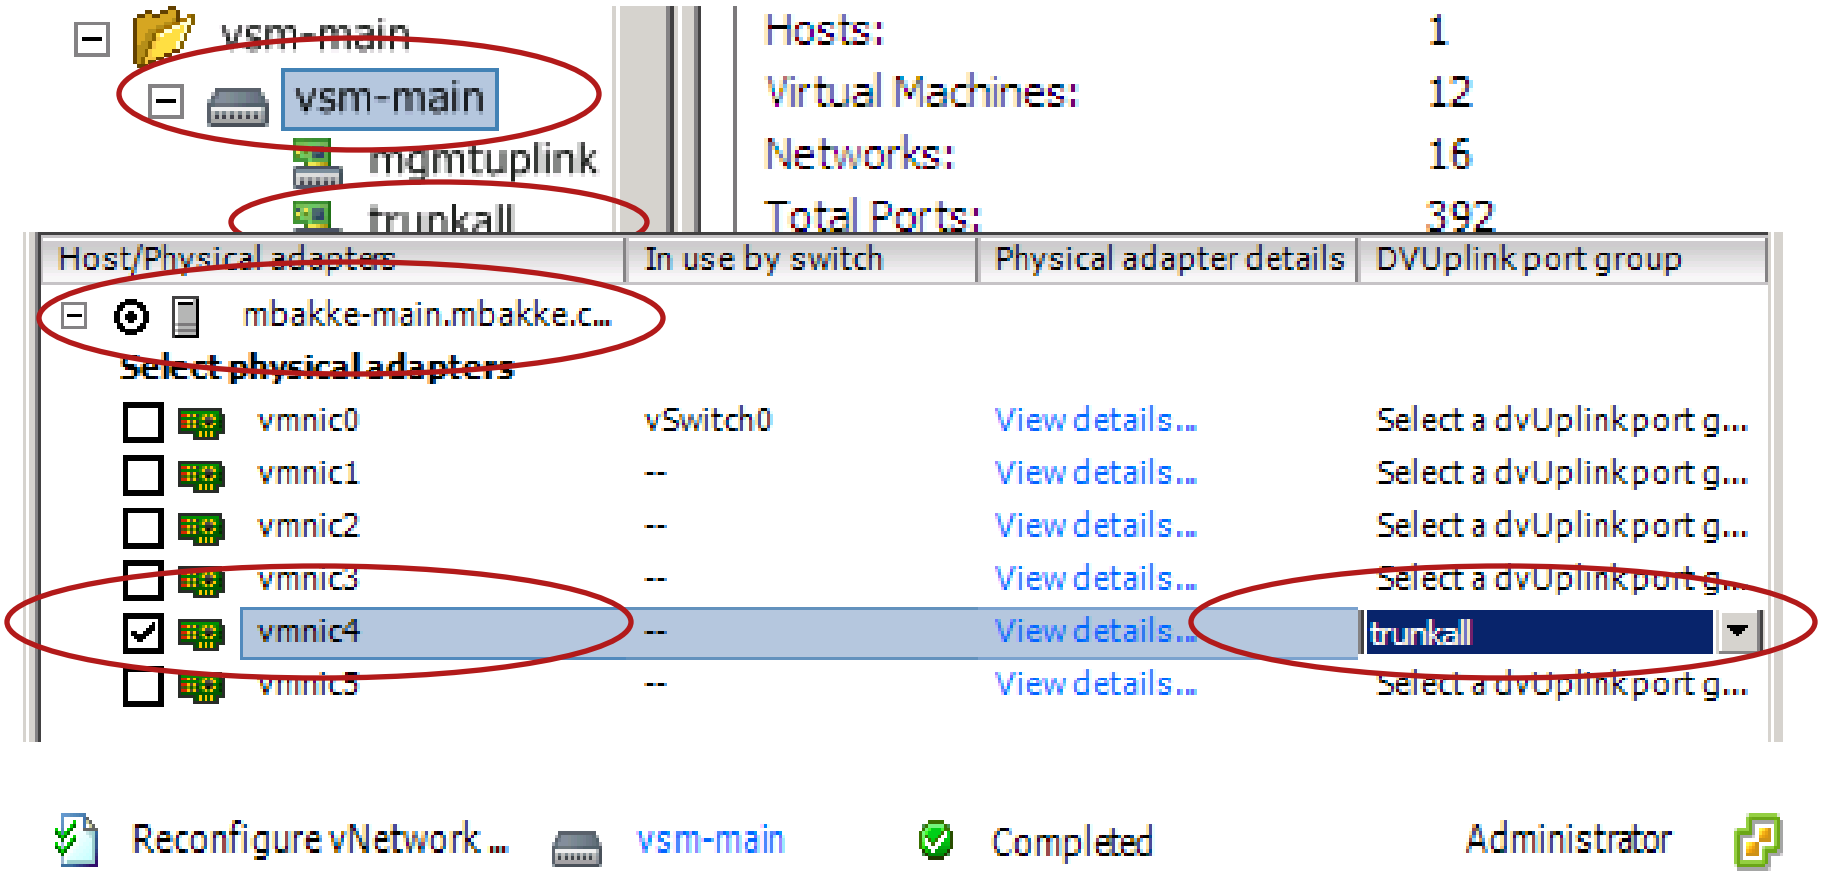


```
vsm-main# 2011 May 18 19:04:50 vsm-main …: Module 3 powered up
```

Port Profiles for Physical Ports

Automatic Port Channels

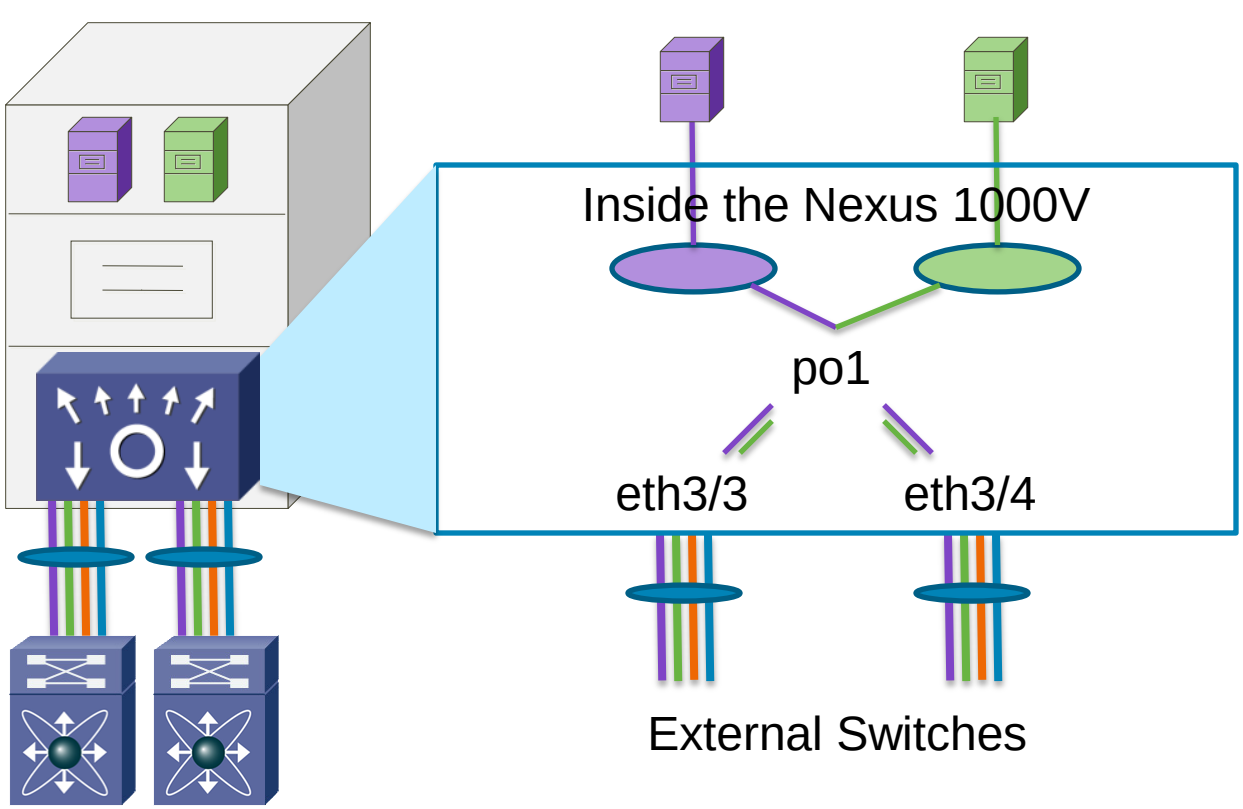

And provided to virtual machines

↑

Connected to internal VLANs

↑

Combined into a port channel

↑

Physical ports Carry VLANs

Server Admin View of Physical Ports

View: Virtual Switch | Distributed Virtual Switch

Networking — Refresh

vsm-main

- data17 — Virtual Machines (0)
- data18 — Virtual Machines (0)
- data19 — Virtual Machines (0)
- data20 — Virtual Machines (0)
- iscsi26 — Virtual Machines (0)
- mgmt — Virtual Machines (0)

trunkall

- UpLink0 (1NIC Adapter)
  - vmnic3 mbakke-ucs.mbakke.cisco.c...
- UpLink1 (1NIC Adapter)
  - vmnic2 mbakke-ucs.mbakke.cisco.c...

Cisco Discovery Protocol

Properties

Version	2
Timeout	0
Time to live	169
Samples	1251
Device ID	mbakke-router.mbakke.cisc
Address	192.168.17.254
Port ID	GigabitEthernet1/0/24
Software Version	Cisco IOS Software, C3750
Hardware Platform	cisco WS-C3750G-48TS
IP Prefix	0.0.0.0
IP Prefix Length	0
VLAN	1

Network Admin View of Physical Ports

Automatic Port Channels

```
vsm-main# show interface brief
…
--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason   Speed    Port
Interface                                                   Ch #
--------------------------------------------------------------------------------
Eth3/3        1       eth  trunk  up      none     1000(D)  1
Eth3/4        1       eth  trunk  up      none     1000(D)  1

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

```
vsm-main# show cdp neighbors interface eth3/3
…
mbakke-r… Eth3/3          173    R S I      WS-C3750G-48T Gig1/0/24
```

Hooking It Up—Virtual Ports on the Host

Management and Control Networks

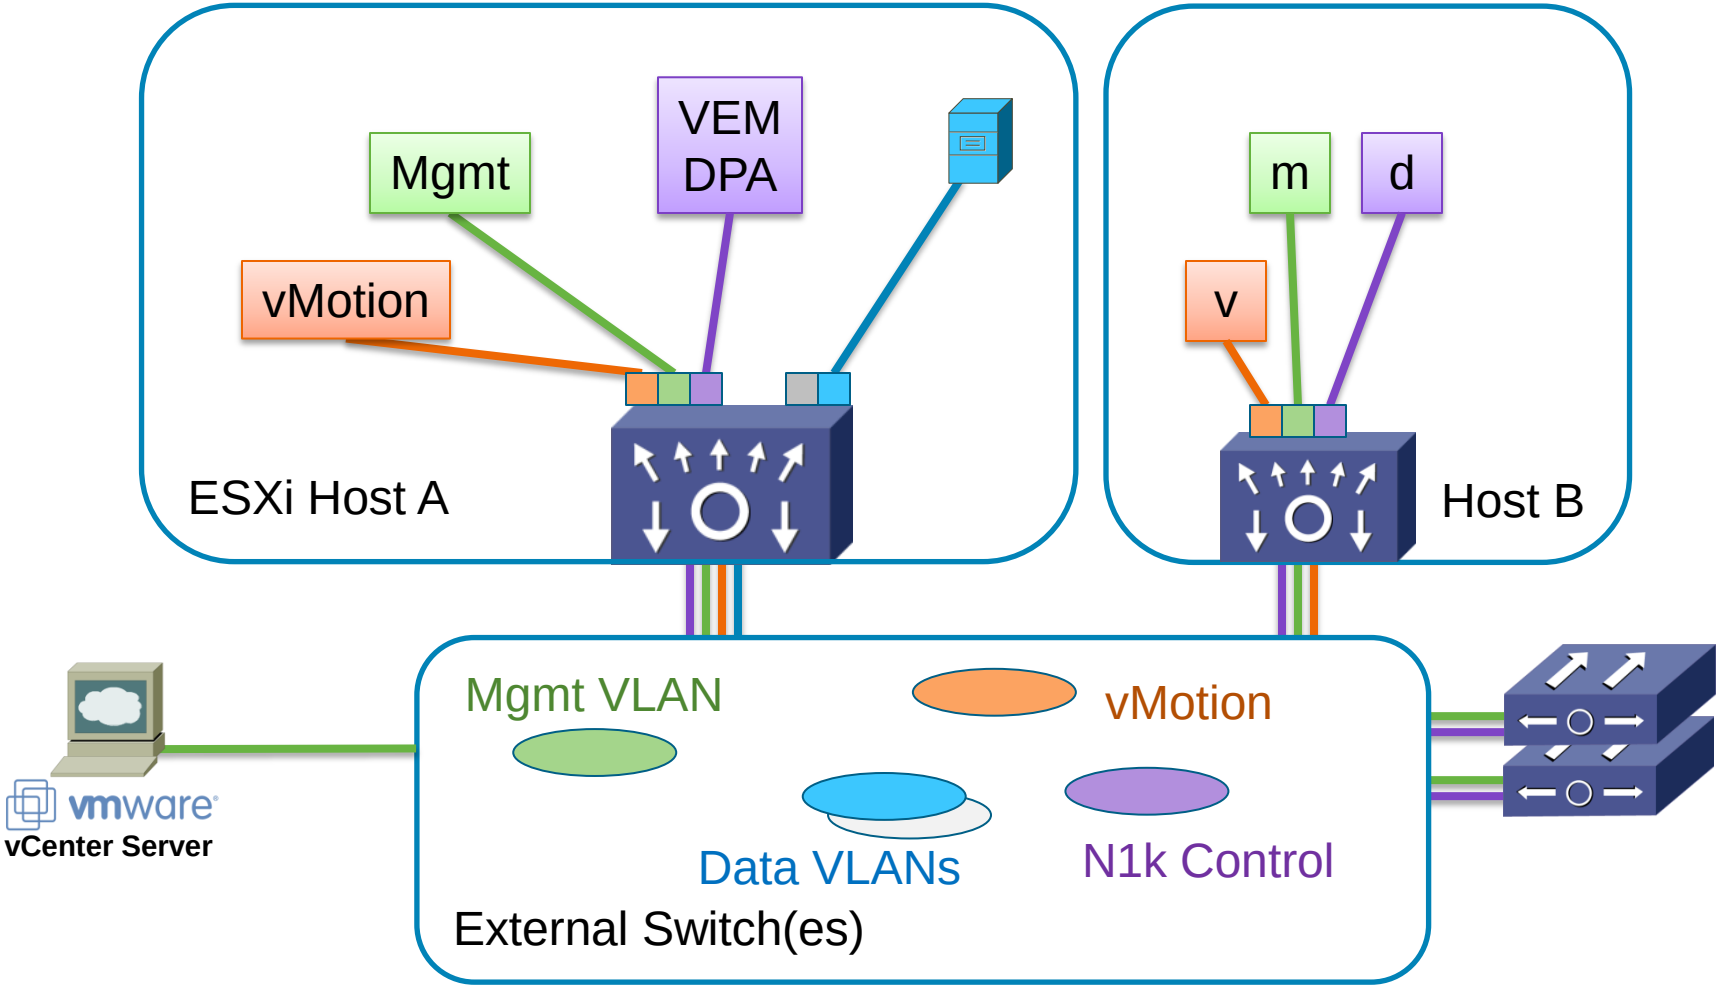

Hooking It Up—Network Storage

NFS and iSCSI Through the Nexus 1000V

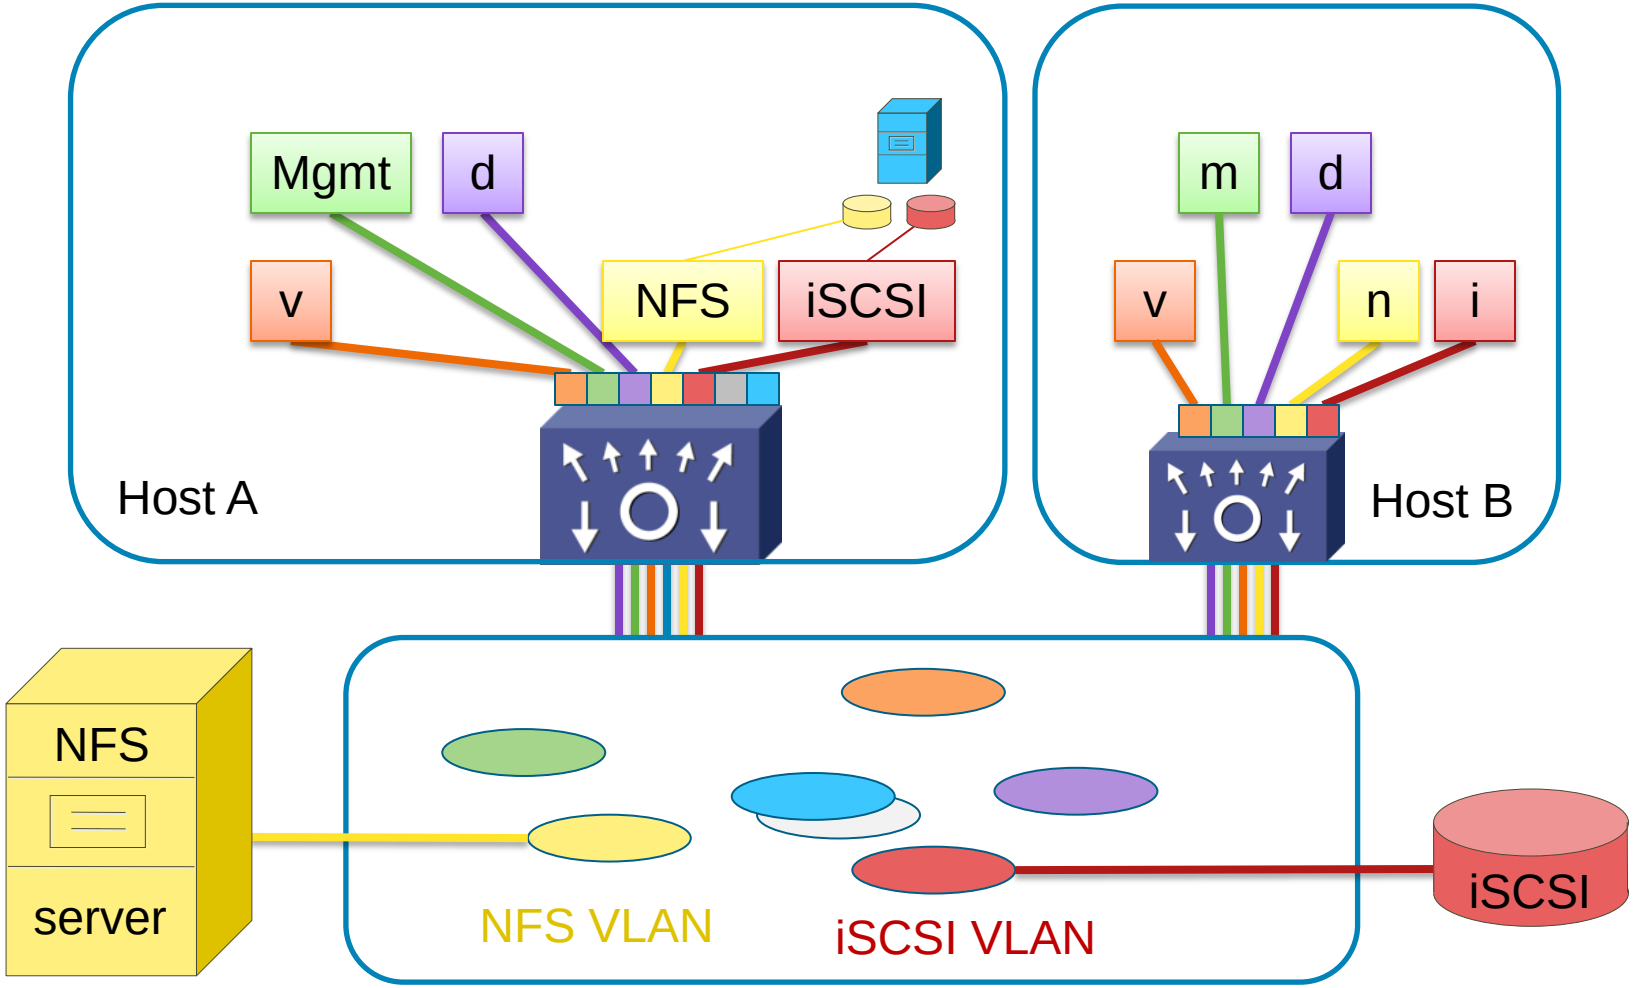

Port Profiles for the Host

```
# port-profile control-vmk
switchport mode access
switchport access vlan 30
system vlan 30
capability l3control
no shutdown
```

```
# port-profile management
switchport mode access
switchport access vlan 16
no shutdown
```

```
# port-profile vmotion
switchport mode access
switchport access vlan 27
no shutdown
```

```
# port-profile nfs
switchport mode access
switchport access vlan 28
no shutdown
```

```
# port-profile iscsi
switchport mode access
switchport access vlan 26
no shutdown
```

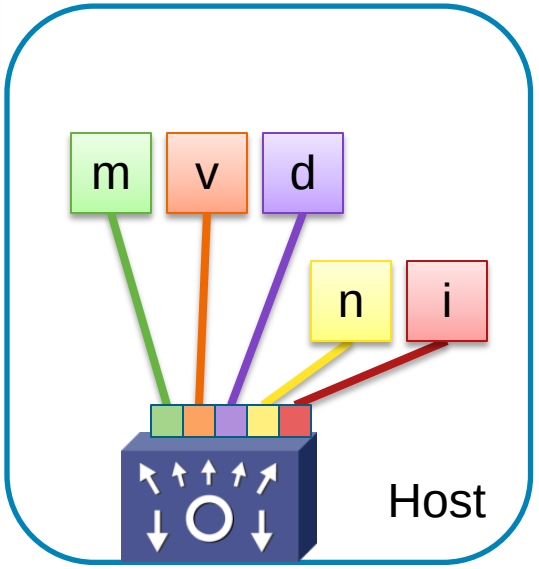

System VLANs for Bootstrapping

Systems VLANs must be used for:
- ESX Management (console) VLAN
  - Both uplink and vethernet
- Nexus 1000V Control VLAN

These VLANs are brought up on their ports before talking with the VSM

Systems VLANs may be used for:
- Storage VLANs (NFS/iSCSI)

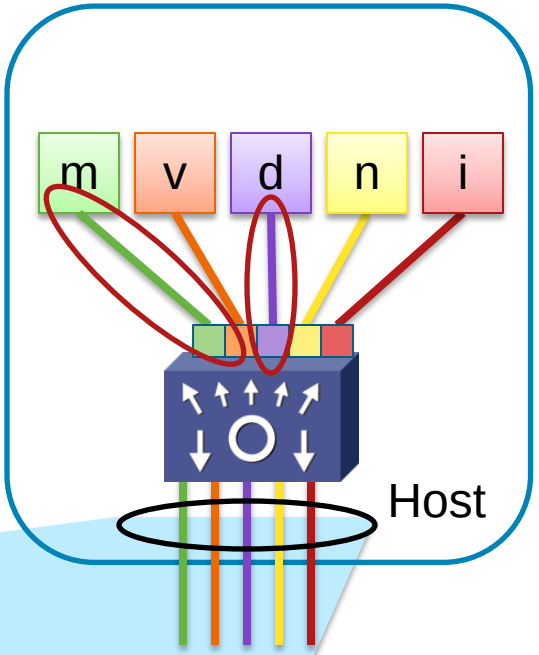


