



CONNECTING THE  
IT COMMUNITY

**INTEROP**<sup>®</sup>

# vSphere Networking for the Network Admin

Jason Nash, Varrow CTO



# Agenda

- What is virtualization?
- How does virtualization impact the network team?
- How should you approach virtualization?
- vSphere 101
- Where is the demarcation point?
- Integrating vSphere in to the physical network
- Common questions



# What is Virtualization?

In computing, **virtualization** refers to the act of creating a virtual (rather than actual) version of something, including (but not limited to) a virtual computer hardware platform, operating system (OS), storage device, or computer network resources.



# Why You Should Care

Normally Network Administrators don't get too involved in end-server configuration.



# Virtualization Changes Things

Server teams will need to collaborate more closely with the network teams.



# Where Does It End?

That doesn't mean you own the hypervisor, but you may own some of the configuration.



# Network Team Demarcation

- Usually the demarcation for the network team is the ToR switch
  - Using integrated vSphere Networking
  - Or VMware NSX
- But the network team may be asked to manage 3<sup>rd</sup> party virtual switches
  - Cisco Nexus 1000v



# VMware Products

- **vSphere** – Hypervisor that runs on physical servers
- **NSX** – Software Defined Networking platform that offers network abstraction and software services
- **vReal Automation Center (vRAC)** – Cloud Management Platform used to automate VM and app deployments





# vSphere Networking 101

- vSphere is the hypervisor that sits on the hardware server and hosts virtual machines (VMs)
- Each vSphere host has its own internal virtual switching system
- Each VM has one or more virtual NICs (vNICs) connected to that switch



