

# Networking Virtual Machines

Jon Hall  
Technical Trainer

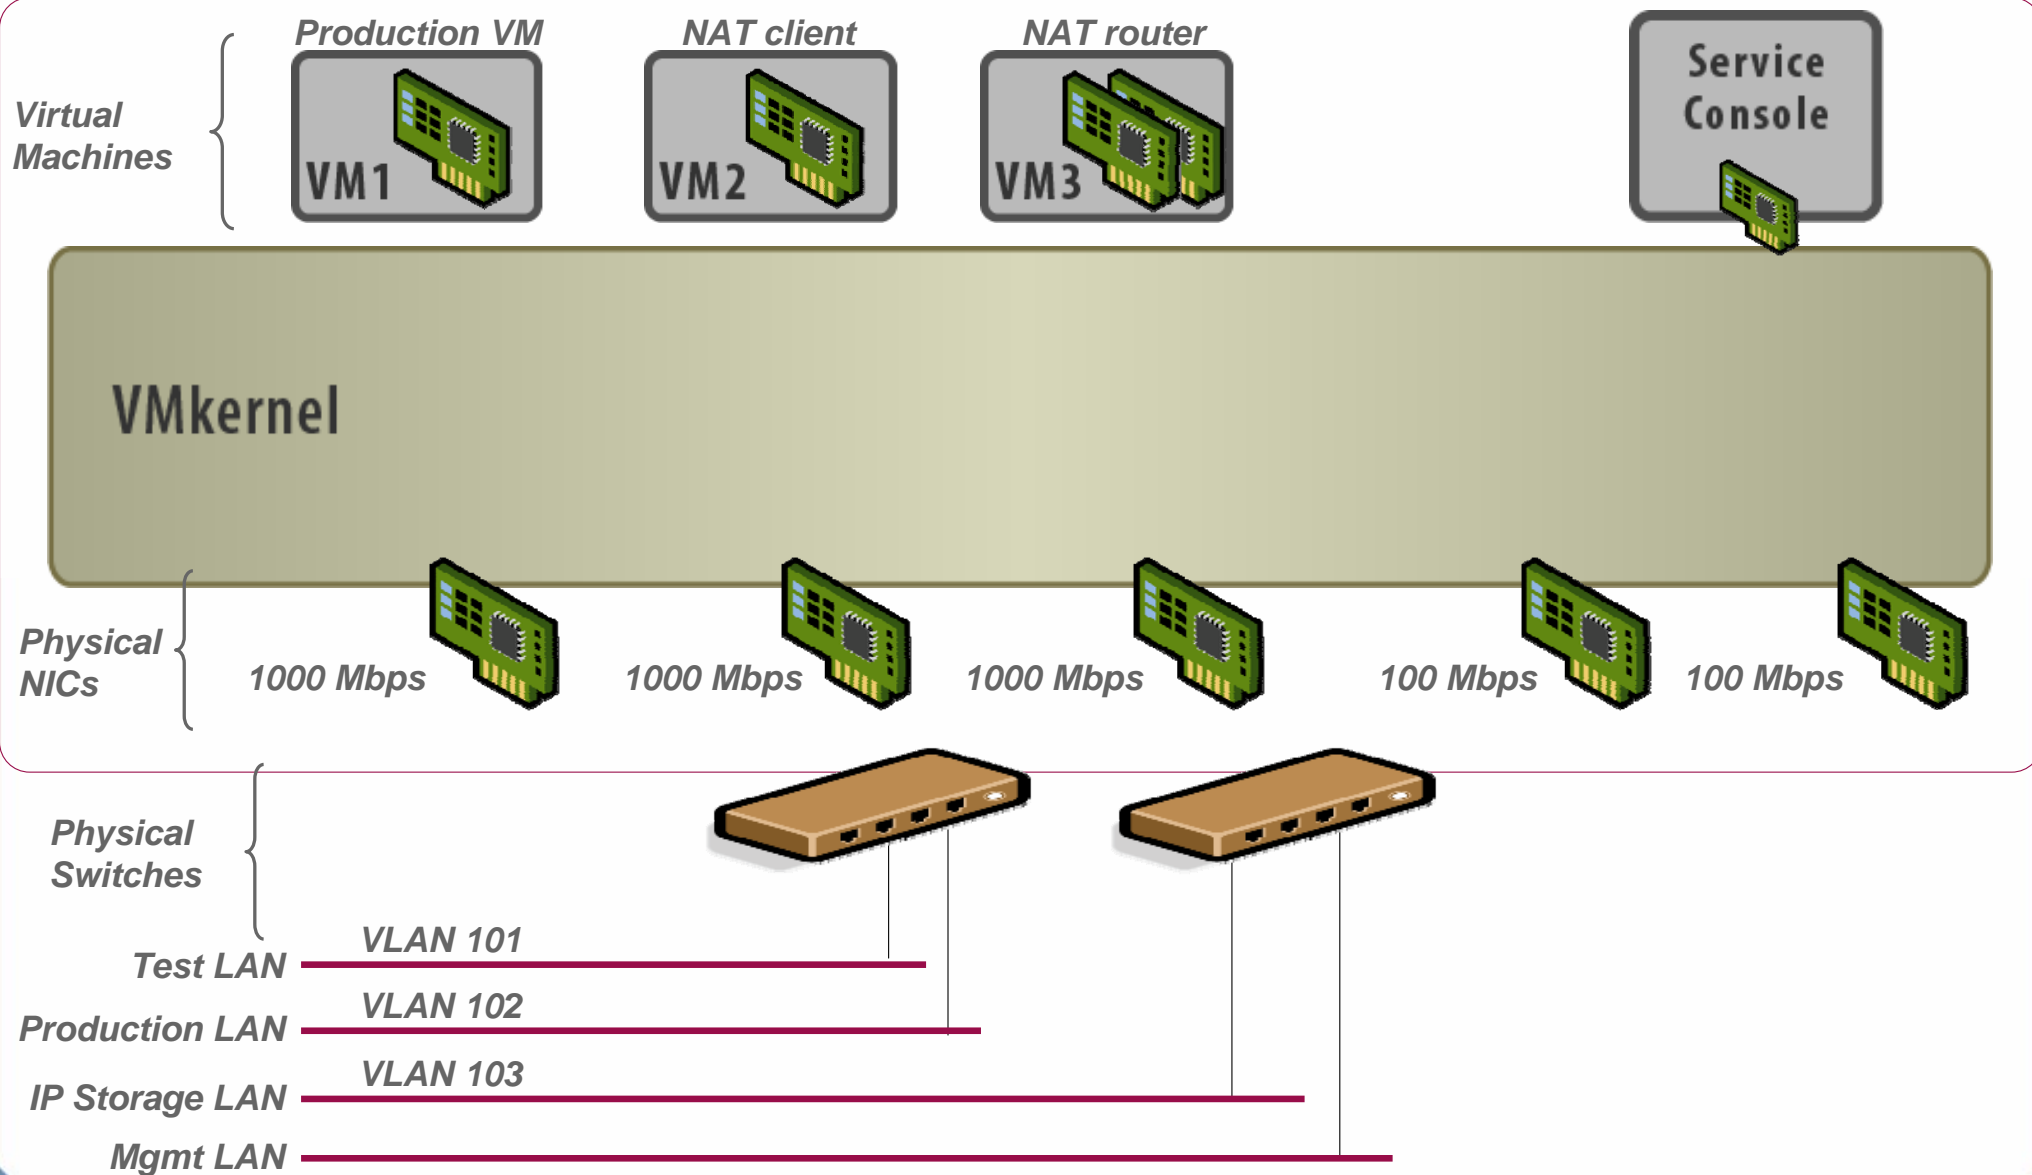