```
# port-profile type ethernet trunkall
switchport mode trunk
switchport trunk allowed-vlan 16,26-28,30
no shutdown
system vlan 16,30
```

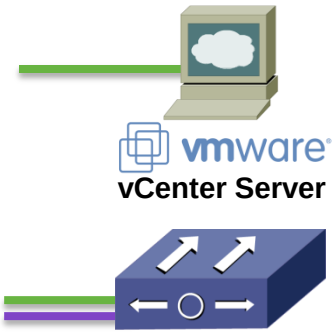

System VLANs for VSM on its own VEM

Systems VLANs must also be used in vethernet port profiles for:

- VSM Management (console) VLAN
- Nexus 1000V Control VLAN

```
# port-profile management
switchport mode access
switch
no shu
system
```

```
# port-profile vsm-control
switchport mode access
switchport access vlan 30
no shutdown
system vlan 30
```

```
# port-profile type ethernet trunkall
switchport mode trunk
switchport trunk allowed-vlan 16,26-28,30
no shutdown
system vlan 16,30
```

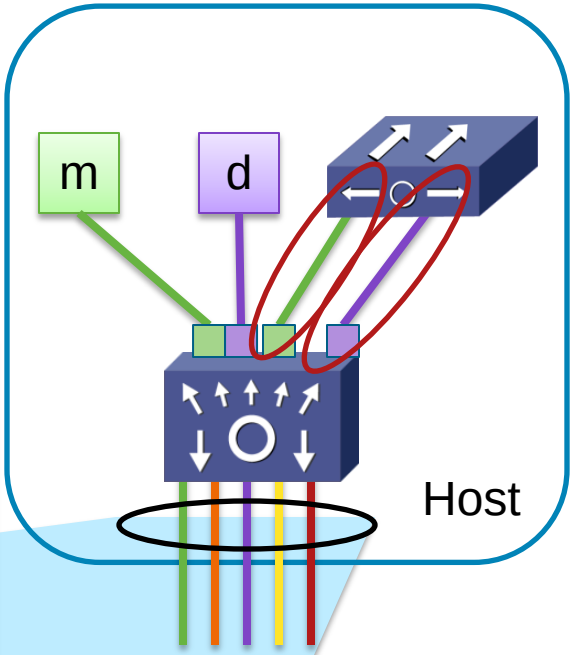

Inside the Nexus 1000V

Virtual NIC Types

Inside the Host…

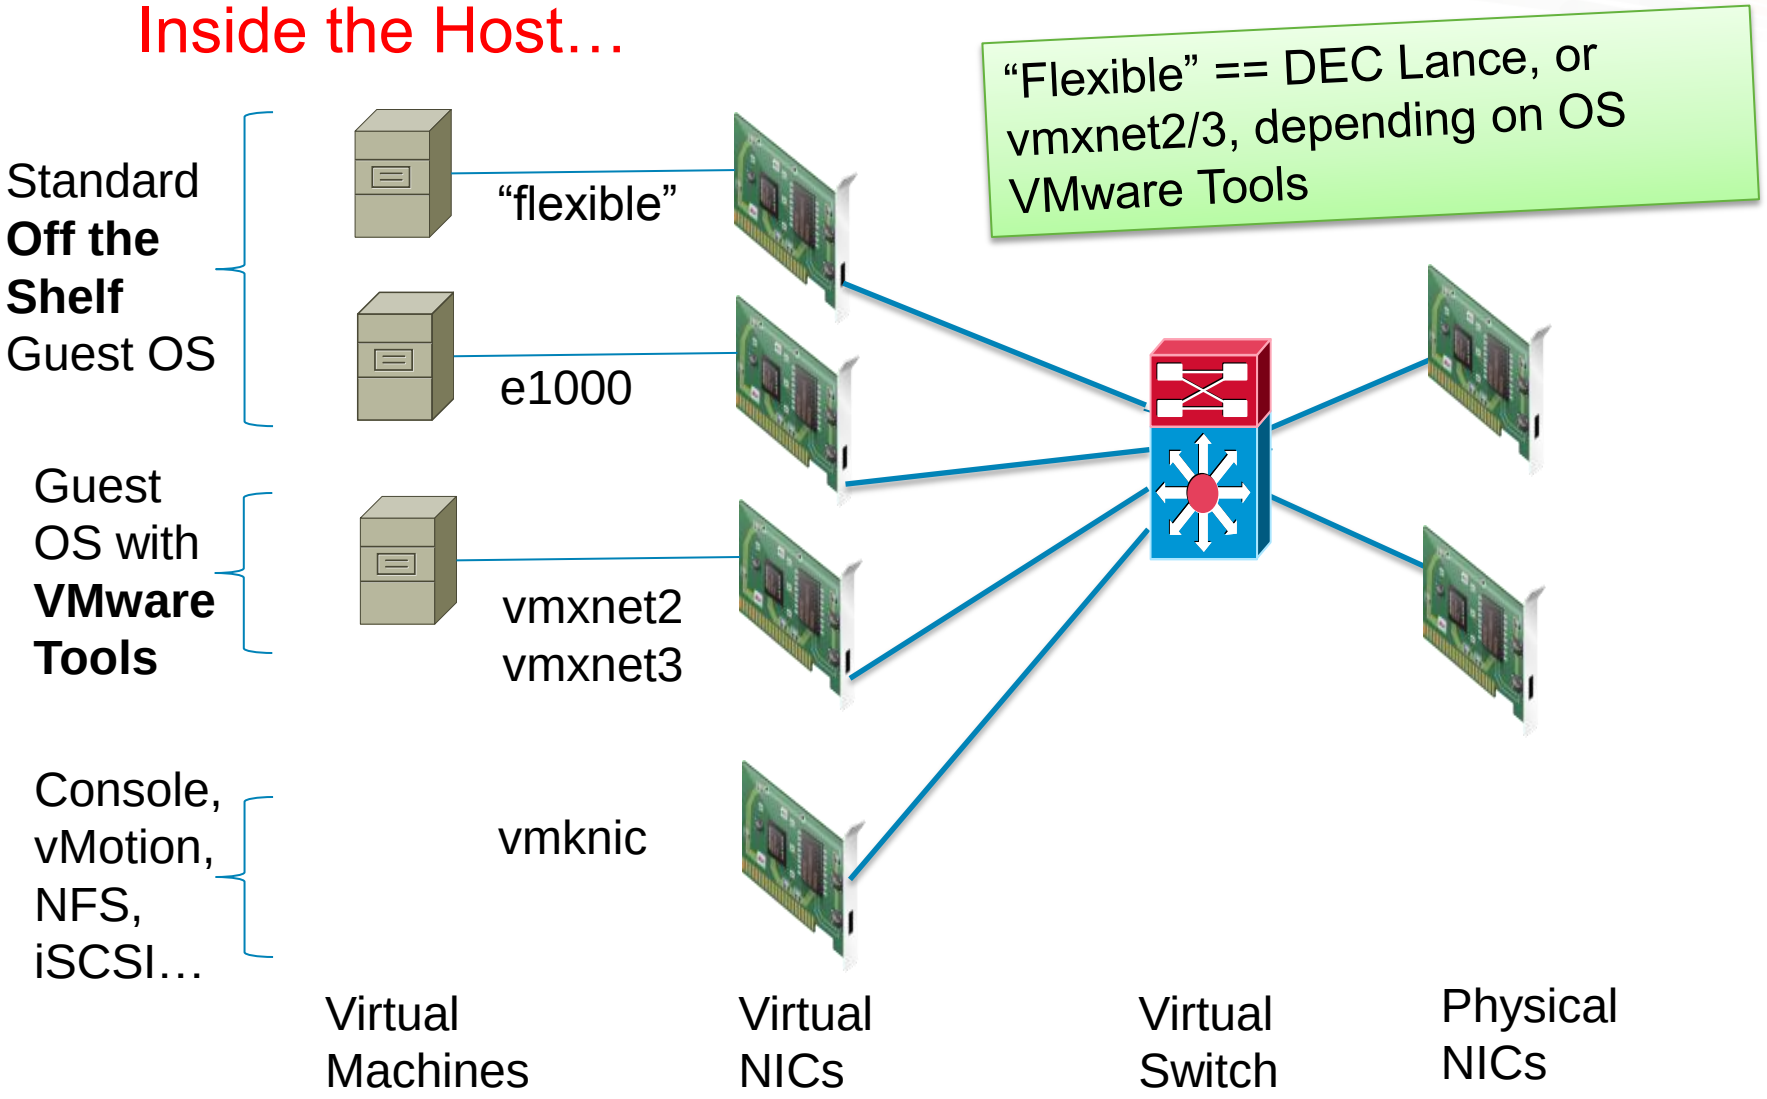

VMware Port IDs and Port Groups

Inside the Virtual Switch…

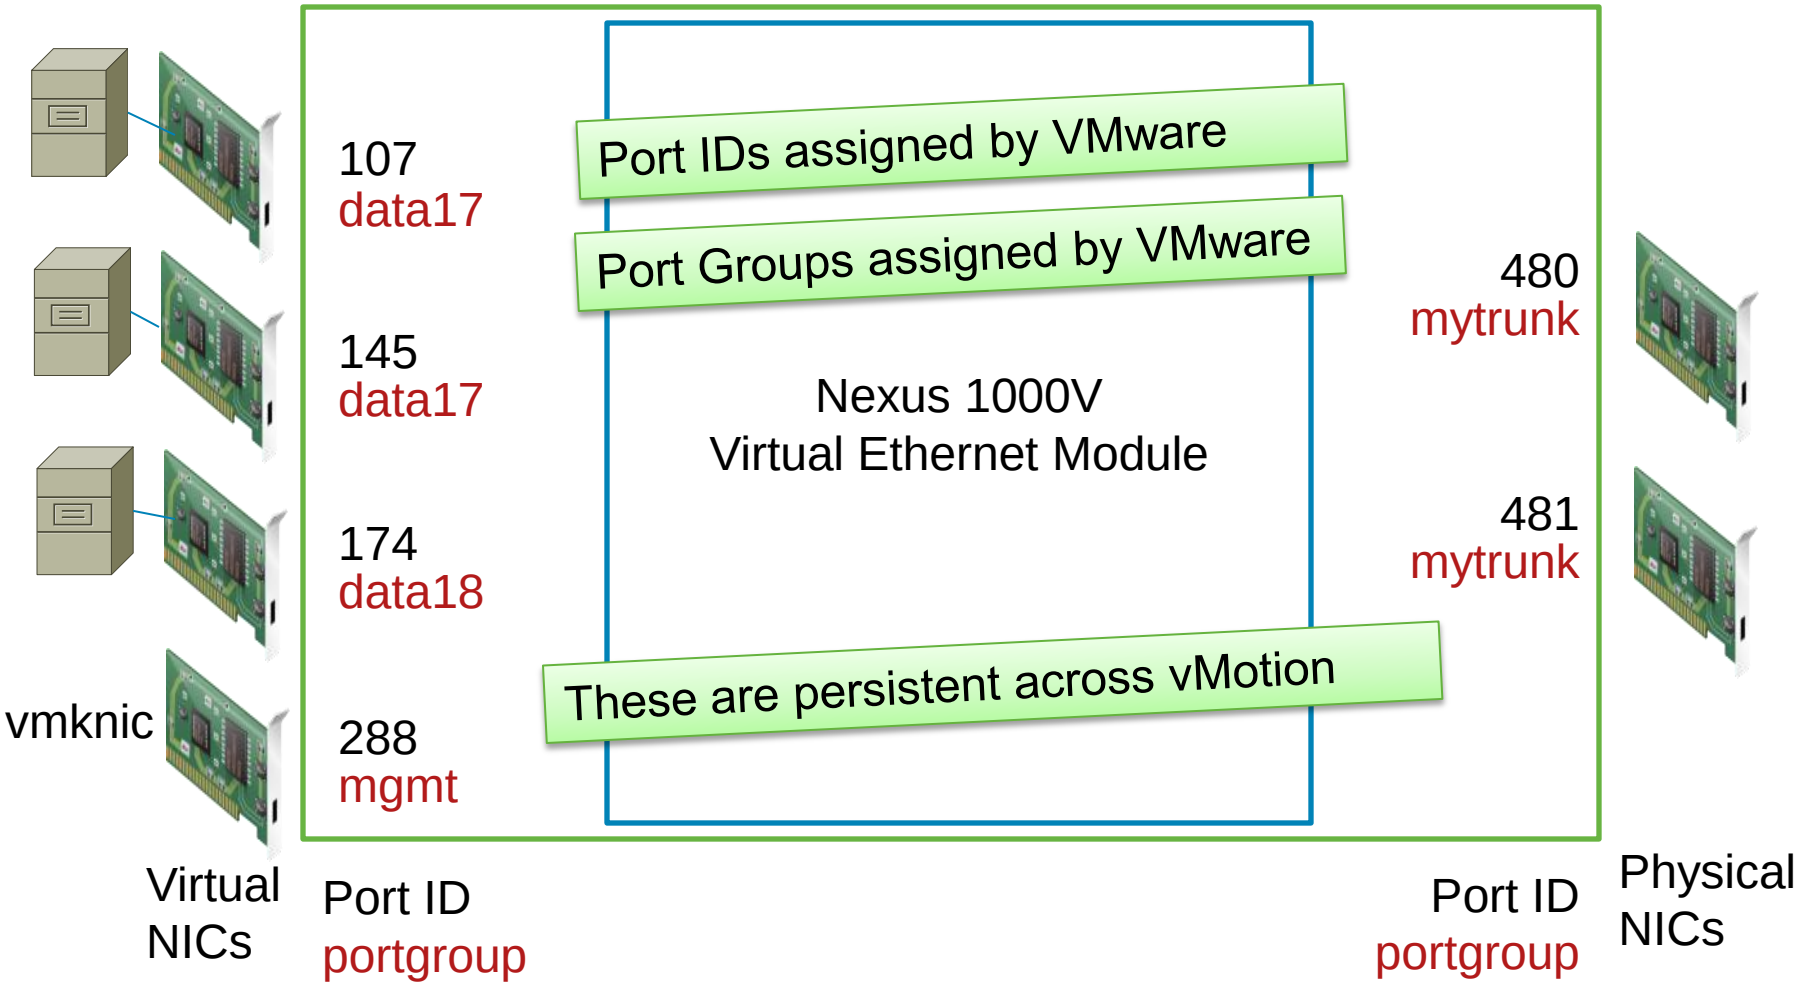

VEM LTLs (port numbers) and veths

Inside the VEM: veths and LTLs

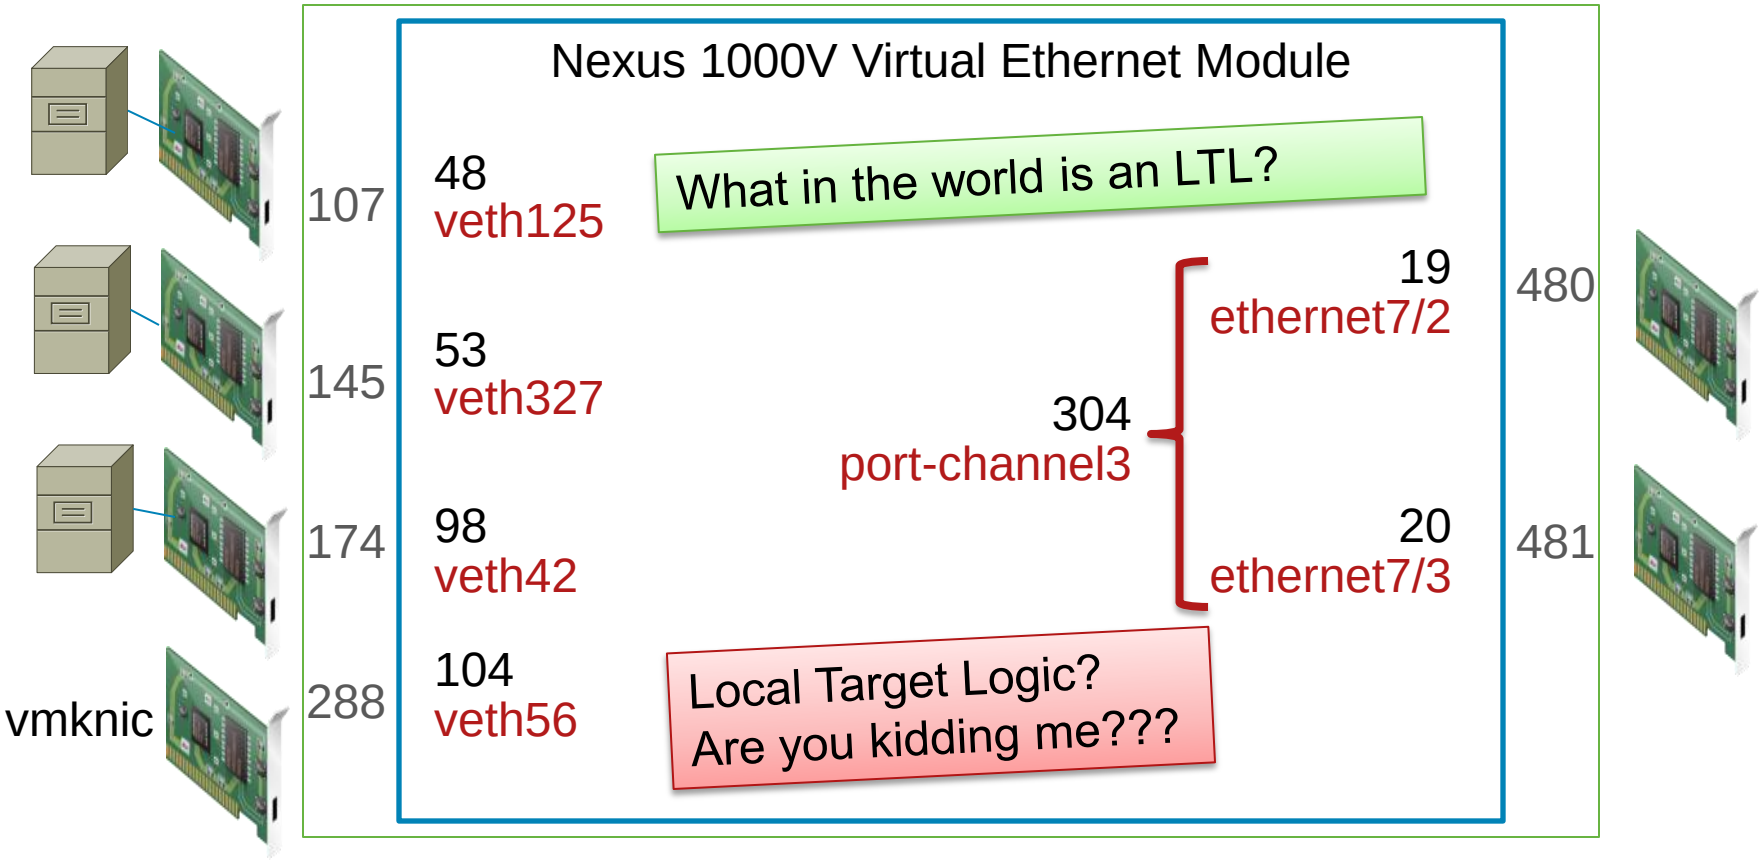

Each port has an LTL number, local to the host

Veth, port-channel and ethernet assigned by VSM

VLANs tie Virtual and Physical Together

Inside the VEM: VLANs

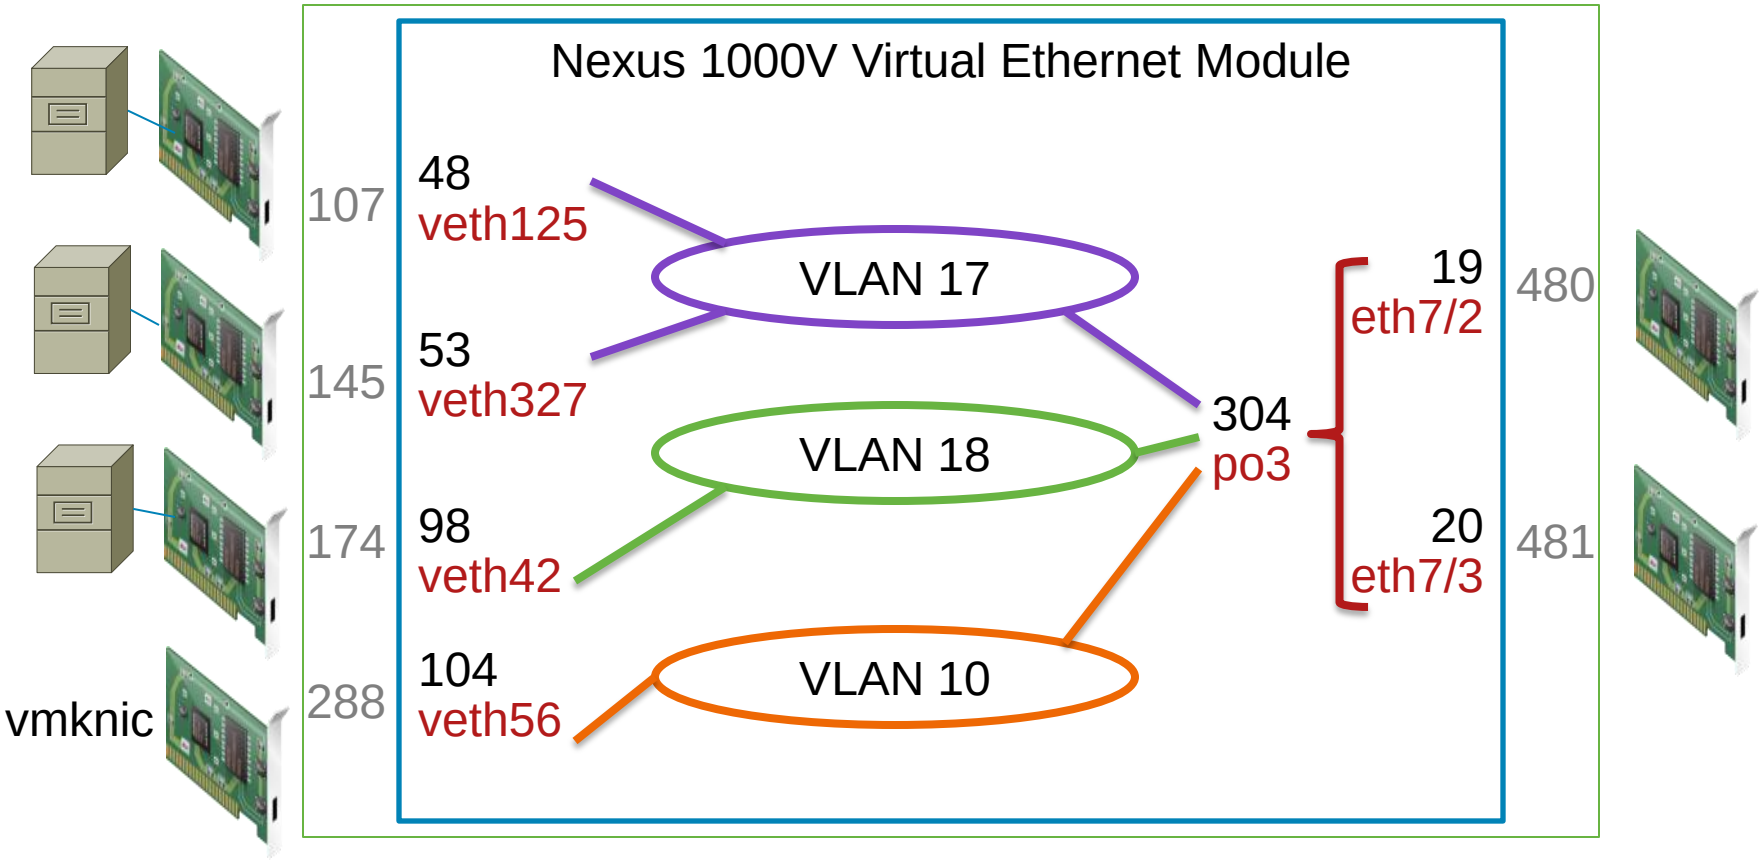