# vSphere Network Terms

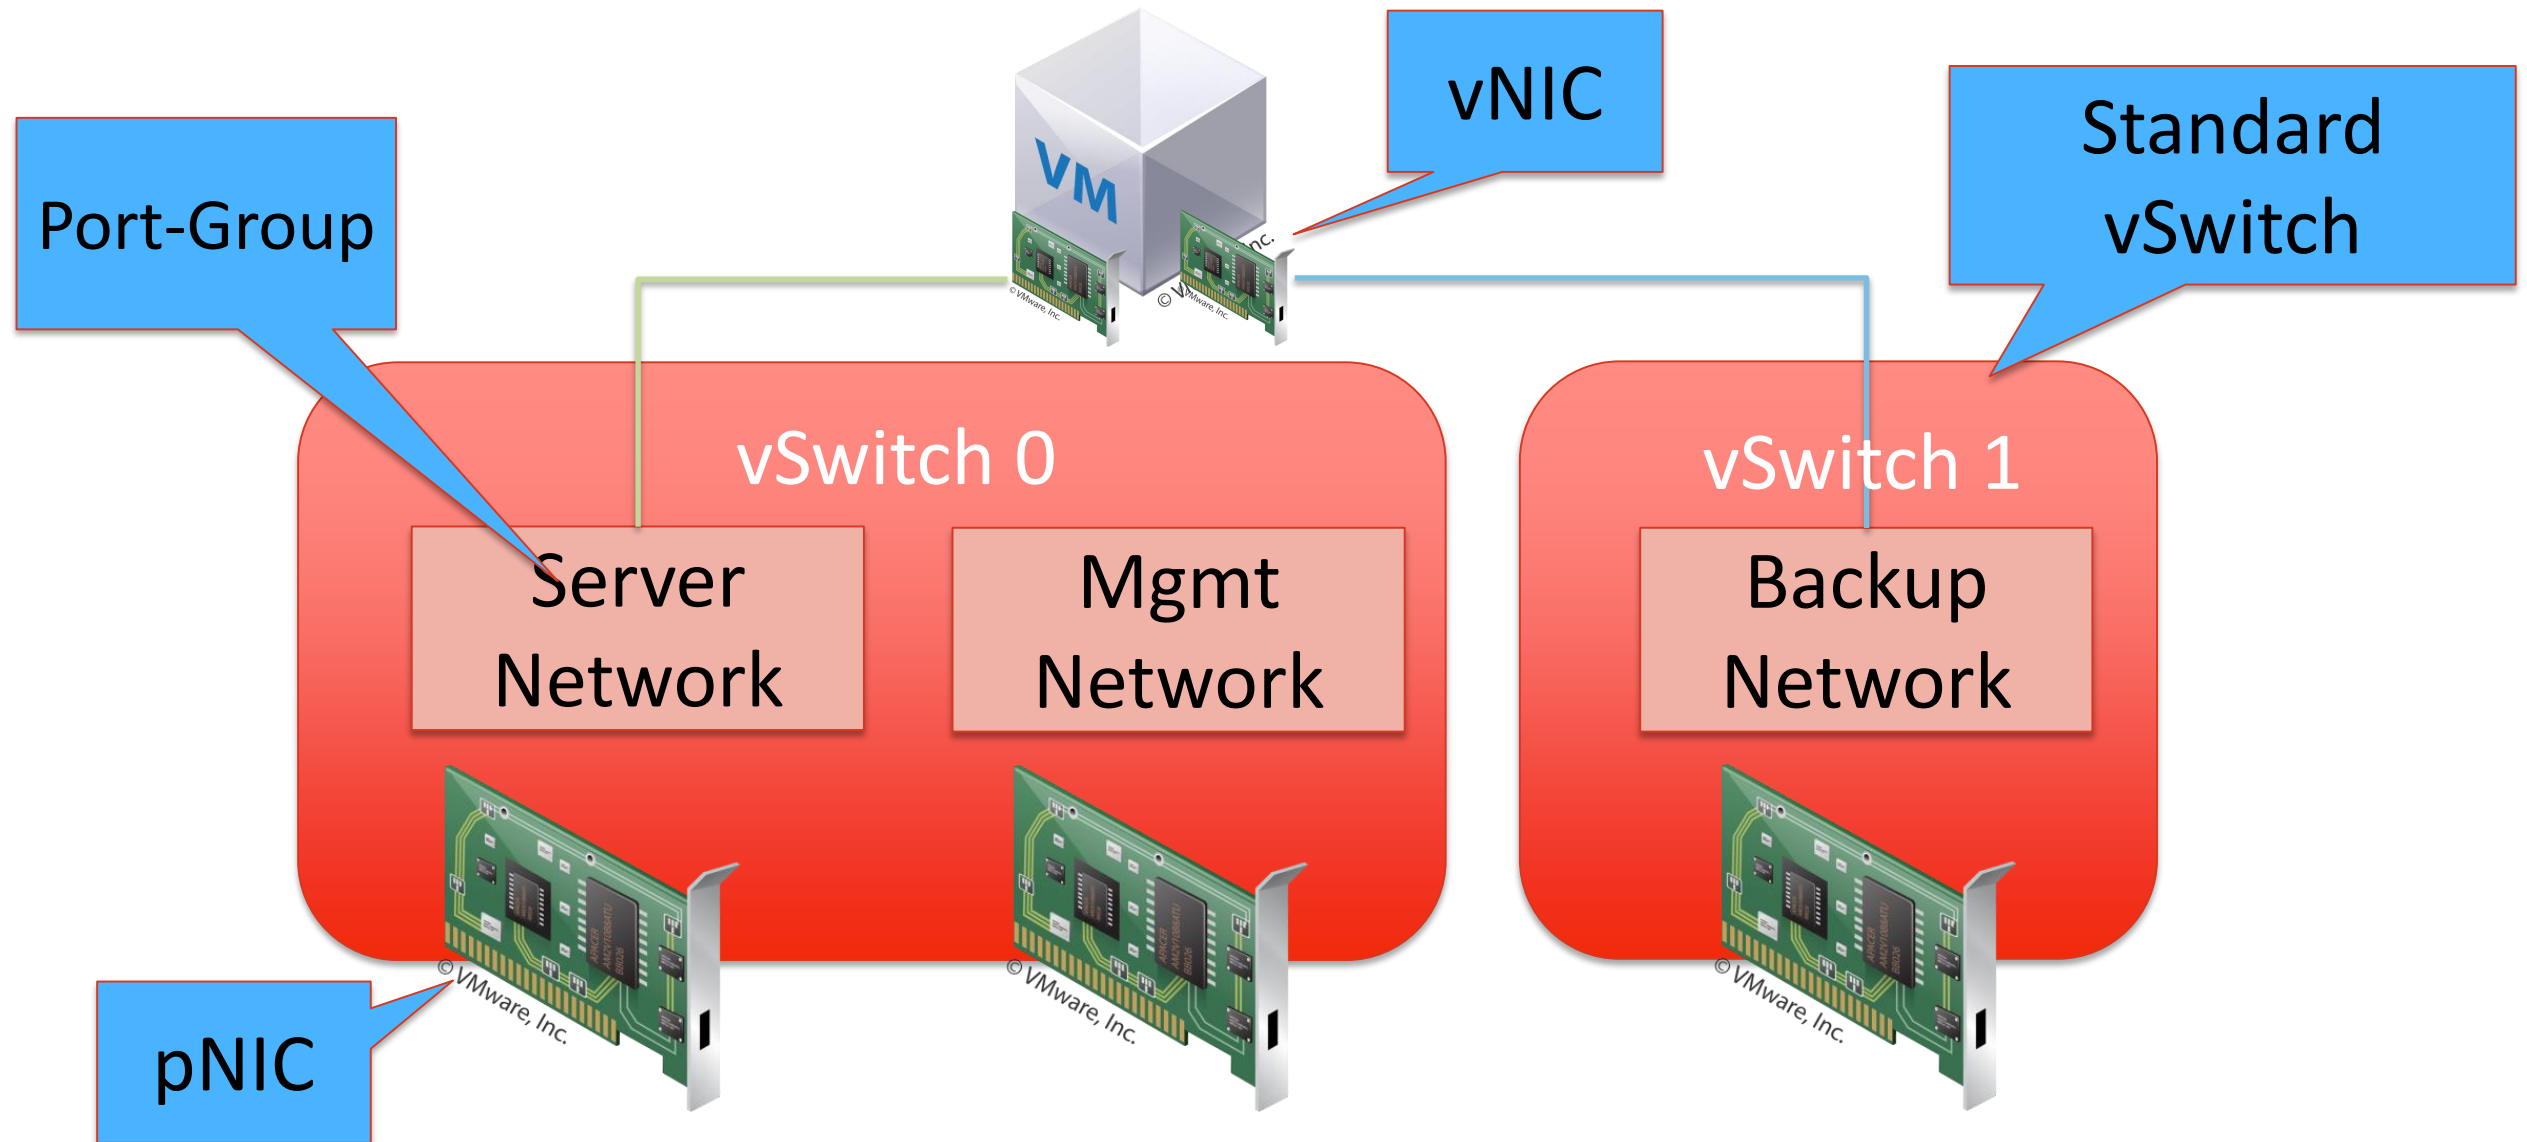
- **vSwitch** – Virtual Switch on each vSphere host
- **vNIC** – Virtual NIC(s) that each guest VM has
- **pNIC** – Physical NICs in the servers that connect the host to the network
- **Port-Group** – Think of them as a virtual network hub
- **Distributed Switch (VDS)** – Shared virtual switch available in Enterprise Plus licensing
- **VMKernel** - Virtual management interface(s) on each vSphere host