**VMWORLD 2006**

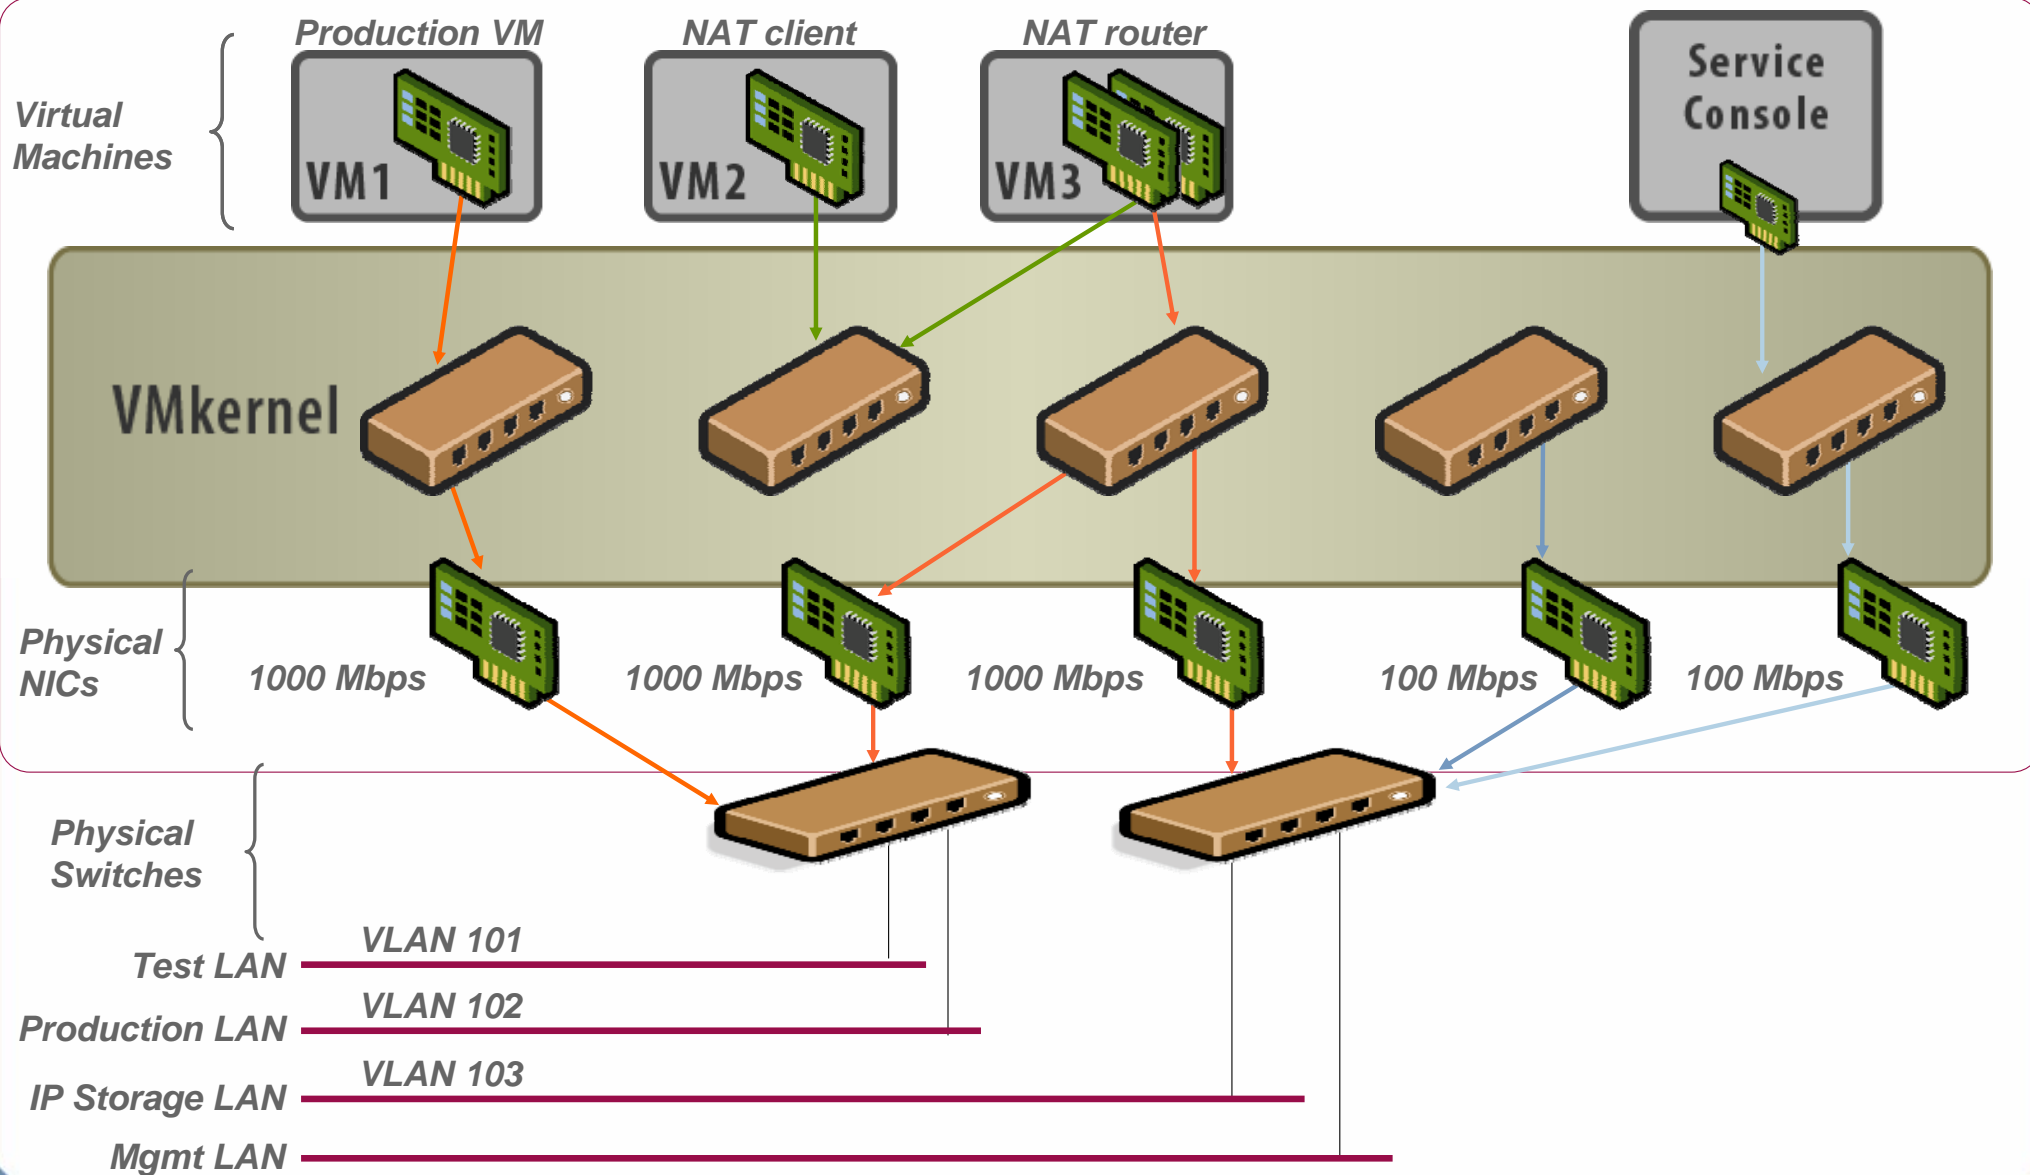
# Agenda

- Introduction
- New networking features
- Networking virtual machines
  - > Virtual Switch Connections
  - > Port Group Policies
- Networking IP Storage
  - > iSCSI
  - > NAS