Most virtual ports are Access Ports

Most physical ports are Trunk Ports

MAC Tables for Dynamic Forwarding

Each VEM has its own MAC Table

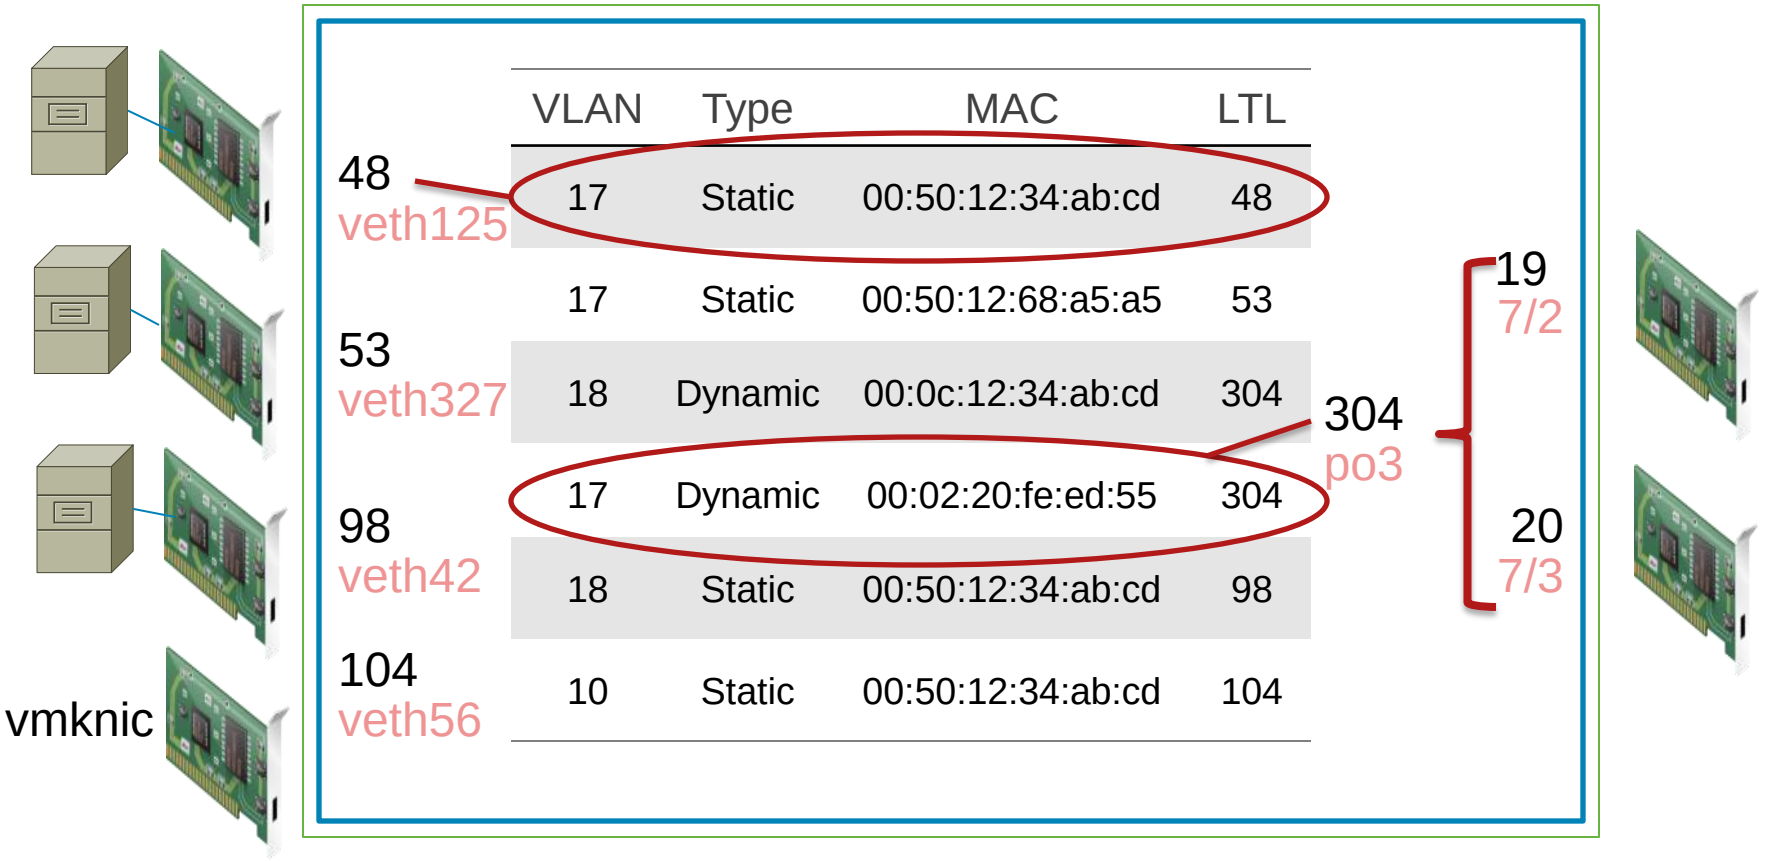

VMware supplies veth MAC addresses (static)

N1k learns dynamic MAC addresses

Distributed L2 Switching

Each VEM is an Independent L2 Switch

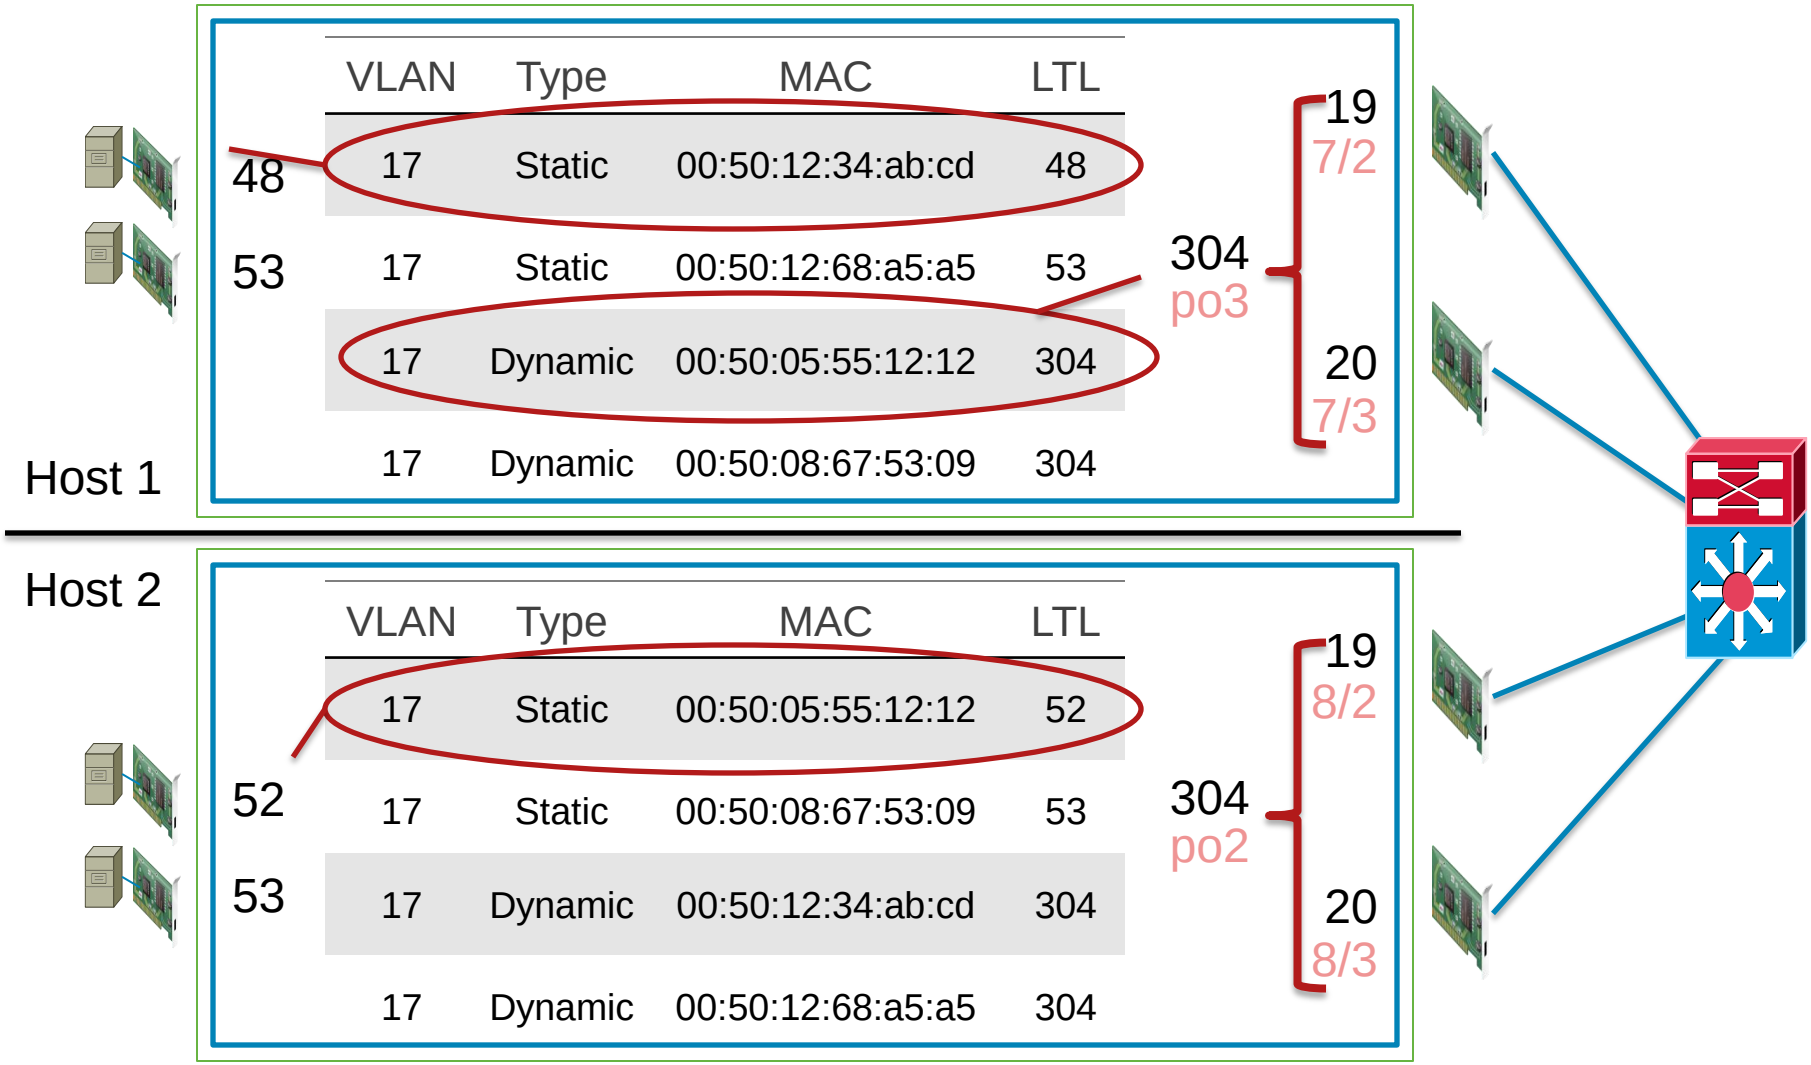

Packet Forwarding

Inside the VEM: L2 Packet Forwarding

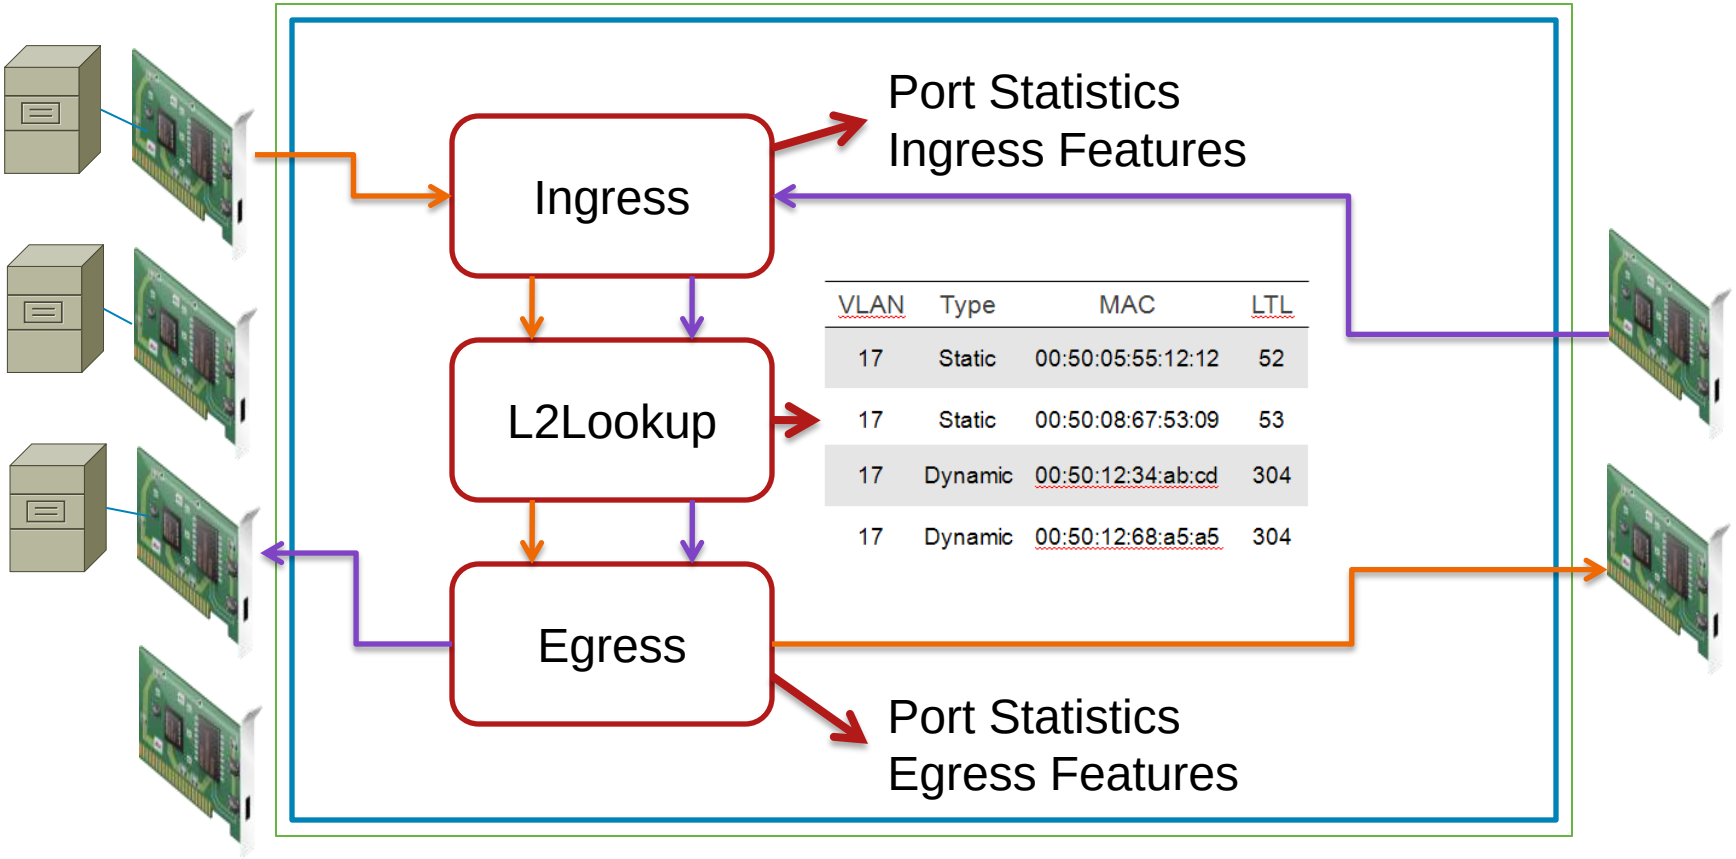

Packet Forwarding

Inside the VEM: Flooding, Multicast and PVLANs

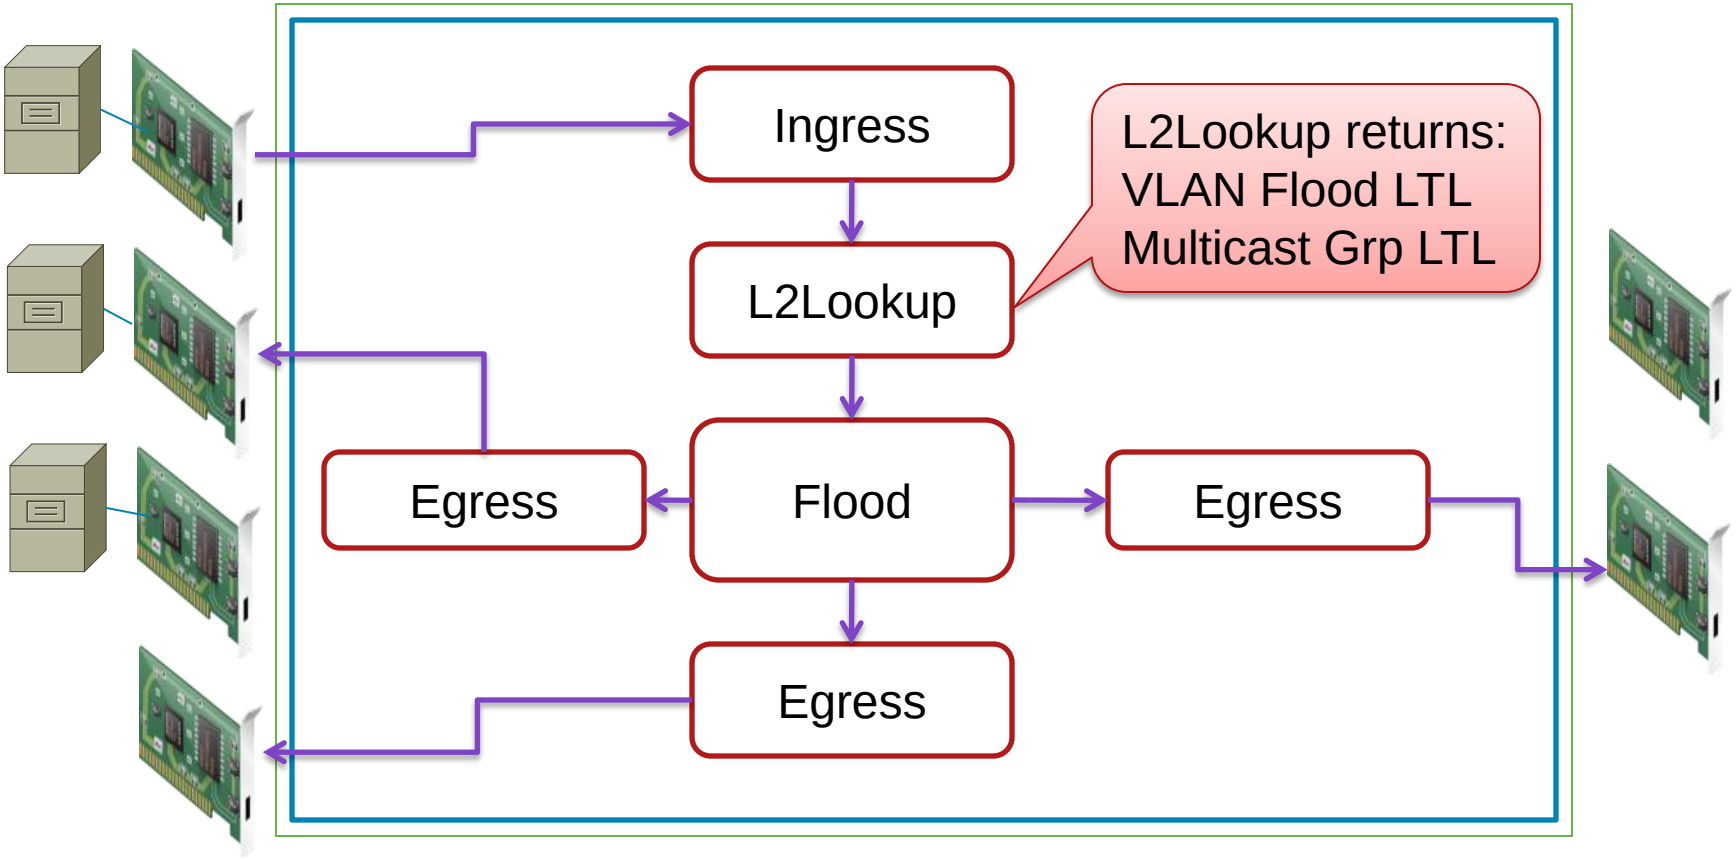

Multicast groups and VLANs have LTLs, too
These contain lists of actual LTLs
Each LTL has its own Egress processing

Packet Forwarding Features

Inside the VEM: Features

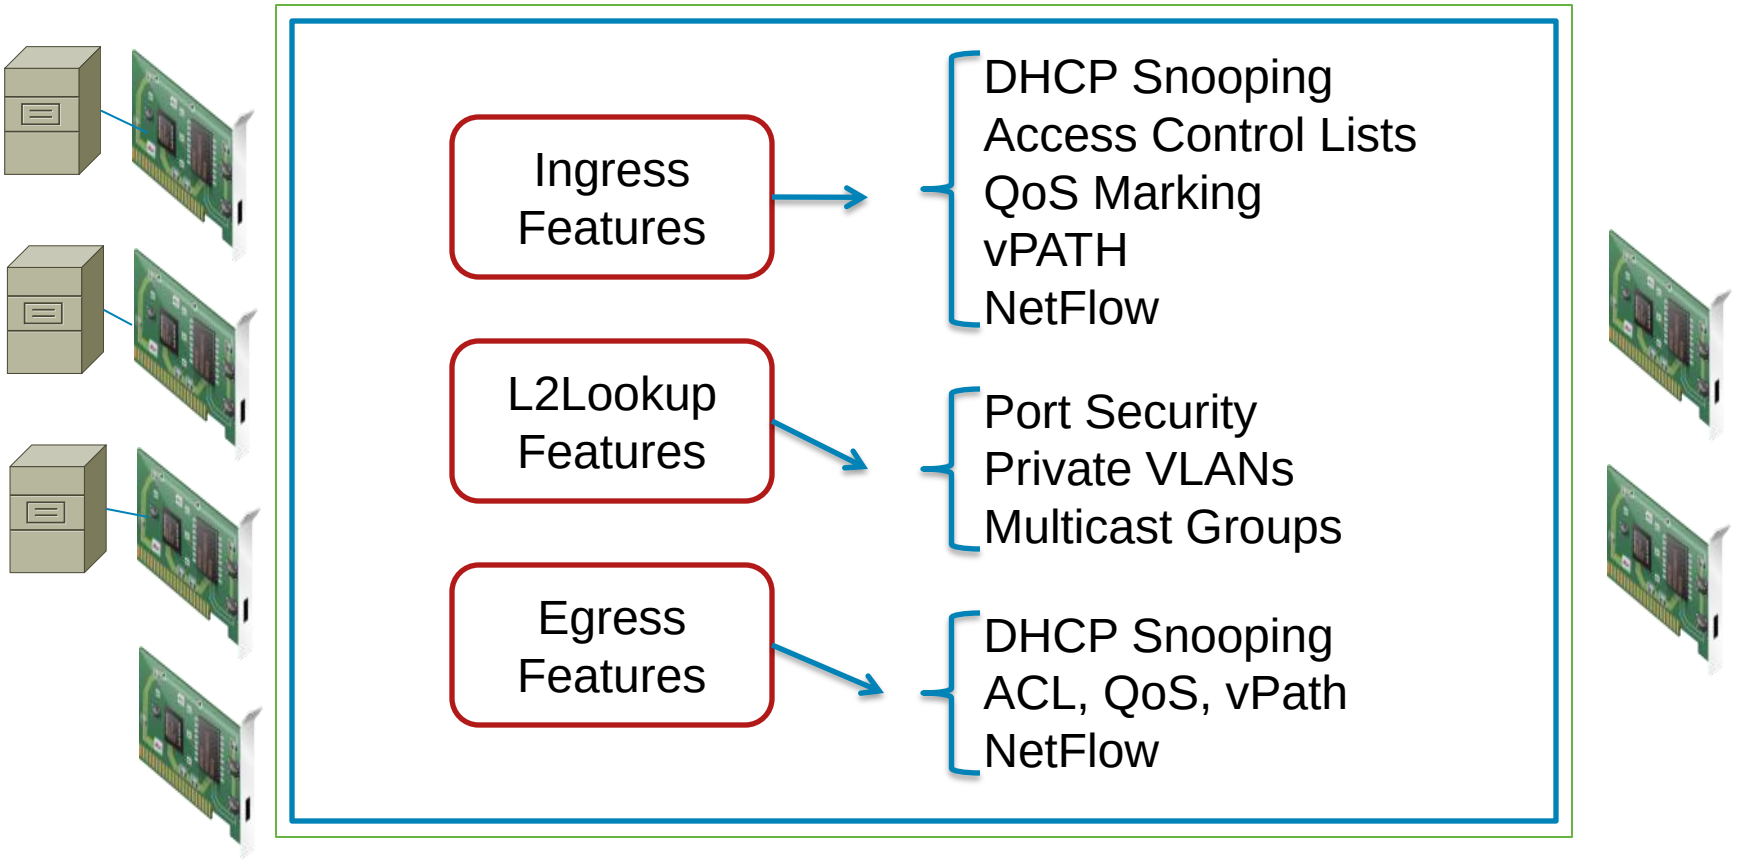

Port Channels and Upstream Switches

Common Port Channel Scenarios

Port Profiles allow vMotion across different host types

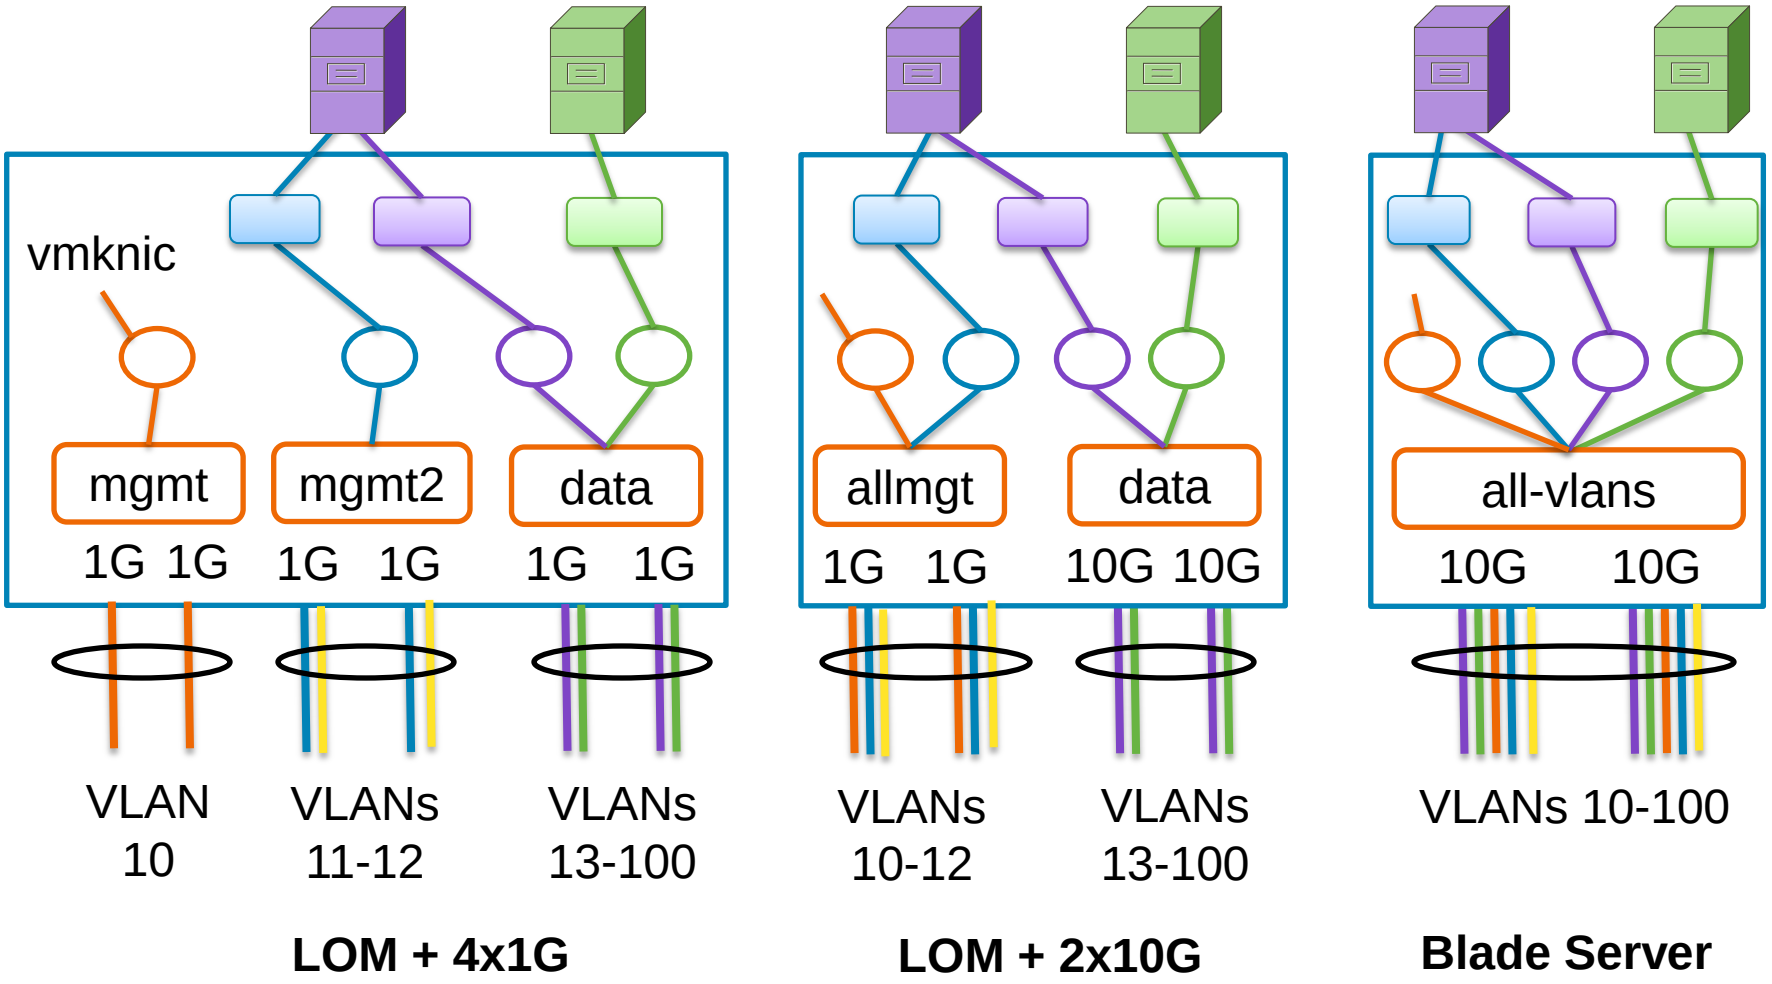

Port Channels with Palo (on UCS)

Nexus 1000V Features + Hardware Segmentation

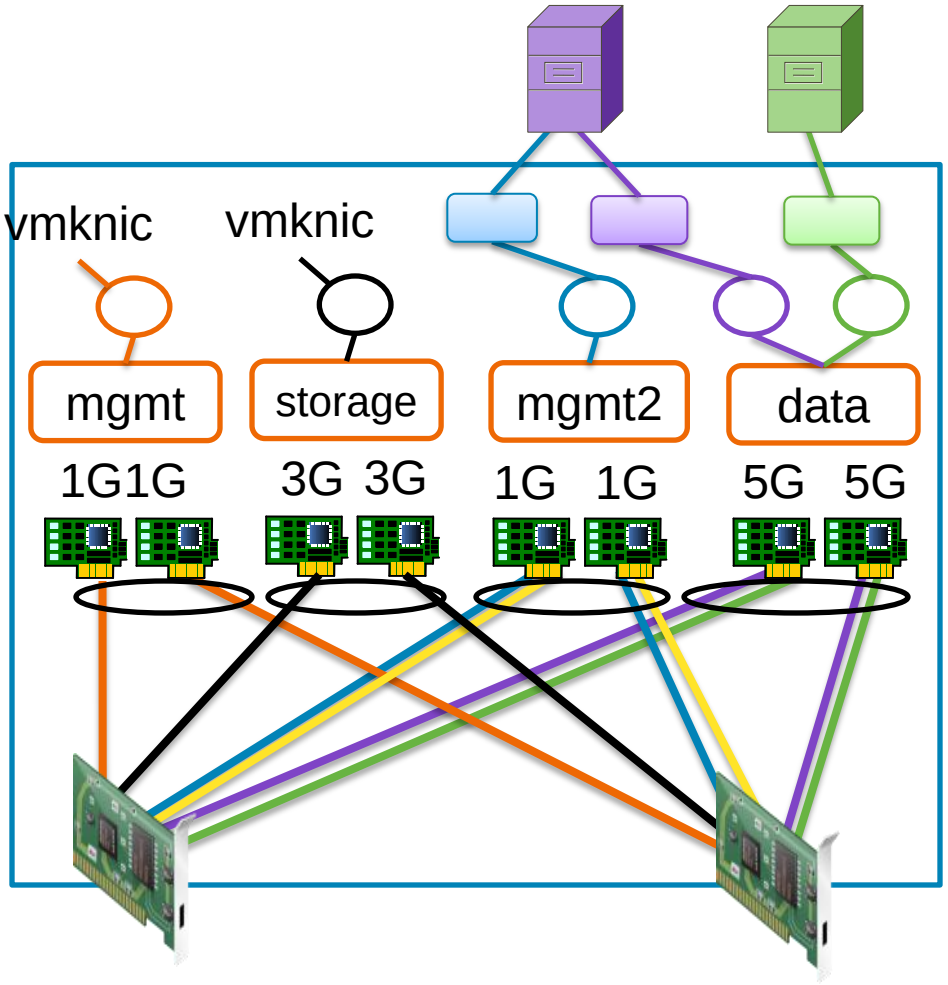

Port channels are used as normal

ESX and Nexus 1000V see Palo's virtual adapters as physical NICs

Each Palo can be provisioned to provide different "virtual" adapters, with bandwidth allocated