# Is a vSwitch a Switch?

- YES! Virtual Switches (vSwitch) are real, but...
  - No Spanning Tree
  - Performs its own loop detection
- It switches frames based on MAC addresses
- Understands VLAN tags
- Can do QoS services and tagging

# Components of a vSwitch Architecture





# Traffic Types – VM Traffic

- Network traffic to/from the virtual machine applications and guest OS
- Characteristics depend greatly on the application
  - vSphere doesn't change this
- Can be distributed using teaming methods
- Often on dedicated NICs



# Traffic Types - Management

- Management traffic to control the vSphere host
  - SSH and vCenter
- Low bandwidth and pretty lenient latency requirements
- Not normally distributed over multiple NICs
  - But should do at least Active/Passive teaming



# Traffic Types - vMotion

- Used to move live servers between hosts
- Very bursty and latency impacts application performance
  - Intra-DC is fine
  - Can go between DCs up to 100ms now
- Keep this traffic as local as possible
  - Is supported to route it under some instances





# Traffic Types – Fault Tolerance

- Used to create mirror VMs on other hosts
  - Kept in complete lock-step for instant failover
- FT traffic is high bandwidth, constant, and latency sensitive
- Dedicate NICs and keep traffic local
  - Some configurations require 10Gb





# Traffic Types – IP Storage

- Same as any other NFS or iSCSI traffic
  - iSCSI seems to be more finicky
- Dedicate NICs and watch for latency and drops
  - May dictate an upgrade in switches
- Can route NFS/iSCSI under some scenarios
  - But try to be as direct as possible



# How Many VLANs?

- Like most things...it depends.
- One or more VLANs for VM traffic
  - Depends on the logical segmentation
- Dedicated VLANs for other types
  - Management
  - vMotion
  - FT
  - IP Storage



# VLAN Span

- Traditionally all VLANs spanned racks for all vSphere hosts in a cluster
  - May have multiple clusters with different requirements
- Seeing more Leaf/Spine deployments
  - Rack is Layer 2 boundary
  - Routing of vMotion/IP Storage
- Gets in to a deeper discussion



# What About Physical Connectivity?

- There are a number of things that dictate the design
  - 1Gb or 10Gb?
  - IP storage or FC storage?
  - Hardware switches, how many and what kind?
  - VM traffic requirements
- For a proper design you need to understand these requirements and constraints
- Make your configuration scalable



# Connectivity Best Practices

- If traffic types will be mixed look at the profile of that traffic
  - Bursty or sustained?
  - High or low?
  - Common to share NICs with vMotion and Management
- Suggested to never share NICs with IP storage
- Fault Tolerance ~~really wants~~ requires dedicated NICs
- Management traffic is very low, can be shared
  - Often with vMotion on a different VLAN



## Topology Suggestions

- The final configuration depends on how many NICs you have per server
  - Common to see 6 or 8 Gb NICs, have seen up to 12
- Utilizing FT and/or IP storage will add additional requirements
  - Both should have dedicated NICs
- Recommended to have 6 NICs for Fibre Channel storage or 8 NICs for IP (iSCSI/NFS) storage