2x10G Palo Adapters

Common Port Channel Problems

Upstream switches have many configurations

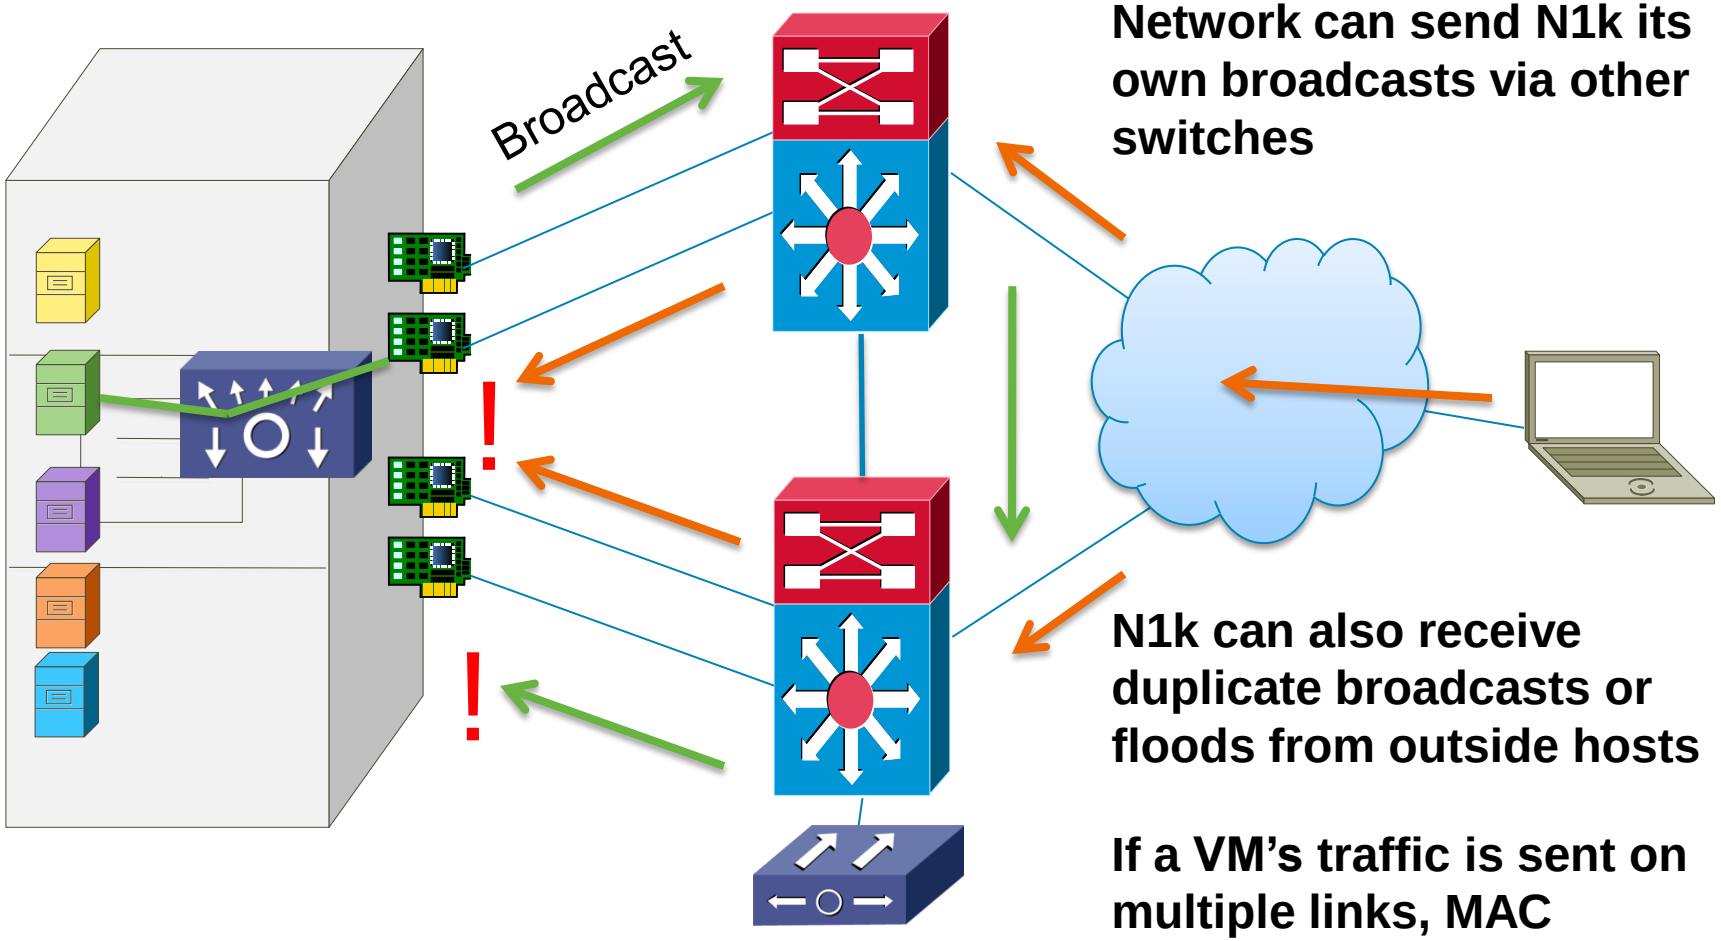

Network can send N1k its own broadcasts via other switches

N1k can also receive duplicate broadcasts or floods from outside hosts

If a VM's traffic is sent on multiple links, MAC addresses could "flap"

Common Port Channel Solutions

Essential to choose the Right One for your network

- LACP for clustered switches (Cisco vPC, VSS)
  - Use LACP (now offloaded to the VEM)
- vPC-HM MAC Pinning
  - Un-clustered switches without port channels (e.g. UCS)
  - Or when using one port per upstream switch
- vPC-HM Subgroup CDP
  - Un-clustered switches with port-channels enabled (e.g. N5k)
  - vPC-HM Manual Subgroups also available (not covered)
- vPC-HM + Network State Tracking
  - Use with HP Virtual Connect when SmartLink not used
  - Handles loops and split networks

LACP for “Clustered” switches

Cisco vPC, VSS, VBS, Stack

802.3ad LACP used by both sides to agree on how to load balance

Flow-based balancing allows more than one physical NIC to be used by a virtual NIC

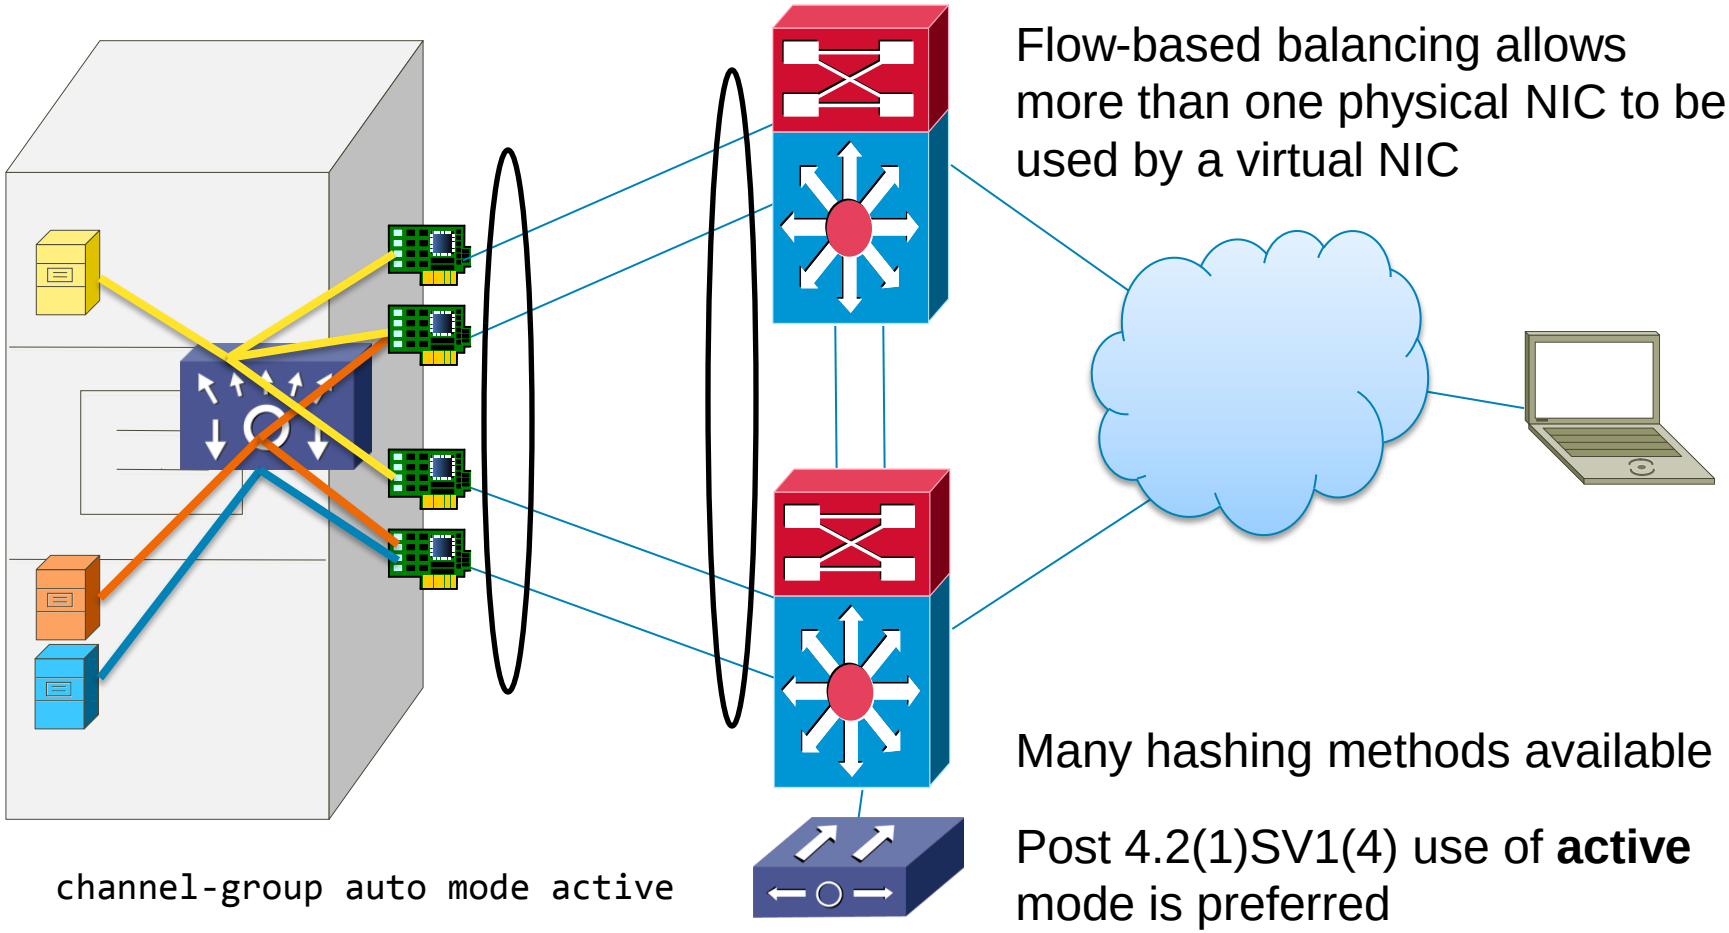

Many hashing methods available

Post 4.2(1)SV1(4) use of **active** mode is preferred

```
channel-group auto mode active
```

LACP Recovery

When NIC fails, LACP recovers

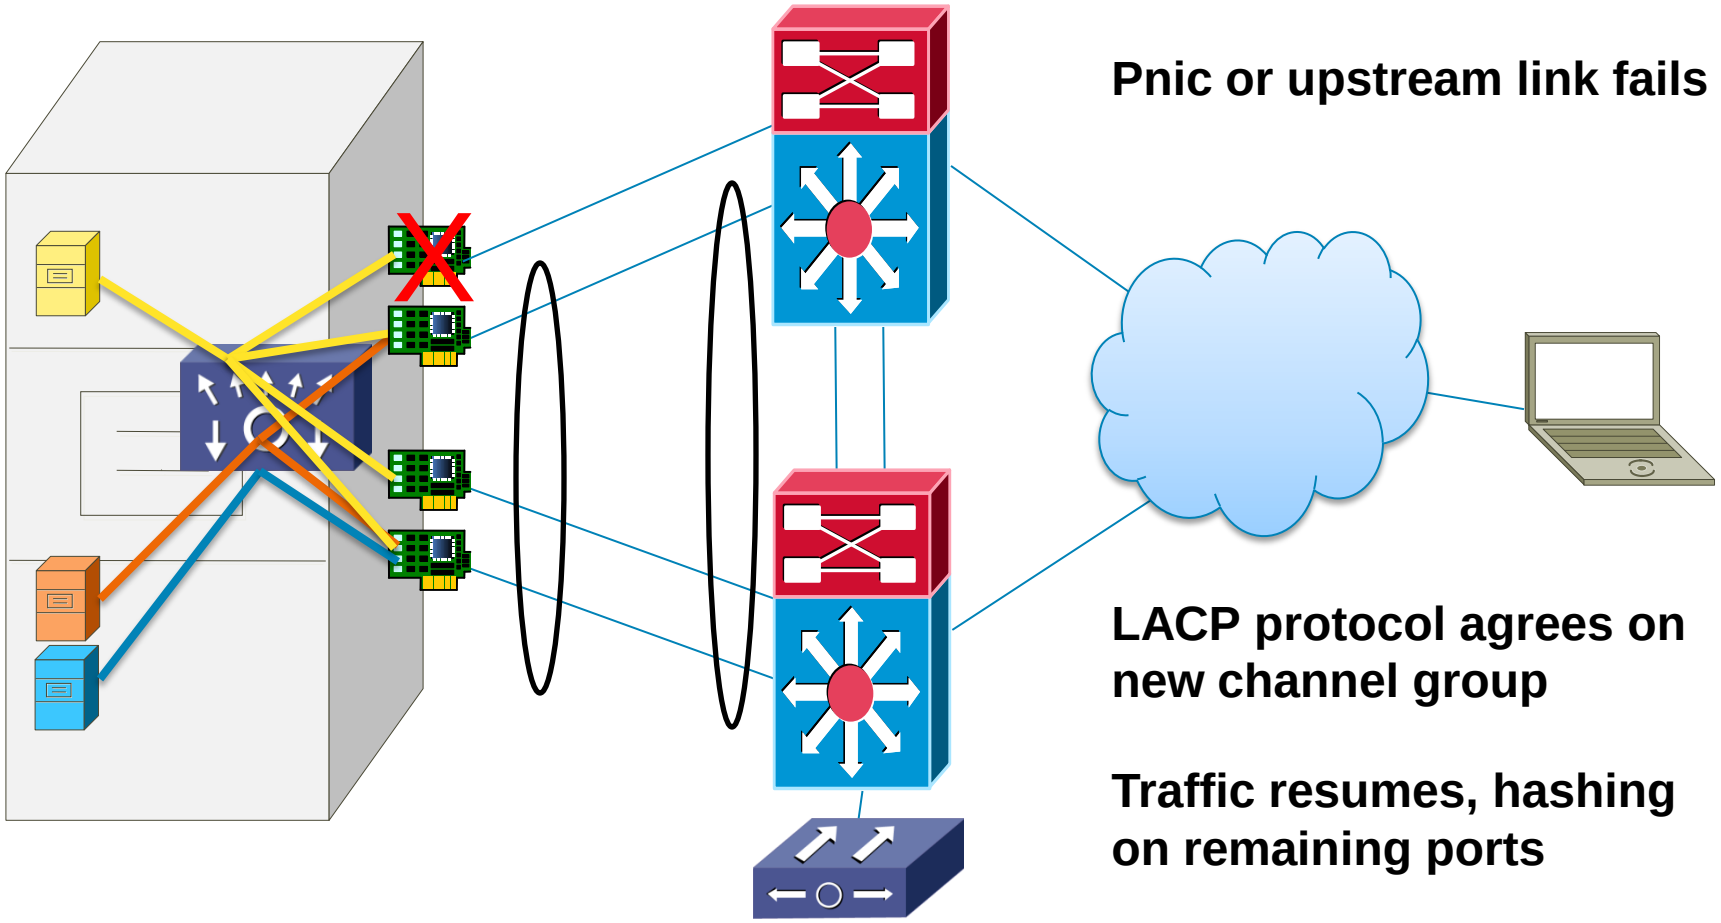

Pnic or upstream link fails

LACP protocol agrees on new channel group

Traffic resumes, hashing on remaining ports

Tip: Use the new “LACP Offload” feature for FCoE or Boot from SAN

MAC Pinning – Keeping it Simple

Simplest configuration; no upstream features required

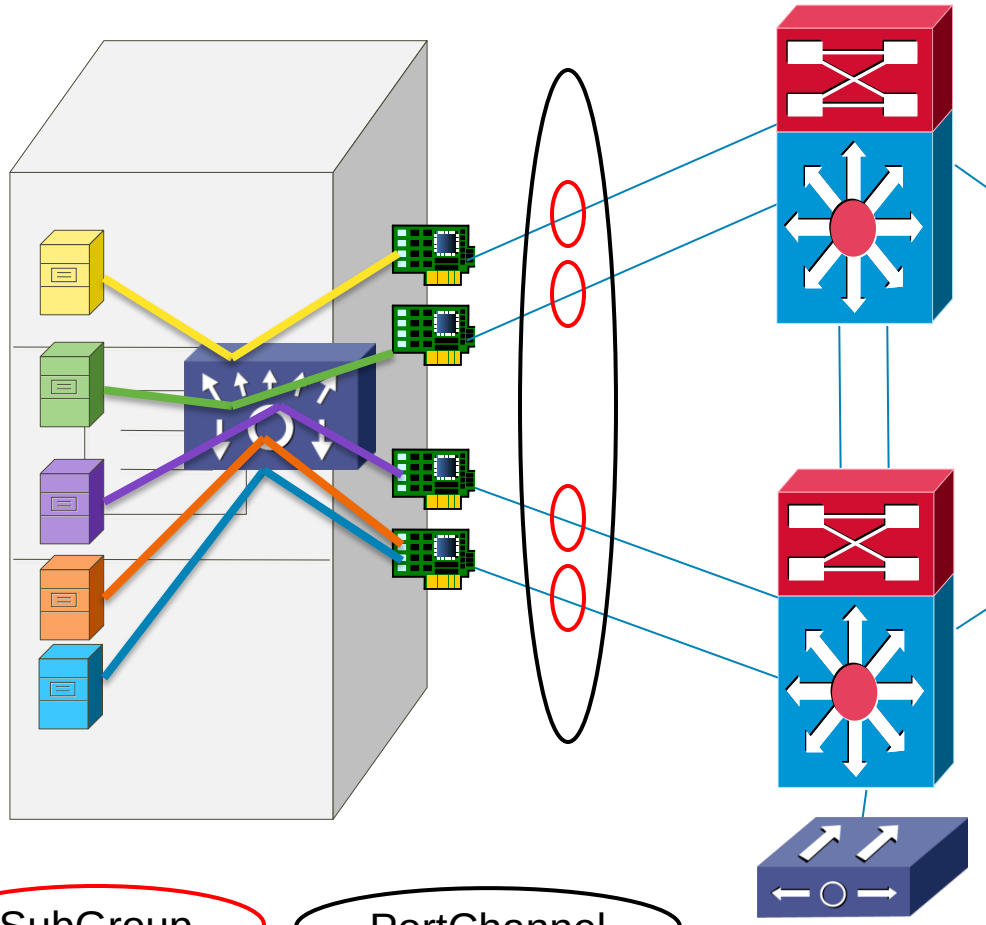

SubGroup

PortChannel

All Pnics with the same port profile create a port channel

Each pnic is formed into its own subgroup

Each vnic is pinned to a particular pnic

vnics balanced across pnics

static pinning also available

MAC Pinning

When NIC fails, Vnics are re-pinned

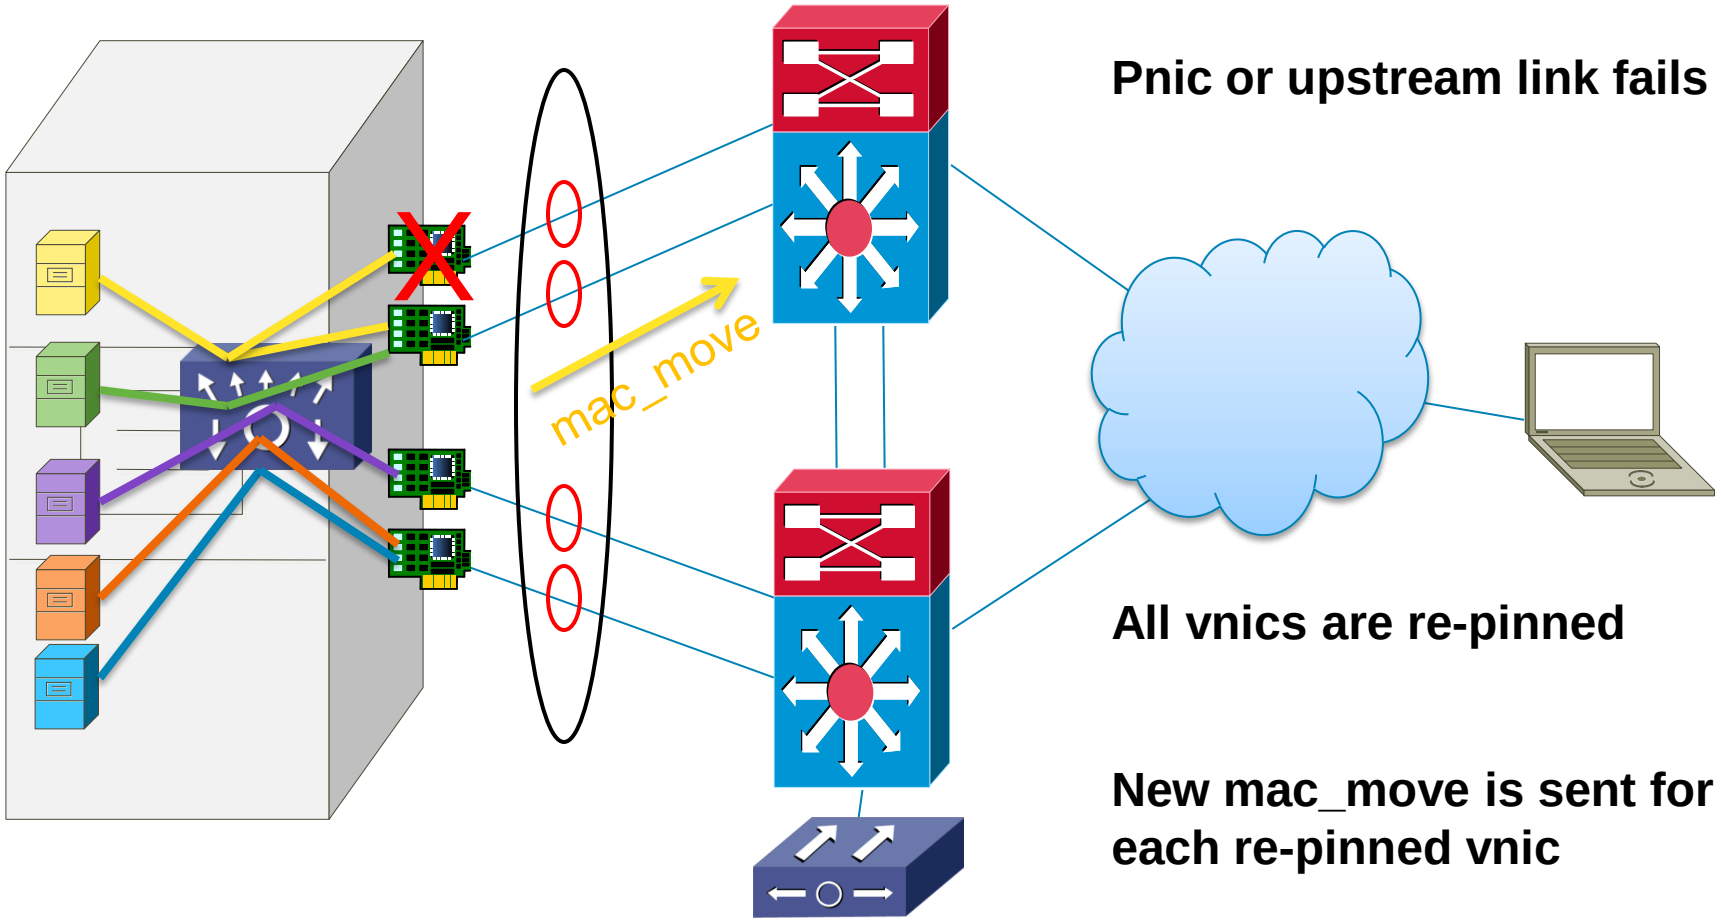

Tip: Do not configure upstream port channel when using MAC pinning

vPC-HM – Subgroup CDP

Use when upstream switches are independent channels

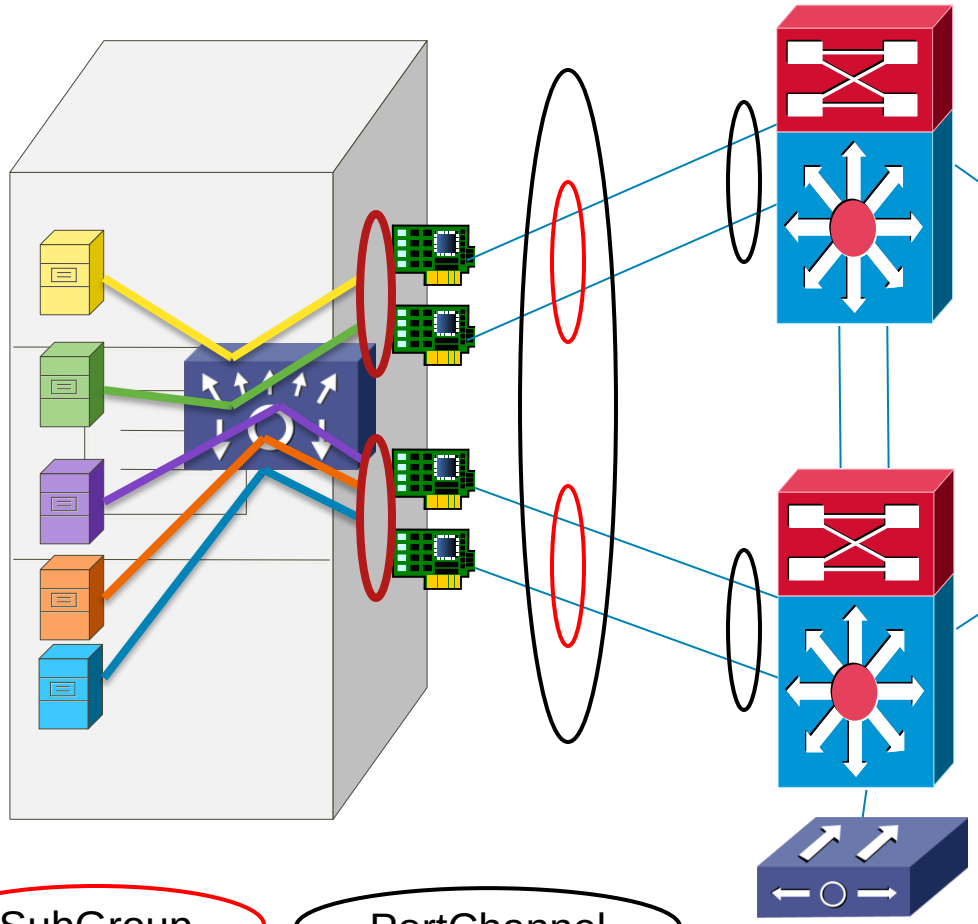

SubGroup

PortChannel

All Pnics with the same port profile create a port channel

Subgroups are formed for each physical switch using CDP info

Each vnic is pinned to a particular subgroup

vPC-HM – Avoiding Duplicates

Use when upstream switches are independent channels

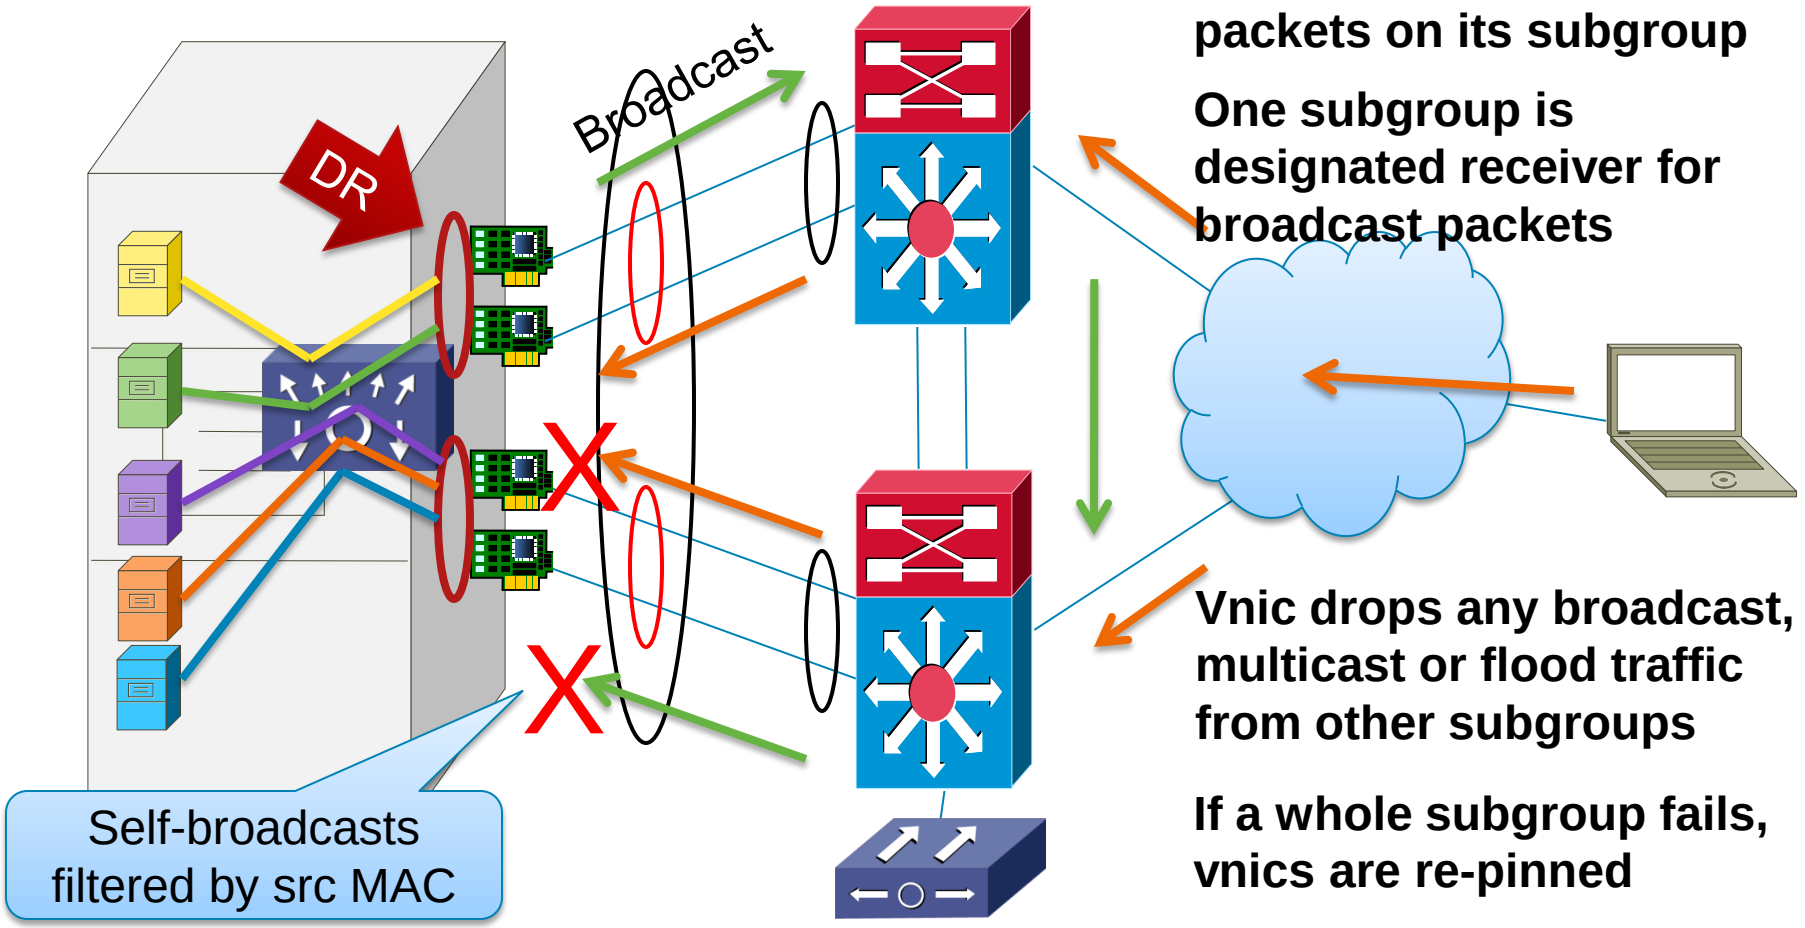

Each vnic only sends packets on its subgroup

One subgroup is designated receiver for broadcast packets

Vnic drops any broadcast, multicast or flood traffic from other subgroups

If a whole subgroup fails, vnics are re-pinned

Tip: Subgroup Manual is also available for non-CDP devices

Network State Tracking

Use with vPC-HM when link detection doesn't work

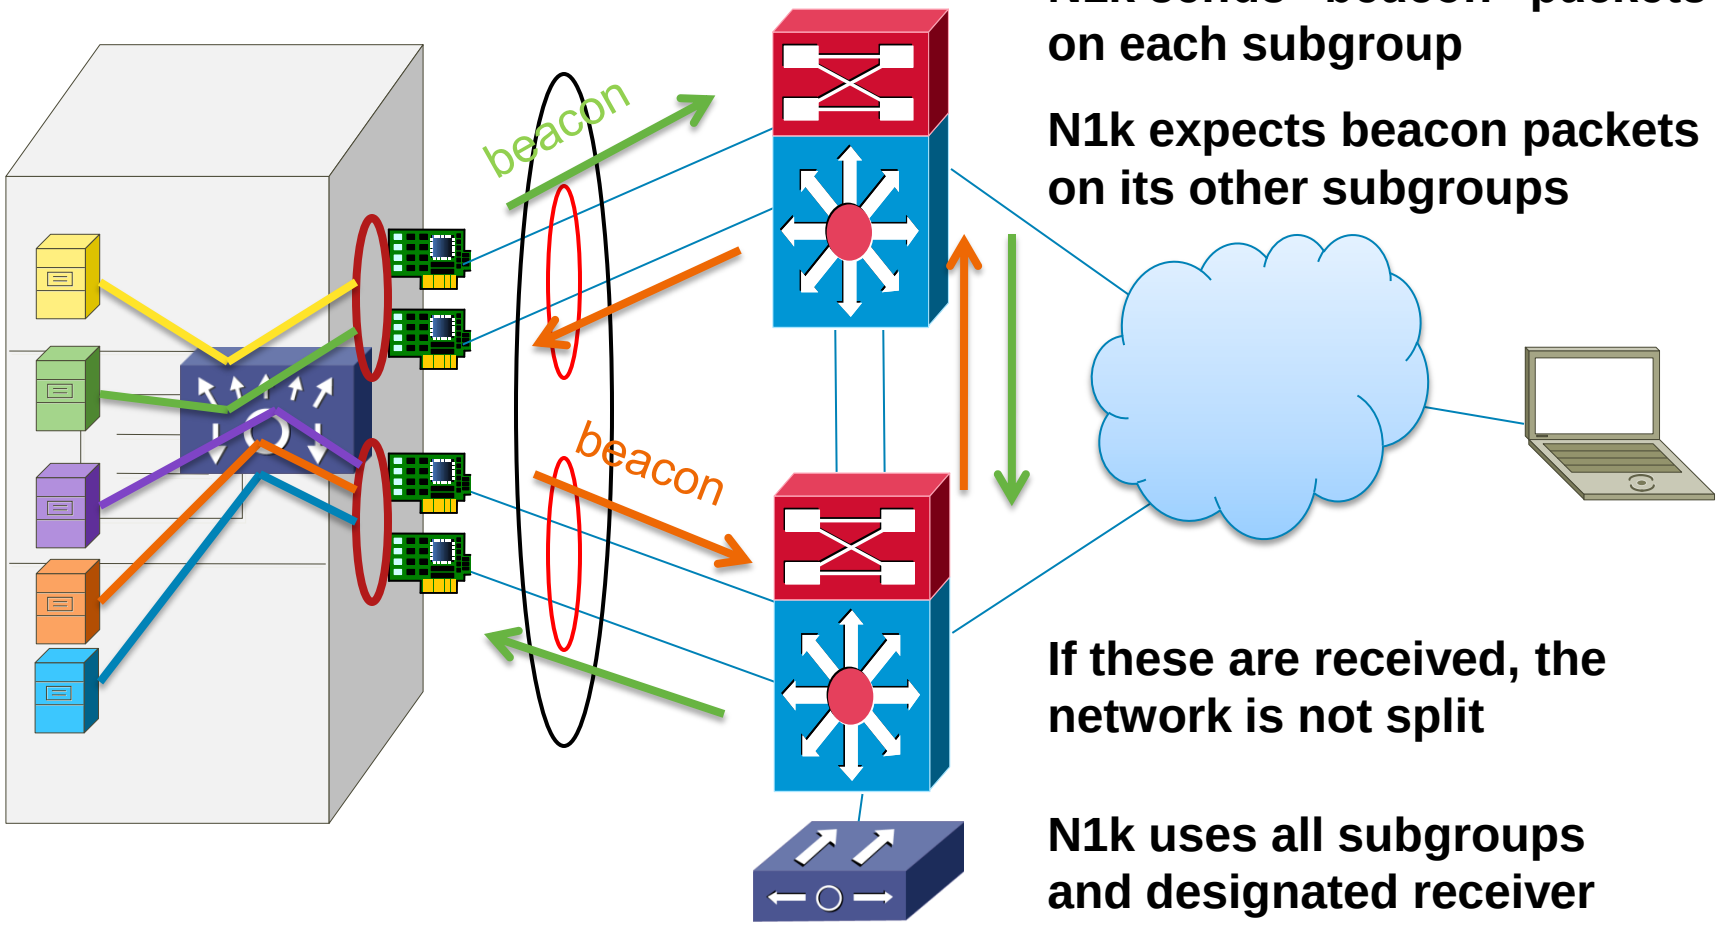

N1k sends "beacon" packets on each subgroup

N1k expects beacon packets on its other subgroups

If these are received, the network is not split

N1k uses all subgroups and designated receiver

Network State Tracking

When a subgroup goes down, network is “split”

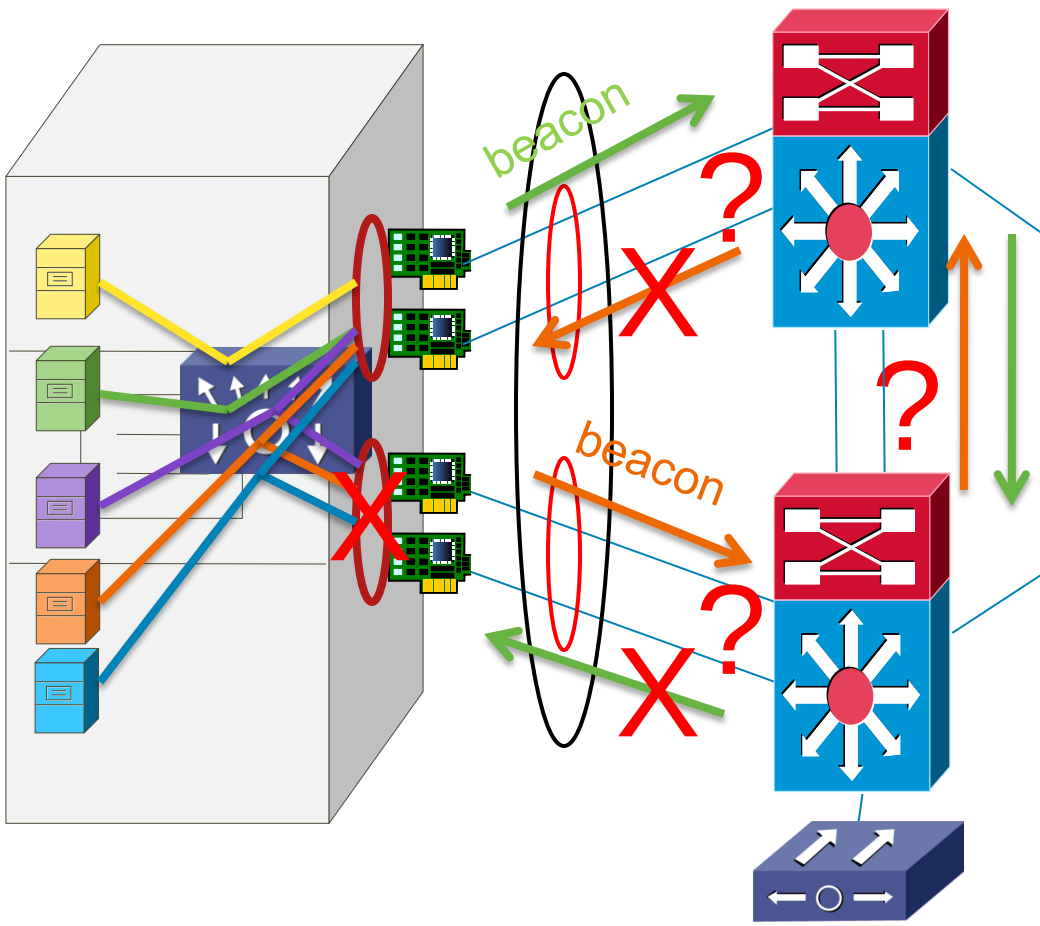

Hardware failure happens

… but we don’t know where!

N1k misses beacon packets on its other subgroups

Since network is split, N1k chooses a subgroup as the active network

Vnics are re-pinned to this subgroup

Tip: NST can be set to just syslog instead

Deploying the Nexus 1000V

Nexus 1000V—How Did It Get There?

- VSM is created as a virtual machine
  - From ISO image, using VSM setup wizard