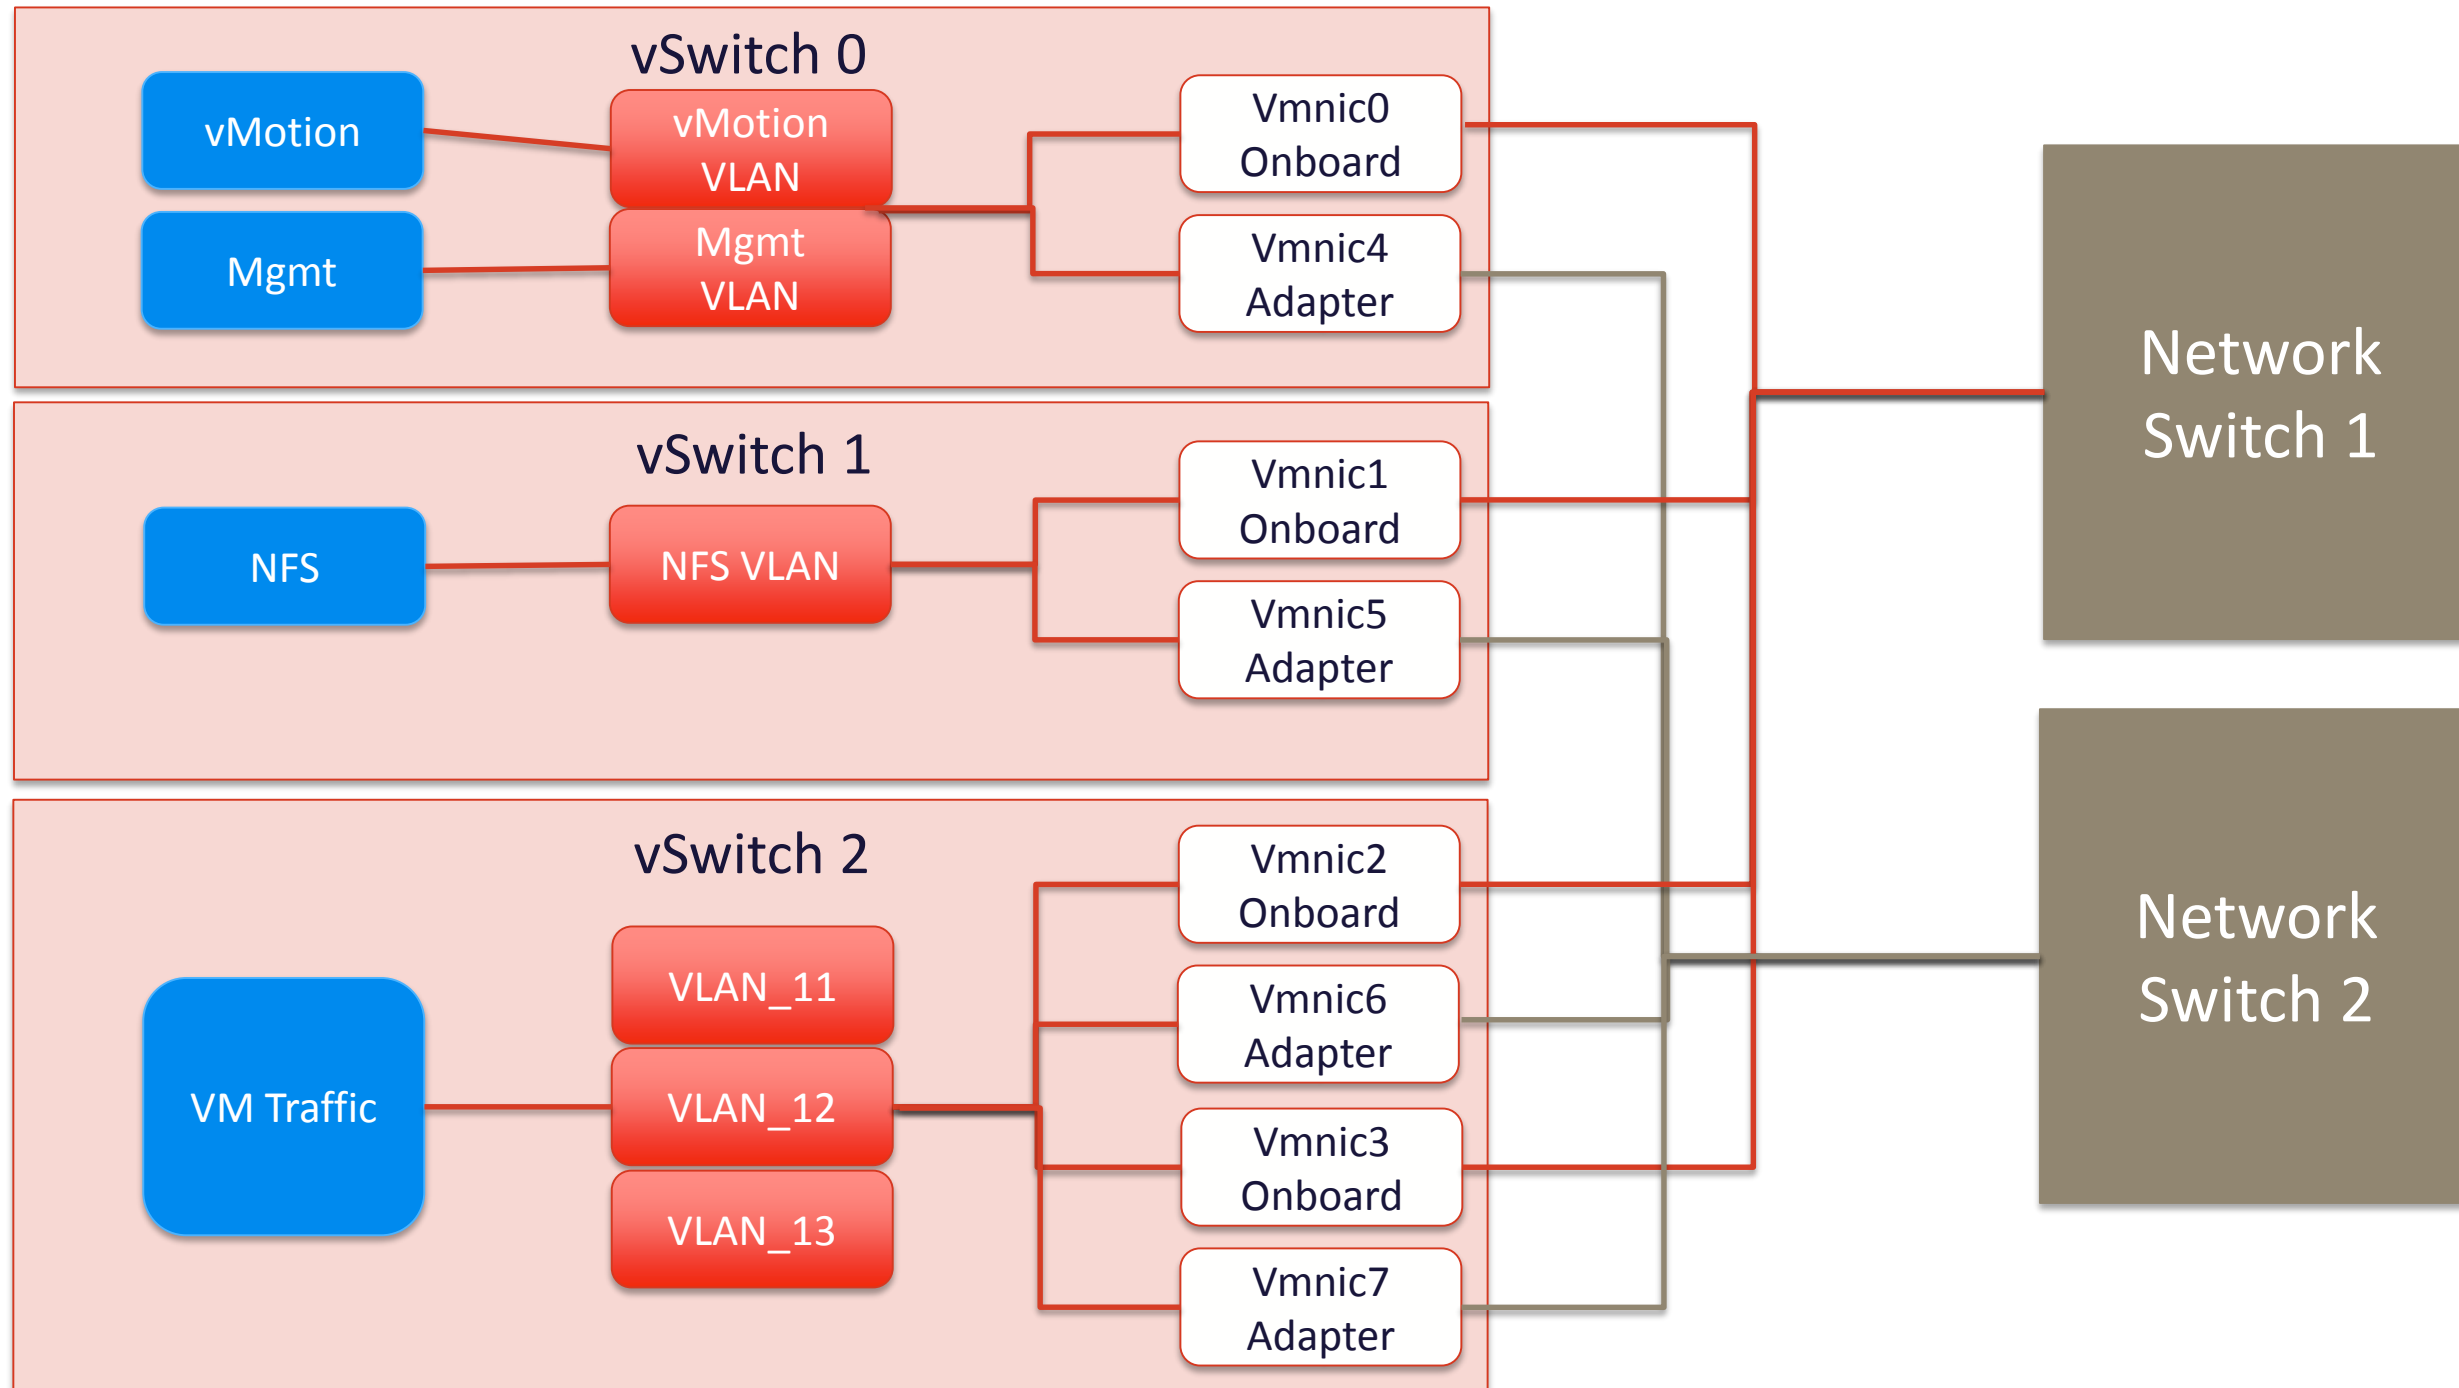


## A Bigger Hammer – 10Gb Ethernet

- Utilizing 10Gb greatly reduces NIC requirements
  - Commonly see just two NICs per host
- Be careful though, with vSphere 4.1 and newer vMotion can use up to 8Gb
- Recommended to use Network I/O Control (NIOC) or other method for traffic shaping with 10Gb
  - Without some form of shaning VMs may experience I/O starvation











# Physical Switch Configuration

- Your physical switch configuration must match the vSwitch configuration
  - VLANs
  - Load-balancing mechanism (discussed later)
- Should also tell the switch that the vSphere system is an “end host” or connected to an “edge port” and not another switch
  - Greatly reduces or eliminates Spanning-Tree delay
  - Tells the port this device does not participate in Spanning-Tree
  - Want failover to be instant



# Load-Balancing and Teaming

- Another very common area of confusion is load-balancing and teaming
- First, there is no load-balancing mechanism that truly balances load