```
  Enter the password for "admin":
  Confirm the password for "admin":
  Enter HA role[standalone/primary/secondary]: primary
  Enter the domain id<1-4095>: 47
…
Would you like to enter the basic configuration dialog? Yes
…
  Enter the switch name: mbakke-47
   Mgmt0 IPv4 address: 192.168.16.236
   Mgmt0 IPv4 netmask: 255.255.255.0
…
  Enter SVS Control mode (L2 / L3) : L3
…
```

Nexus 1000V—How Did It Get There?

… or from “ovf” package with new installer application

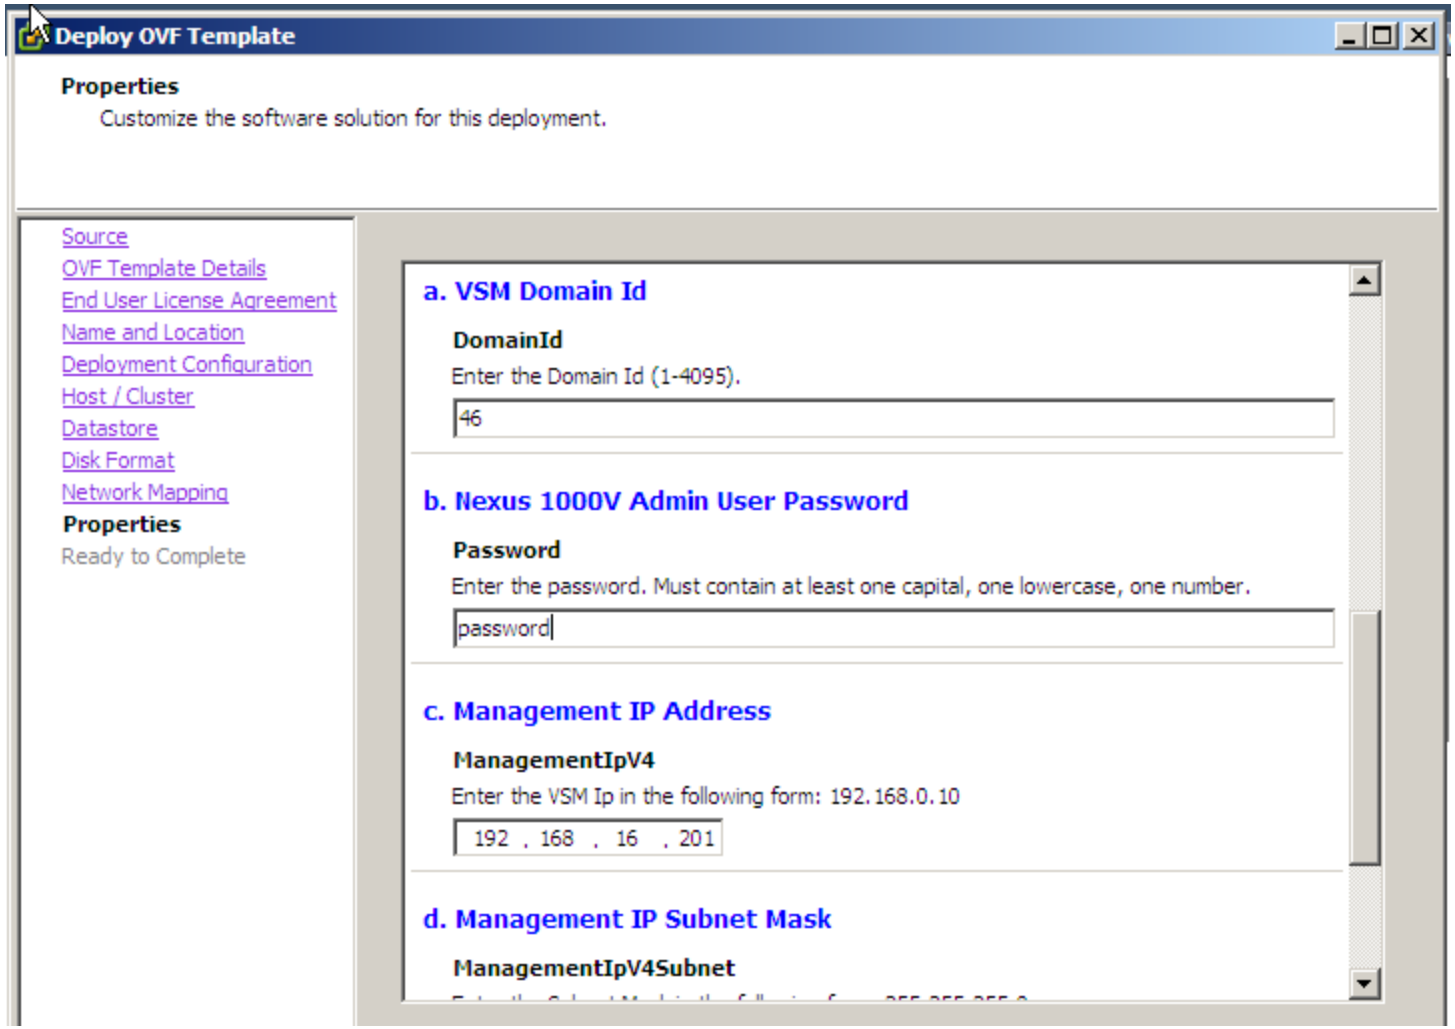

Nexus 1000V in vCenter

VSM shown in vCenter as a Distributed Virtual Switch

Distributed Virtual Switch

Nexus 1000 Virtual Switch
-Locations of VEM images
-Switch data for VEMs
-VSM certificate

```
vsm-main# show svs connec
connection vc:
    ip address: 192.168.1
    certificate: default
    datacenter name: mbak
    config status: Enable
```

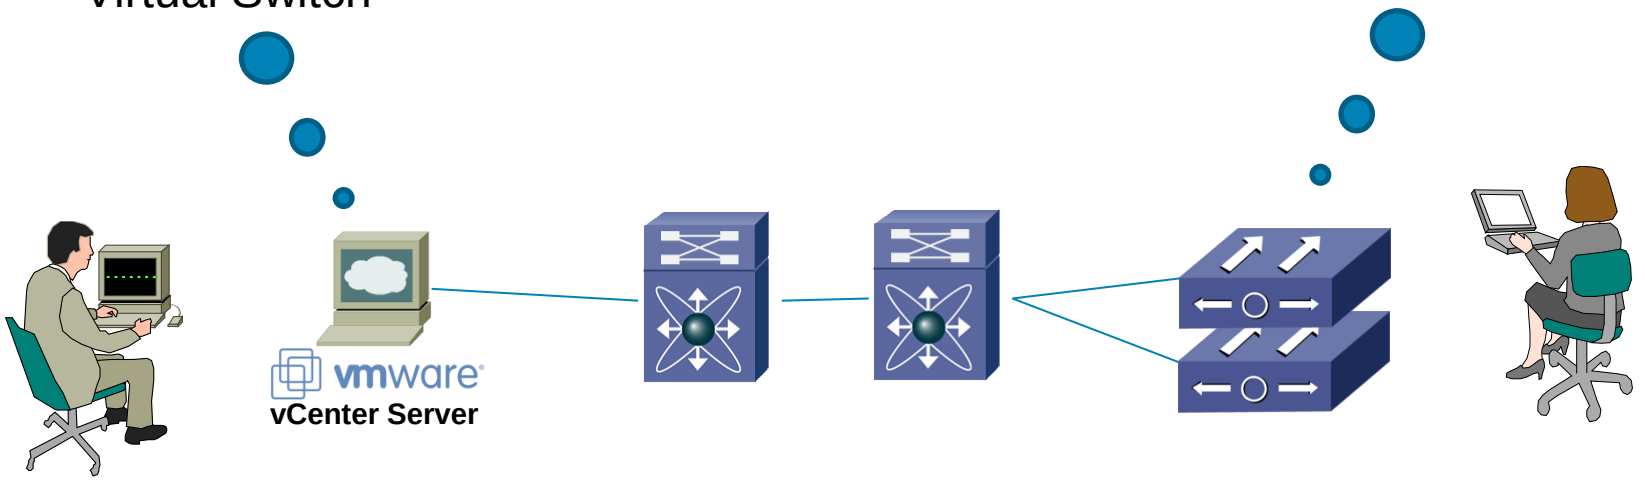

VEM Software Manual Installation

VEM Can Be Installed Manually on Each Host

```
PS C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>
./vihostupdate.pl --server mbakke-ucs --username root --install
 --bundle …\VEM\cisco-vem-v130-4.2.1.1.4.0.0-2.0.1.zip
```

Using the vSphere CLI

… or on the host with ssh