  - Why I often call it load-distribution instead of balancing
- Which method you choose depends on several things
  - Type of physical switches
  - Traffic profile going to/from your VMs
  - VMware license level
  - vSwitch type



# Simple Team Options

- Simplest are
  - Hash by Virtual Port ID
  - Hash by MAC Address
- No configuration needed on switch
  - No LACP, EtherChannel, etc
- Virtual Port ID is default for a reason



# IP Hashing

- Simple forced port-channel
  - Configure your switches
- Hashes on source/destination IP
  - Only option available
- Useful to balance traffic from one VM across multiple pNICs



# LACP

- LACP was added in recent versions of vSphere
  - Much improved in v5.5
- Bundle pNICs together and designate one of the 23 supported algorithms
  - Just match it on the physical side
- Works fine with stacks and vPCs



# Load-Based Teaming

- Requires Enterprise Plus licensing and VDS
- Default VM/pNIC placement is Virtual Port ID
- If a pNIC is 75% utilized for 30 seconds VMs are moved
  - Positioned to the least loaded pNIC
- Does not distribute single “conversation” over multiple pNICs



# Network I/O Control (NIOC)

- Feature in the VDS that lets you define priority, shares, and limits to network traffic
  - Designated by traffic types and port-groups
- Easy way to manage 10Gb+ interfaces
  - So that vMotion/FT don't cause I/O starvation
- Can also tag outgoing traffic with 802.1p tag



# Monitoring

- VDS supports NetFlow
  - V10/IPFIX
- Also supports SPAN/ERSPAN
  - Call Port Mirroring
- Many network monitoring tools understand vSphere vSwitches





# The Future

- Software Defined Networking (SDN) rapidly maturing
  - VMware's offering is NSX
- Allows for
  - Distributed Firewalling
  - Distributed Routing
  - Edge Services
  - Network Overlay
- Heavily used in cloud/automation systems

# Questions?



Jason Nash  
[jason@varrow.com](mailto:jason@varrow.com)  
@TheJasonNash  
<http://www.jasonnash.com>

Feel free to reach out after the session for any questions.