```
 # esxupdate -b cross_cisco-vem-v130-4.2.1.1.4.0.0-2.0.1.vib update
Unpacking cross_cisco-vem-v13..
########################################## [100%]
Installing packages :cross_ci..
########################################## [100%]
Running [ipkg -f /tmp/ipkg.conf-n1kv remove cisco-vem-v124-esx]...
ok.
Running [/usr/sbin/vmkmod-install.sh]...
ok.
```

VMware Update Manager

Or Automatically When a Physical Port Is Selected

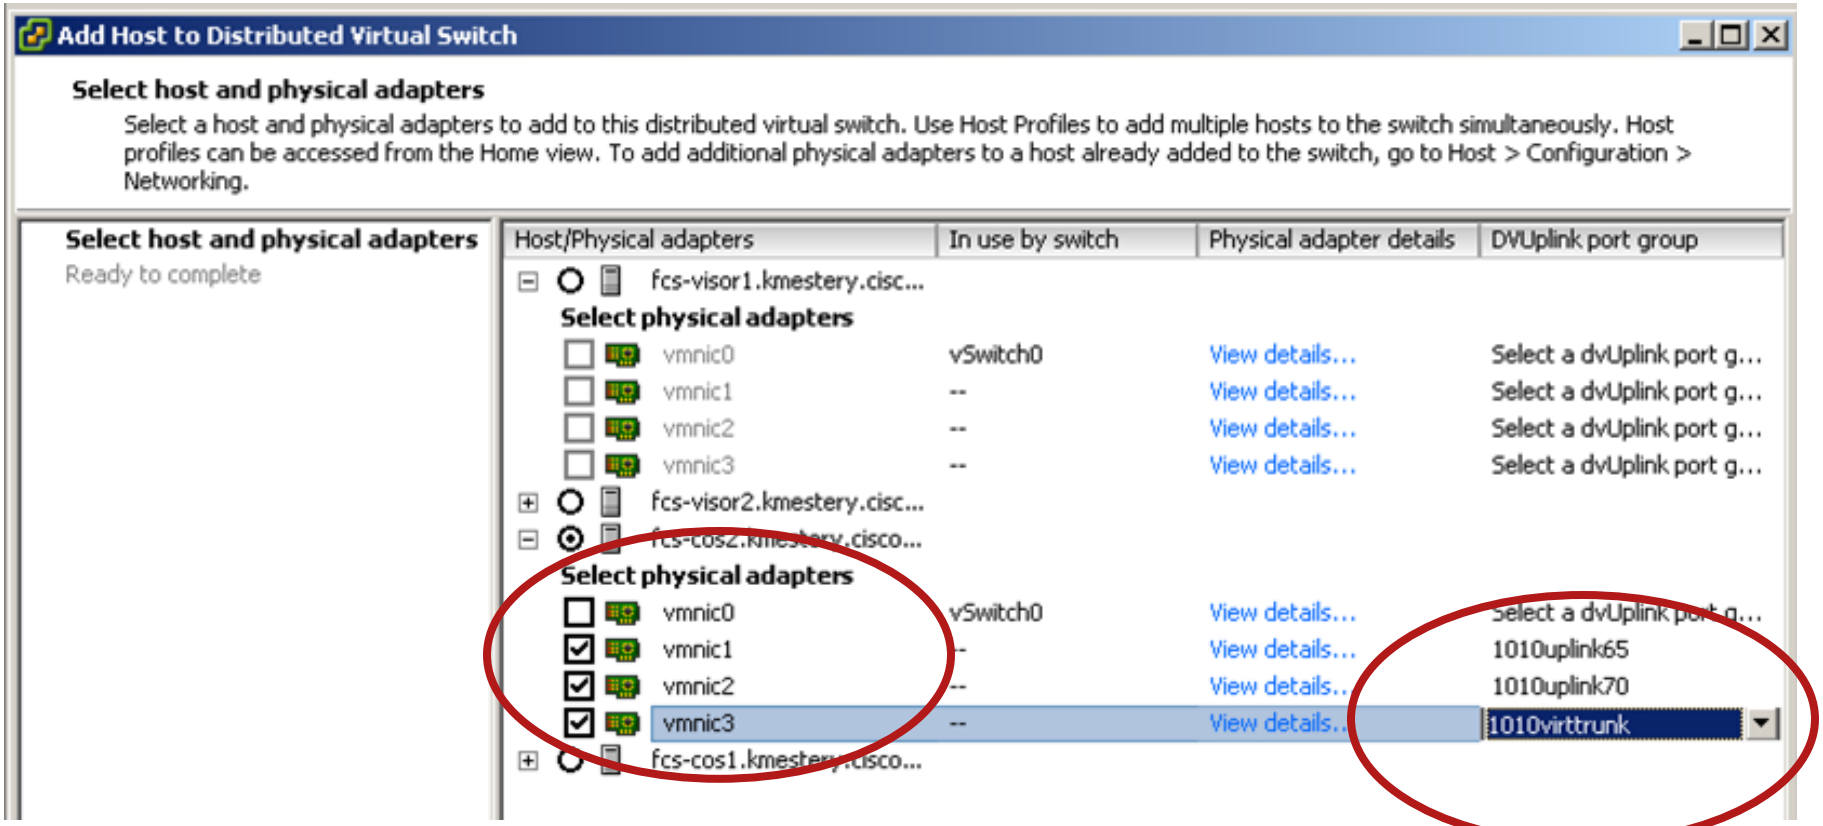

Ports and Uplink Port Profiles are Selected

… then VEM software is installed in the background and automatically started

VEM Automatically Found the VSM!

Wait a Minute; How Did It Know?

- Nexus 1000V DVswitch has custom “opaque data”
- Contains VSM’s contact information
- And initial physical port configuration

Switch-domain 40

VSM IP 192.168.30.7

Uplink System VLANs

"com.cisco.svs.switch.config"

"data-version 1.0 switch-domain 40 switch-name vsm-main cp-version 4.2(1)SV1(4) control-vlan 1 system-primary mac 00:0c:29:62:25:df active-vsm packet mac 00:0c:29:62:25:f3 active-vsm mgmt mac 00:0c:29:62:25:e9 standby-vsm ctrl mac ffff-ffff-ffff inband-vlan 1 svs-mode L3 l3control-ipaddr 192.168.30.7 upgrade state 0 mac ffff-ffff-ffff l3control-ipv4 null profile dvportgroup-562 access 30 profile dvportgroup-562 capability l3control profile dvportgroup-258 access 16 profile dvportgroup-277 access 29 profile dvportgroup-276 access 30 profile dvportgroup-255 access 26 profile dvportgroup-247 trunk 17-18,26,29-30 profile dvportgroup-254 access 17 profile dvportgroup-257 access 19 profile dvportgroup-256 access 16 end-version 1.0 "

Deploying Large Numbers of Hosts

No Network Admin Actions Required!

- VUM for VEM installation
- Set up a host

  Complete with port profiles!
- Create a host profile
- Add hosts using host profile
- Nexus 1000V is Added!

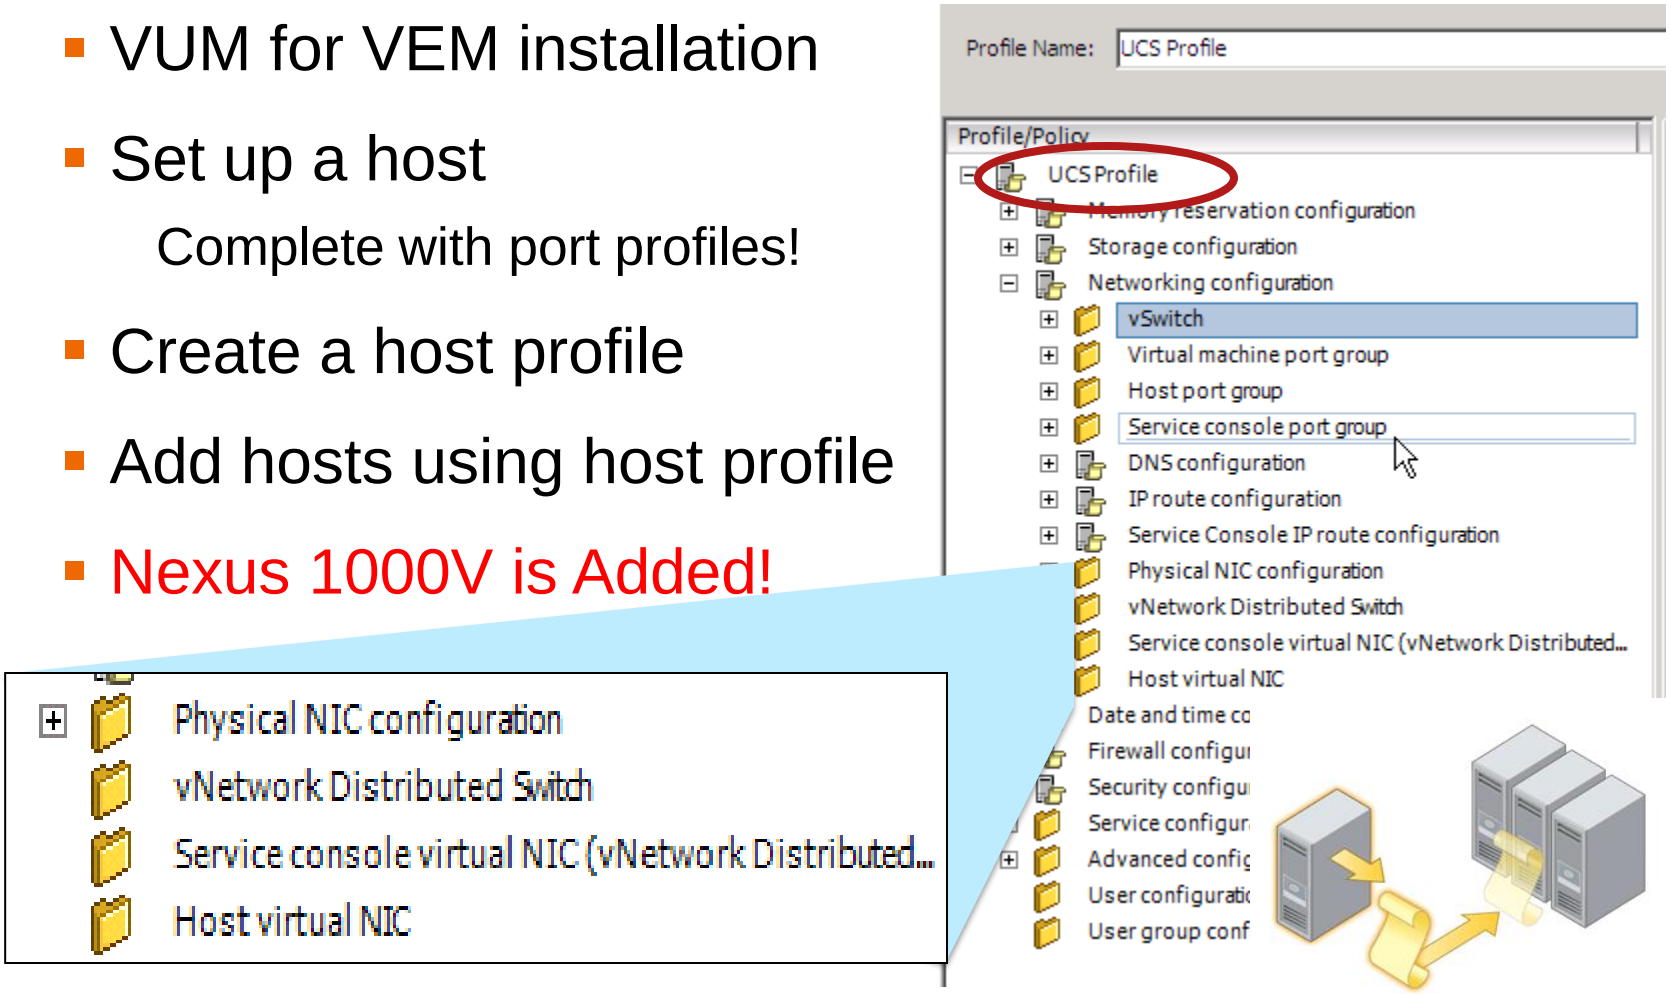

Images from VMware vSphere

Migrating to Nexus 1000V

- VMware vSwitch and Nexus 1000V can coexist
  - On the same servers
- Migrate some networks and VMs at first
  - Then migrate the rest later
- No need to shut down VMs to migrate
  - Simply use vCenter to change the port group
  - But will momentarily disrupt traffic
  - Very much like unplugging a cable and moving it to a port on another switch

Upgrading the Nexus 1000V Software

Keeping the Boundaries

3. Server admin upgrades VEMs

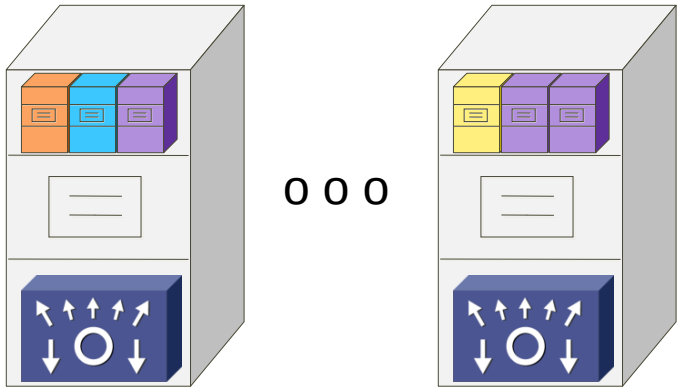

Server admin still owns the "hardware"

1. Network admin upgrades VSMs

2. VSM makes new VEM version available

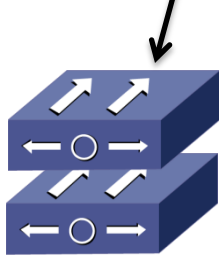

Upgrading the Nexus 1000V Software

Special Procedure for 1.3 to 1.4

1. Server admin upgrades VEMs

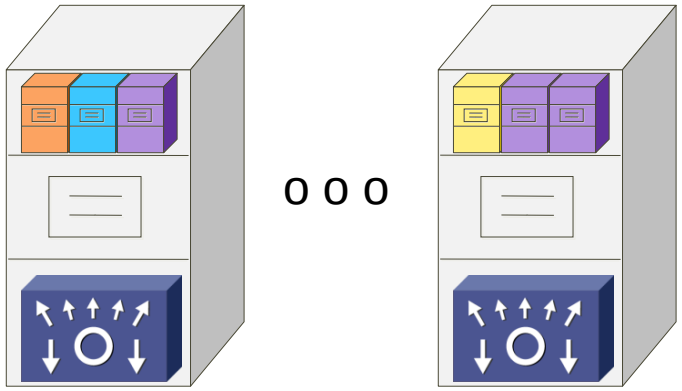

Server admin still owns the "hardware"

3. Network admin upgrades VSMs

2. 1.3 VSM works with new VEM version

4. 1.4 VSM enables new features

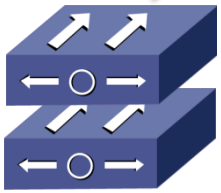

What's New with Nexus 1000V?

Version 1.4

New Features

- vPath for Virtual Services
  - VSG – firewall (more in BRKVIR-2011)
  - vWAAS
- Security features for Virtual Desktops (VDI)
- Long-Distance vMotion
  - Split Nexus 1000V between two locations using OTV
- Support for VMware vCloud Director (multi-tenancy)

Version 1.4

Enhanced Features

- Enhanced Scalability – 2K VLANs, 2K port profiles
- Quality of Service: Class-based Weighted Fair Queuing
  - Differentiated SLA via 8 pre-defined traffic classes
  - Queuing configured via Cisco modular QoS CLI (MQC)
- SPAN/ER-SPAN enhancements – now on a per port profile basis
- Simpler installation
- Enhancements for robustness, interoperability

vPath

Decision Caching for Smart Appliances

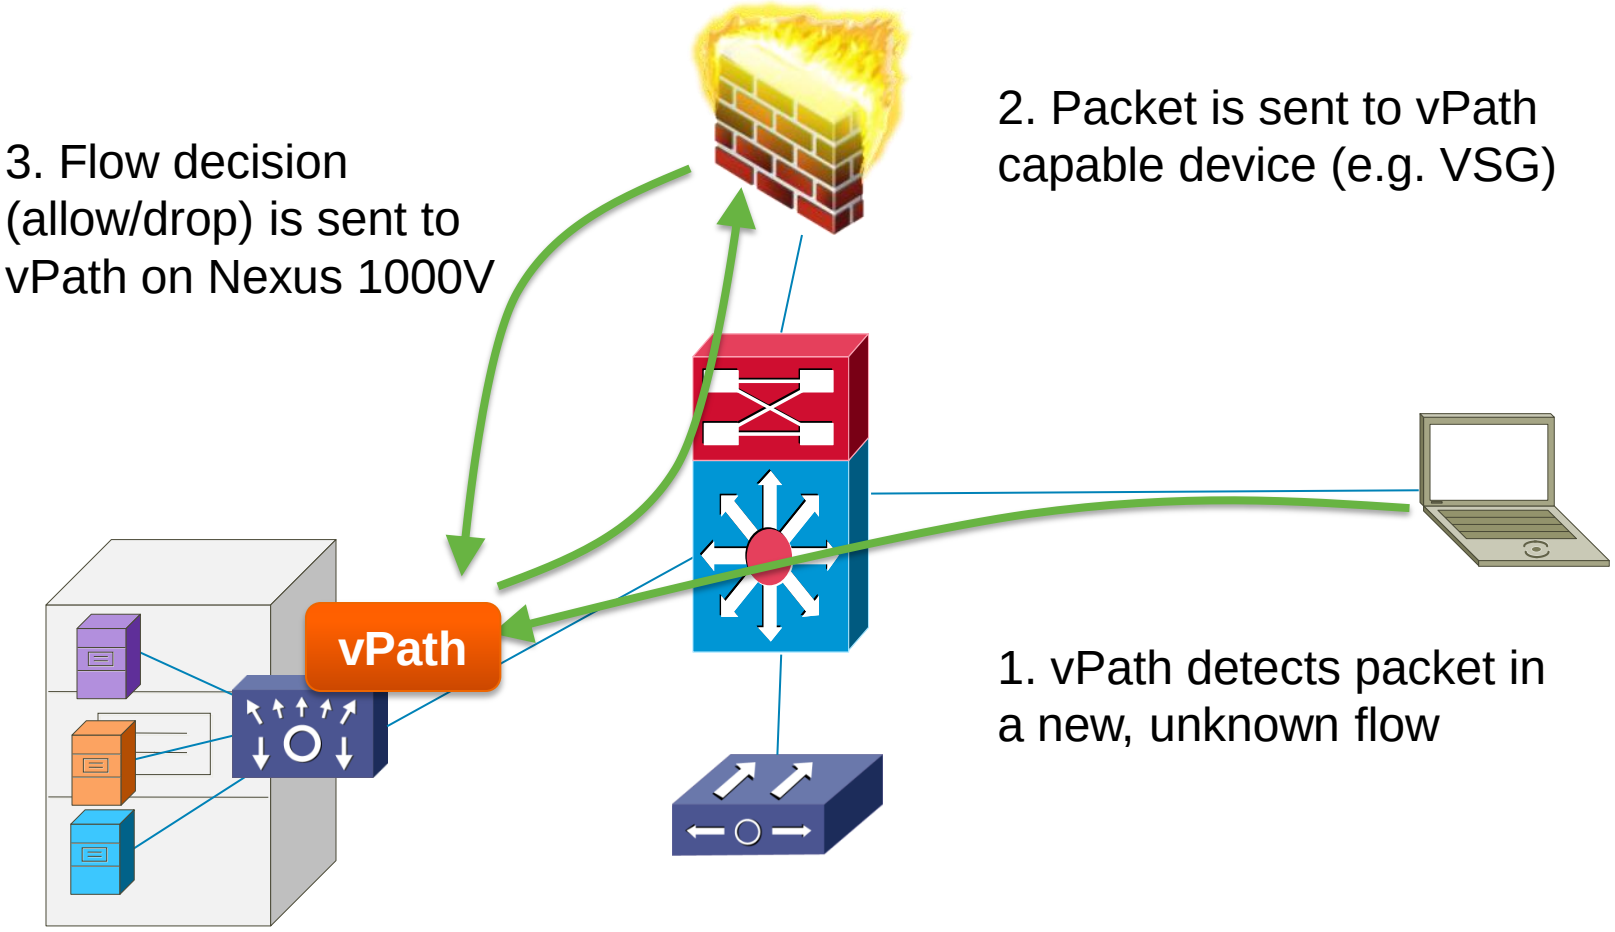

vPath

Decision Caching for Smart Appliances

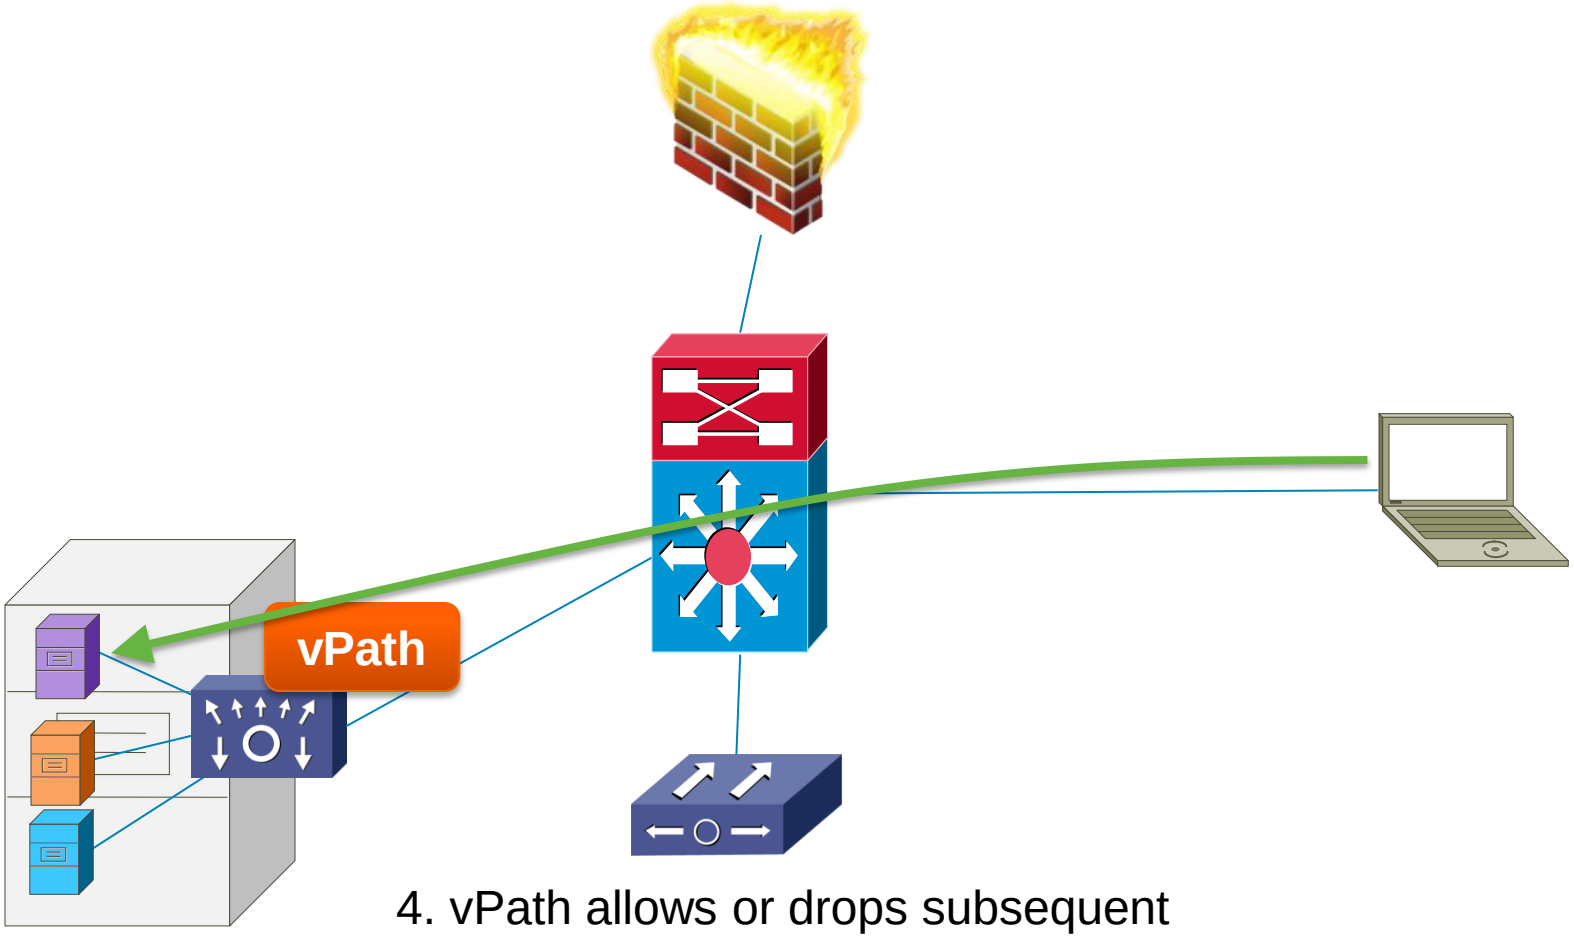

4. vPath allows or drops subsequent packets in flow without intervention

VDI Security Features

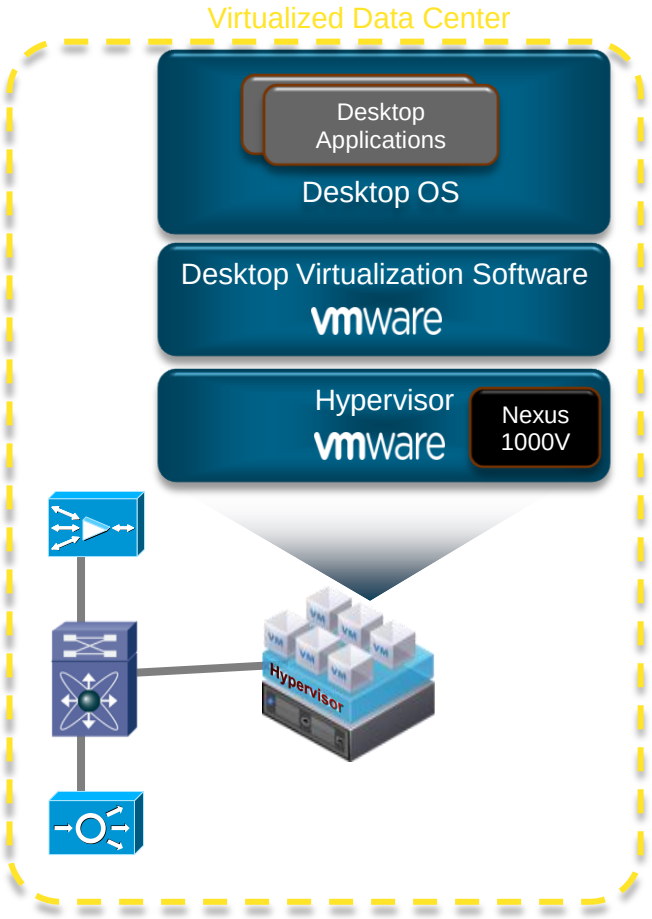

1000V Security Features for VDI

- Access Control List
- Port Security
- Private VLAN
- DHCP Snooping
- Dynamic ARP Inspection
- IP Source Guard

WAAS: Wide Area Application Service
ACE: Application Control Engine

Long Distance vMotion

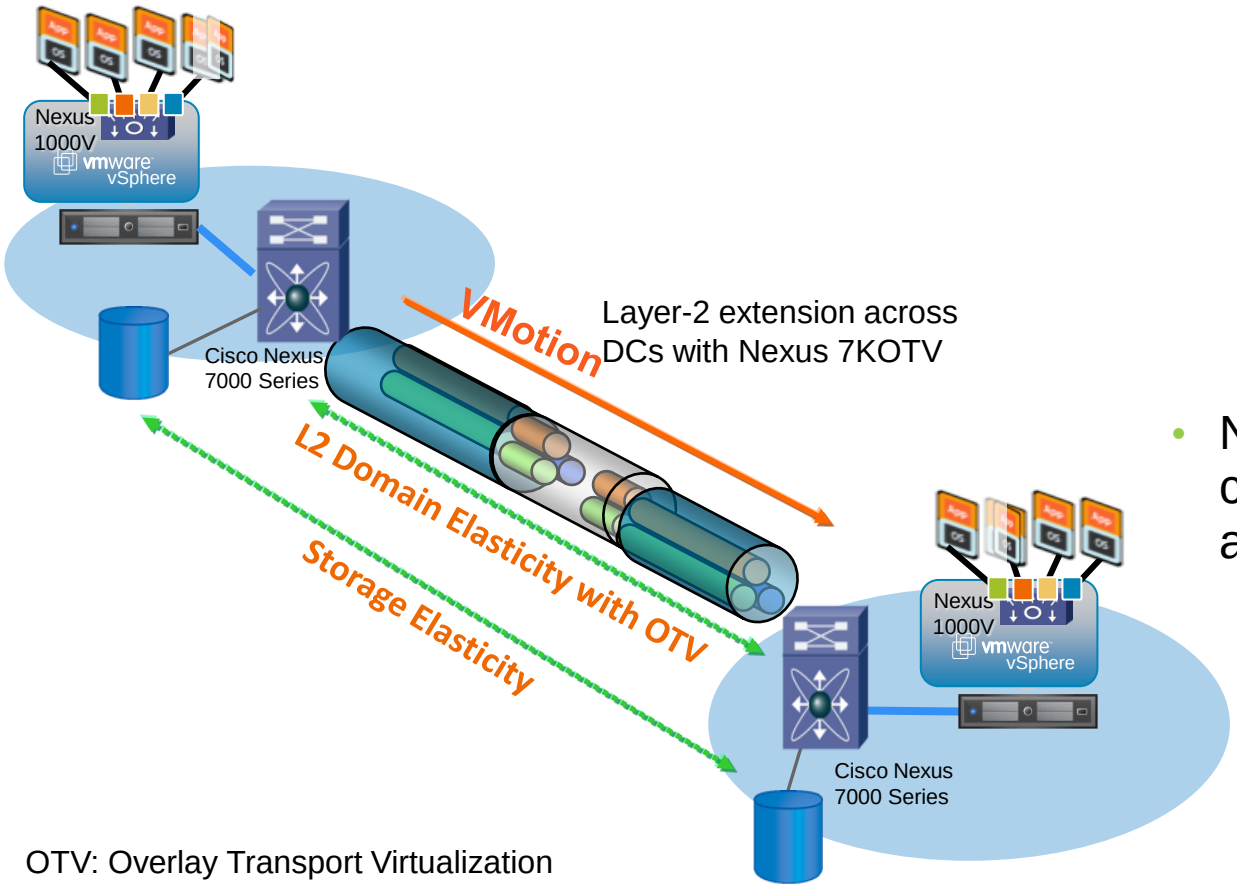

OTV: Overlay Transport Virtualization

- Network integrity is critical to long distance vMotion
  - Security
  - Quality of Service
  - Network Monitoring
  - Troubleshooting
- Nexus 1000V provides these critical network functions across data centers

vCloud Director and the Nexus 1000V

- Nexus 1000V

  Currently integrated in vCloud Director

  Provides VLAN isolation with Portgroup-backed network pools

  Cisco and VMWare consider Cisco Nexus 1000V an integral component of VMware's vSphere and vCloud product lines and are committed to delivering interoperable solutions for current and future versions of these products, including scalable network segmentation technologies.

- Virtual Security Gateway

  Cisco and VMware are partnering to integrate Cisco Virtual Security Gateway into VMware vCloud Director, leveraging the points of integration that vShield Manager provides.

Covered in BRKVIR-3005

Cisco Nexus 1000V Troubleshooting

- Installation
- VSM and VEM connectivity issues
- Port profiles
- Port channels
- VSM HA scenarios
- Upgrades
- Known issues

Related Presentations

BRKVIR-2002	Deploying Virtual Desktop Infrastructure (VDI)
BRKVIR-2008	UCS and Nexus 1000V Virtualization for IaaS Could
BRKVIR-2010	Cisco VXI: An End-to-End Virtualization System
BRKVIR-2011	Deploying Services in a Virtualized Environment
BRKVIR-2931	End-to-End Data Center Virtualization
BRKVIR-3013	Deploying and Troubleshooting the Nexus 1000V Virtual Switch

Keeping Your Sanity

- Virtualization brings new networking challenges
  - As well as opportunities, with more of both to come!
- Nexus 1000V brings **sanity** to the virtual network
  - Clear, appropriate administrative boundaries
  - Port profiles for deployment, scaling and motion
- Nexus 1000V preserves **freedom** for servers
  - Safely deploy hosts and VMs independently
- How to keep your own sanity
  - Follow the configuration guides for Nexus 1000V
  - Don’t panic —use the troubleshooting presentation

Complete Your Online Session Evaluation

- Receive 25 Cisco Preferred Access points for each session evaluation you complete.
- Give us your feedback and you could win fabulous prizes. Points are calculated on a daily basis. Winners will be notified by email after July 22nd.
- Complete your session evaluation online now (open a browser through our wireless network to access our portal) or visit one of the Internet stations throughout the Convention Center.
- Don’t forget to activate your Cisco Live and Networkers Virtual account for access to all session materials, communities, and on-demand and live activities throughout the year. Activate your account at any internet station or visit www.ciscolivevirtual.com.

Learn. Connect. Collaborate. together.

Visit the Cisco Store for Related Titles

http://theciscostores.com

Thank you.

References

Resources

- CCO Links

  1000V: www.cisco.com/go/1000v

  1010: www.cisco.com/go/1010

  VSG: www.cisco.com/go/vsg

  VNMC: www.cisco.com/go/vnmc

- My Cisco Community

  https://communities.cisco.com/community/technology/datacenter/nexus1000v

- Deployment Guides

  Nexus 1000V Deployment Guide

  Nexus 1000V on UCS – Best Practices

  Nexus 1010 Deployment Guide

  VSG Deployment Guide

- White paper

  Nexus 1000V and vCloud Director

Validated Designs

- vBlock with Nexus 1000V
- FlexPOD with Nexus 1000V and Nexus 1010
- Virtual Multi-tenant Data Center with Nexus 1000V
- Virtual Desktop
  - 1000V and VMware View
  - 1000V and VSG in VXI Reference Architecture

Common Virtual NIC Questions

Do I Need Redundant Virtual NICs?

Failover is Always Good, Right?

- Virtual NICs do not need redundancy
  - It's the same bus, memory and software either way
  - Duplicating adds complexity without value

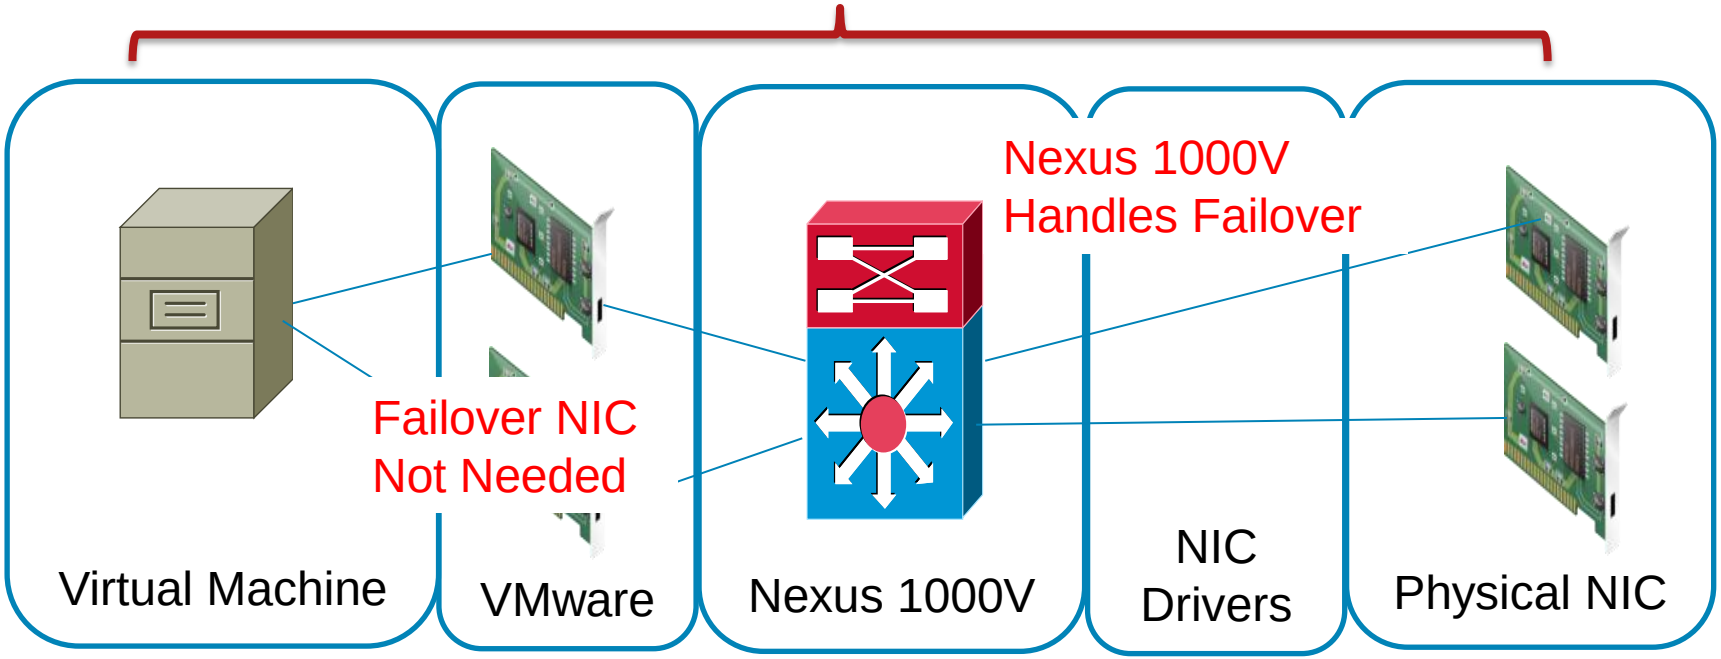

Do I need to Load Balance Virtual NICs?

But What About Performance?

- Speed and duplex do not matter for virtual ports

  Bandwidth not limited by 100 Mb, 1G, 10G, etc.

- No NIC teaming needed on the guest OS

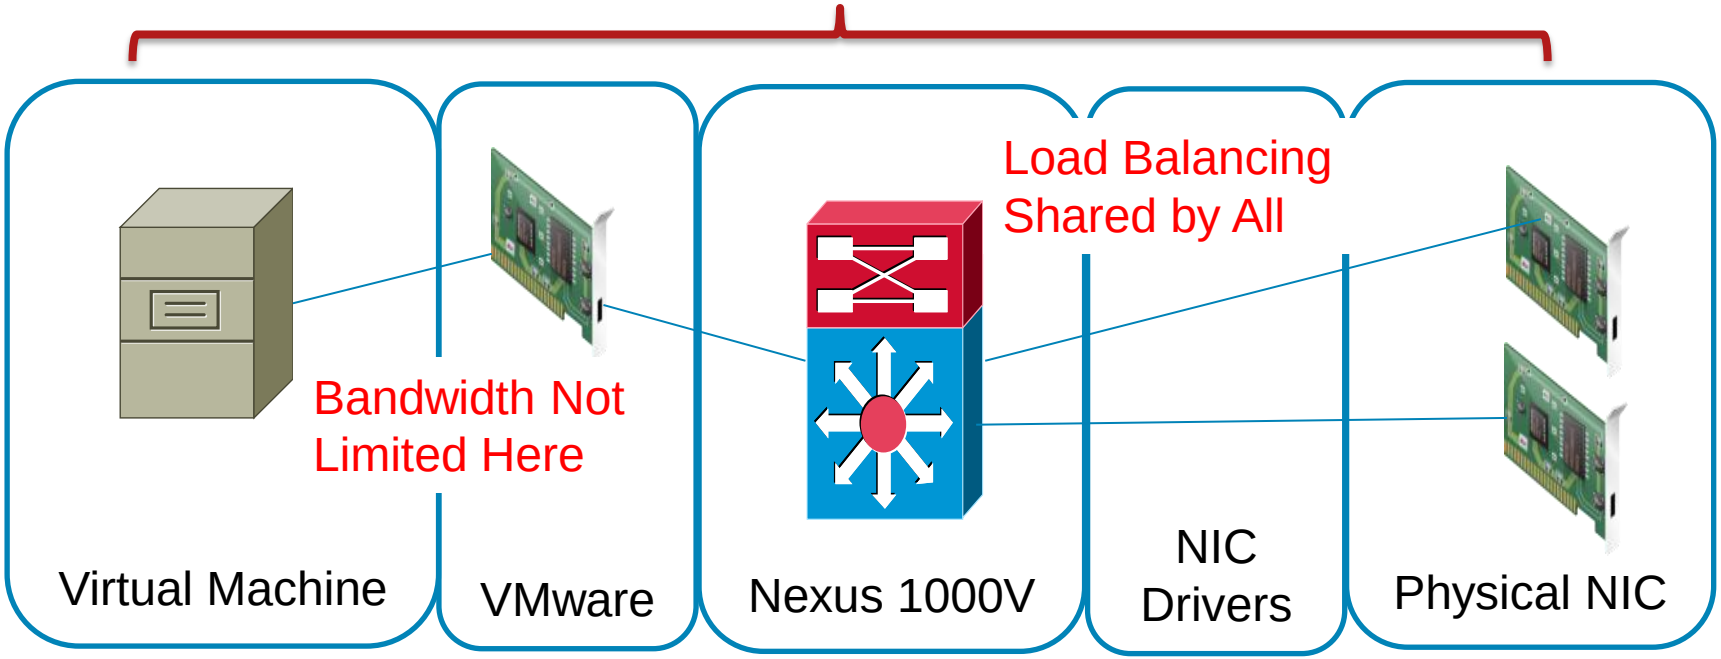

Do I Need Trunk Ports for Two VLANs?

Only for Special Cases

- Trunk ports harder to configure and control on VM
- Virtual wires cheap—up to 10 access ports per VM
- Expensive physical connections shared by 1000V

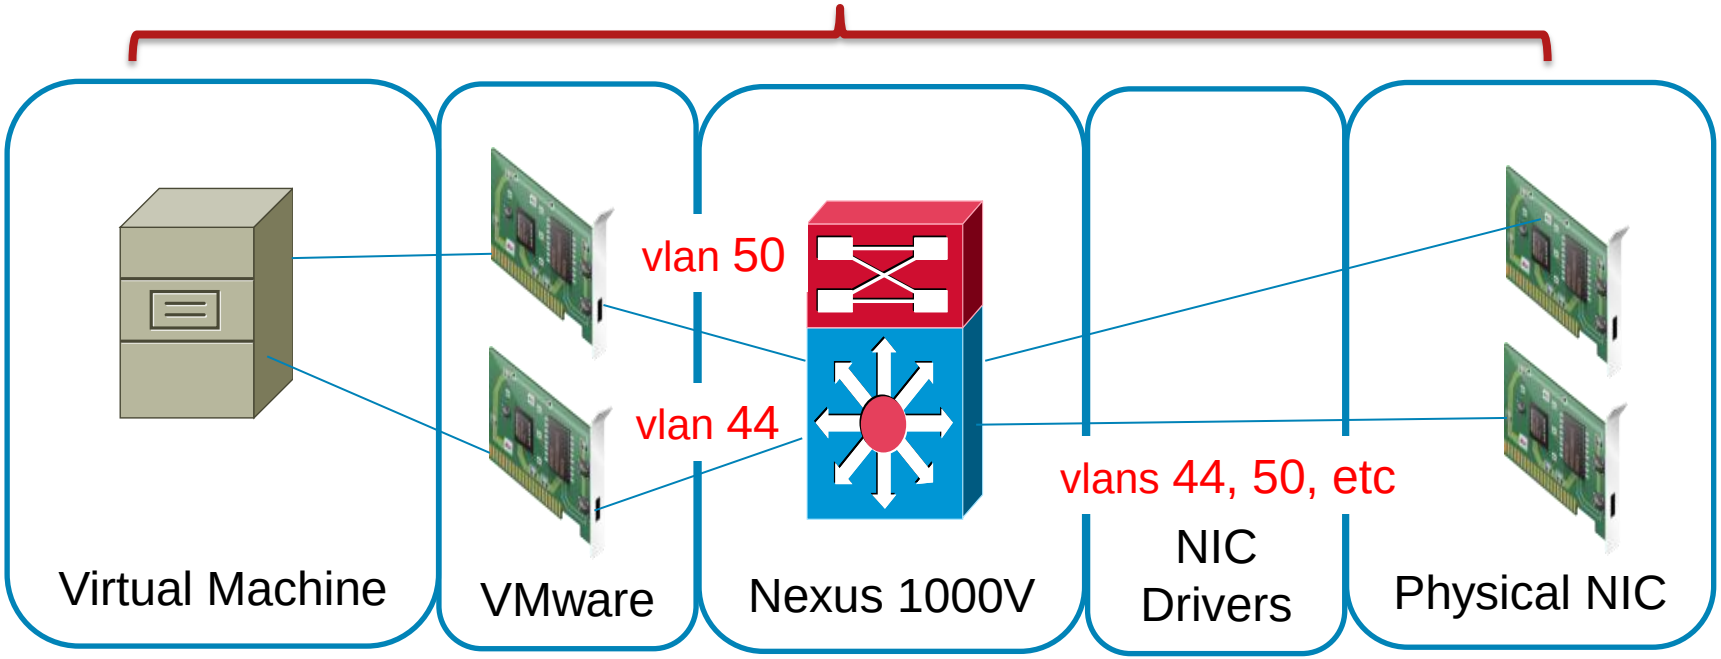

Simplified Networking for VMs

- No teaming drivers on the guest OS
  - Teaming drivers are usually NIC vendor-dependent
  - Nexus 1000V can port channel across any set of NICs
    - Failover
    - Load balancing
- Drivers for trunk ports are not normally needed
  - With very few exceptions
  - Easier to change and maintain VLANs on access ports
  - Guest OS and applications unaware of VLANs

vCloud Director and the Nexus 1000V

- Nexus 1000V

  Currently integrated in vCloud Director

  Provides VLAN isolation with Portgroup-backed network pools

  Cisco and VMWare consider Cisco Nexus 1000V an integral component of VMware's vSphere and vCloud product lines and are committed to delivering interoperable solutions for current and future versions of these products, including scalable network segmentation technologies.

- Virtual Security Gateway

  Cisco and VMware are partnering to integrate Cisco Virtual Security Gateway into VMware vCloud Director, leveraging the points of integration that vShield Manager provides.