# A networking Scenario




# A Networking Scenario




## vSwitch - No Physical Adapters (Internal Only)

- Each switch is an internal LAN, implemented entirely in software by the VMkernel
- Provides networking for the VMs on a single ESX Server system only
- Zero collisions
- Up to 1016 ports per switch
- Traffic shaping is not supported




## vSwitch - One Physical Adapter

- Connects a virtual switch to one specific physical NIC
- Up to 1016 ports available
  - Zero collisions on internal traffic
- Each Virtual NIC will have its own MAC address
- Outbound bandwidth can be controlled with traffic shaping




# Combining Internal And External vSwitches



- Virtual switch with one outbound adapter acts as a DMZ
- Back-end applications are secured behind the firewall using internal-only switches

## vSwitch – Multiple Physical Adapters (NIC Team)

- Can connect to an 802.3ad NIC team
- Up to 1016 ports per switch
  - Zero collisions on internal traffic
- Each Virtual NIC will have its own MAC address
- Improved network performance by network traffic load distribution
- Redundant NIC operation
- Outbound bandwidth can be controlled with traffic shaping






# Connections



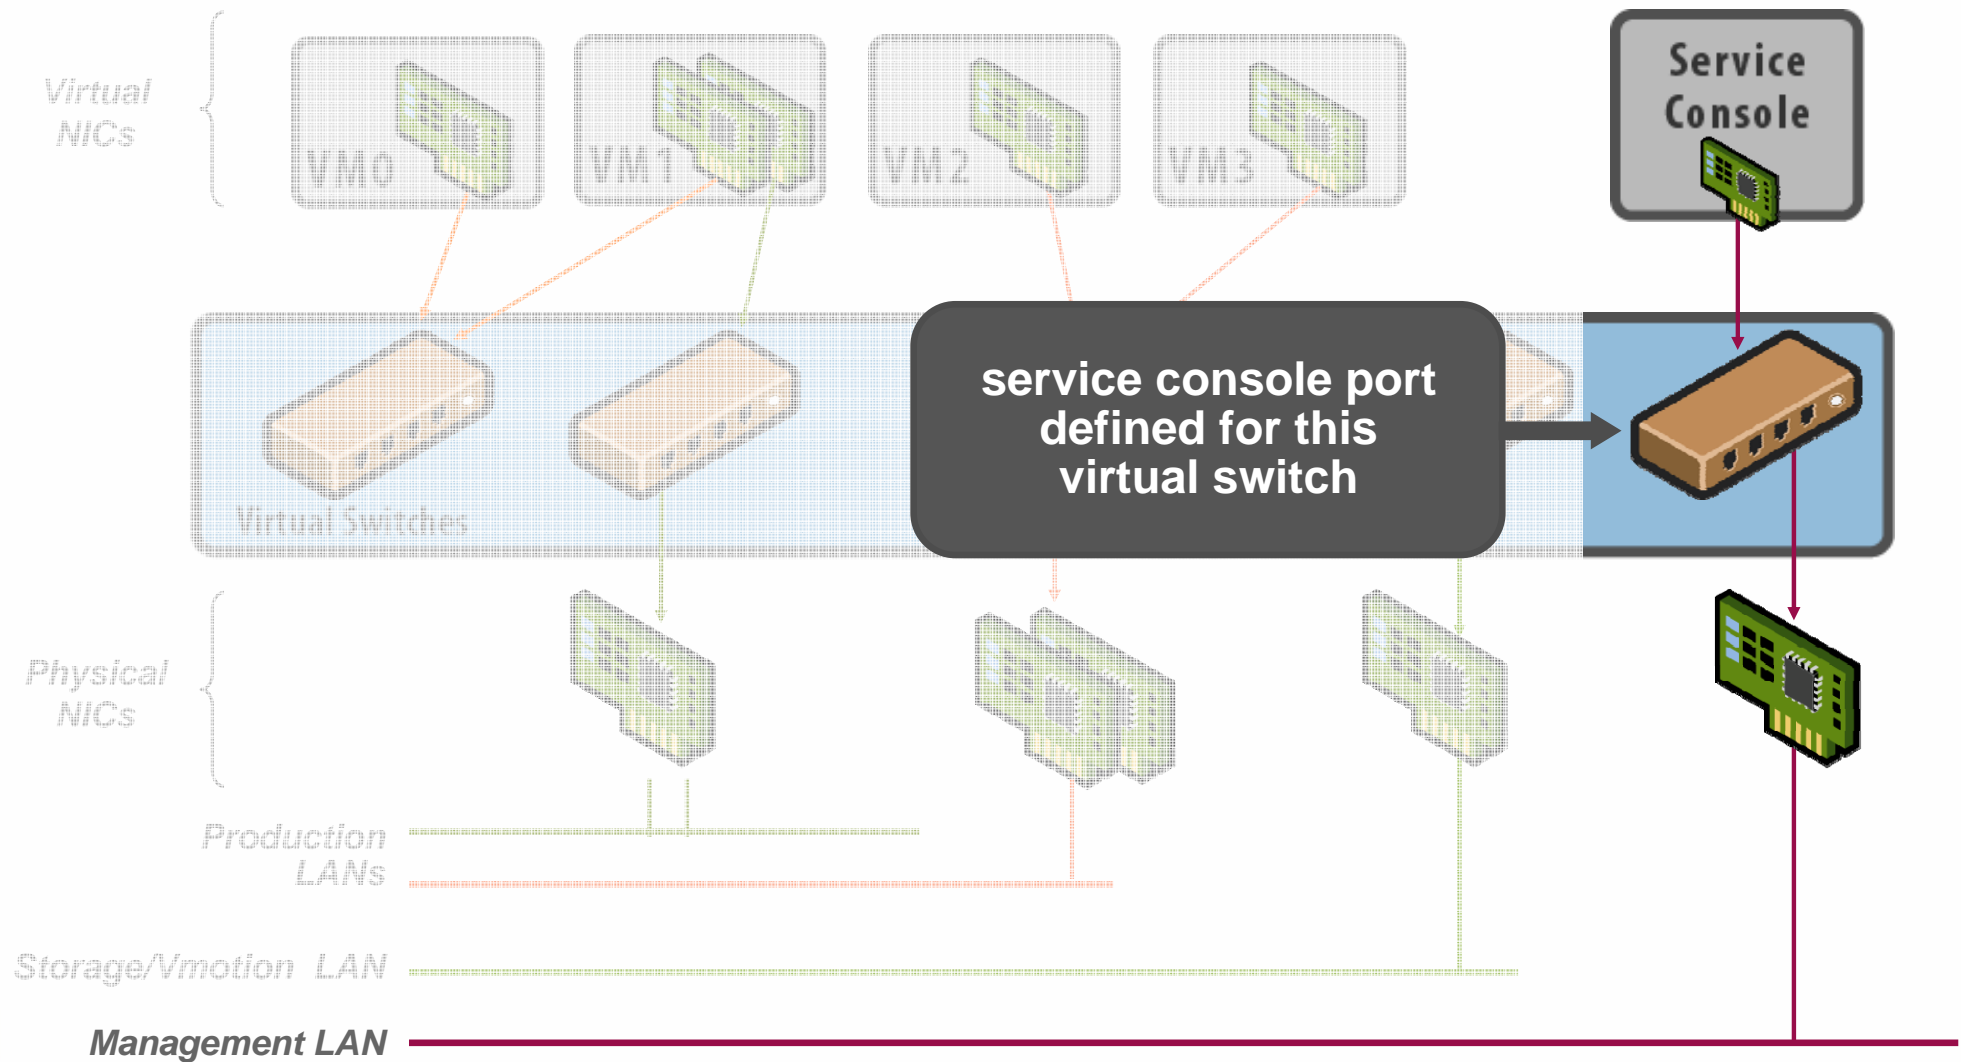
**VMWORLD 2006**

# Network Connections

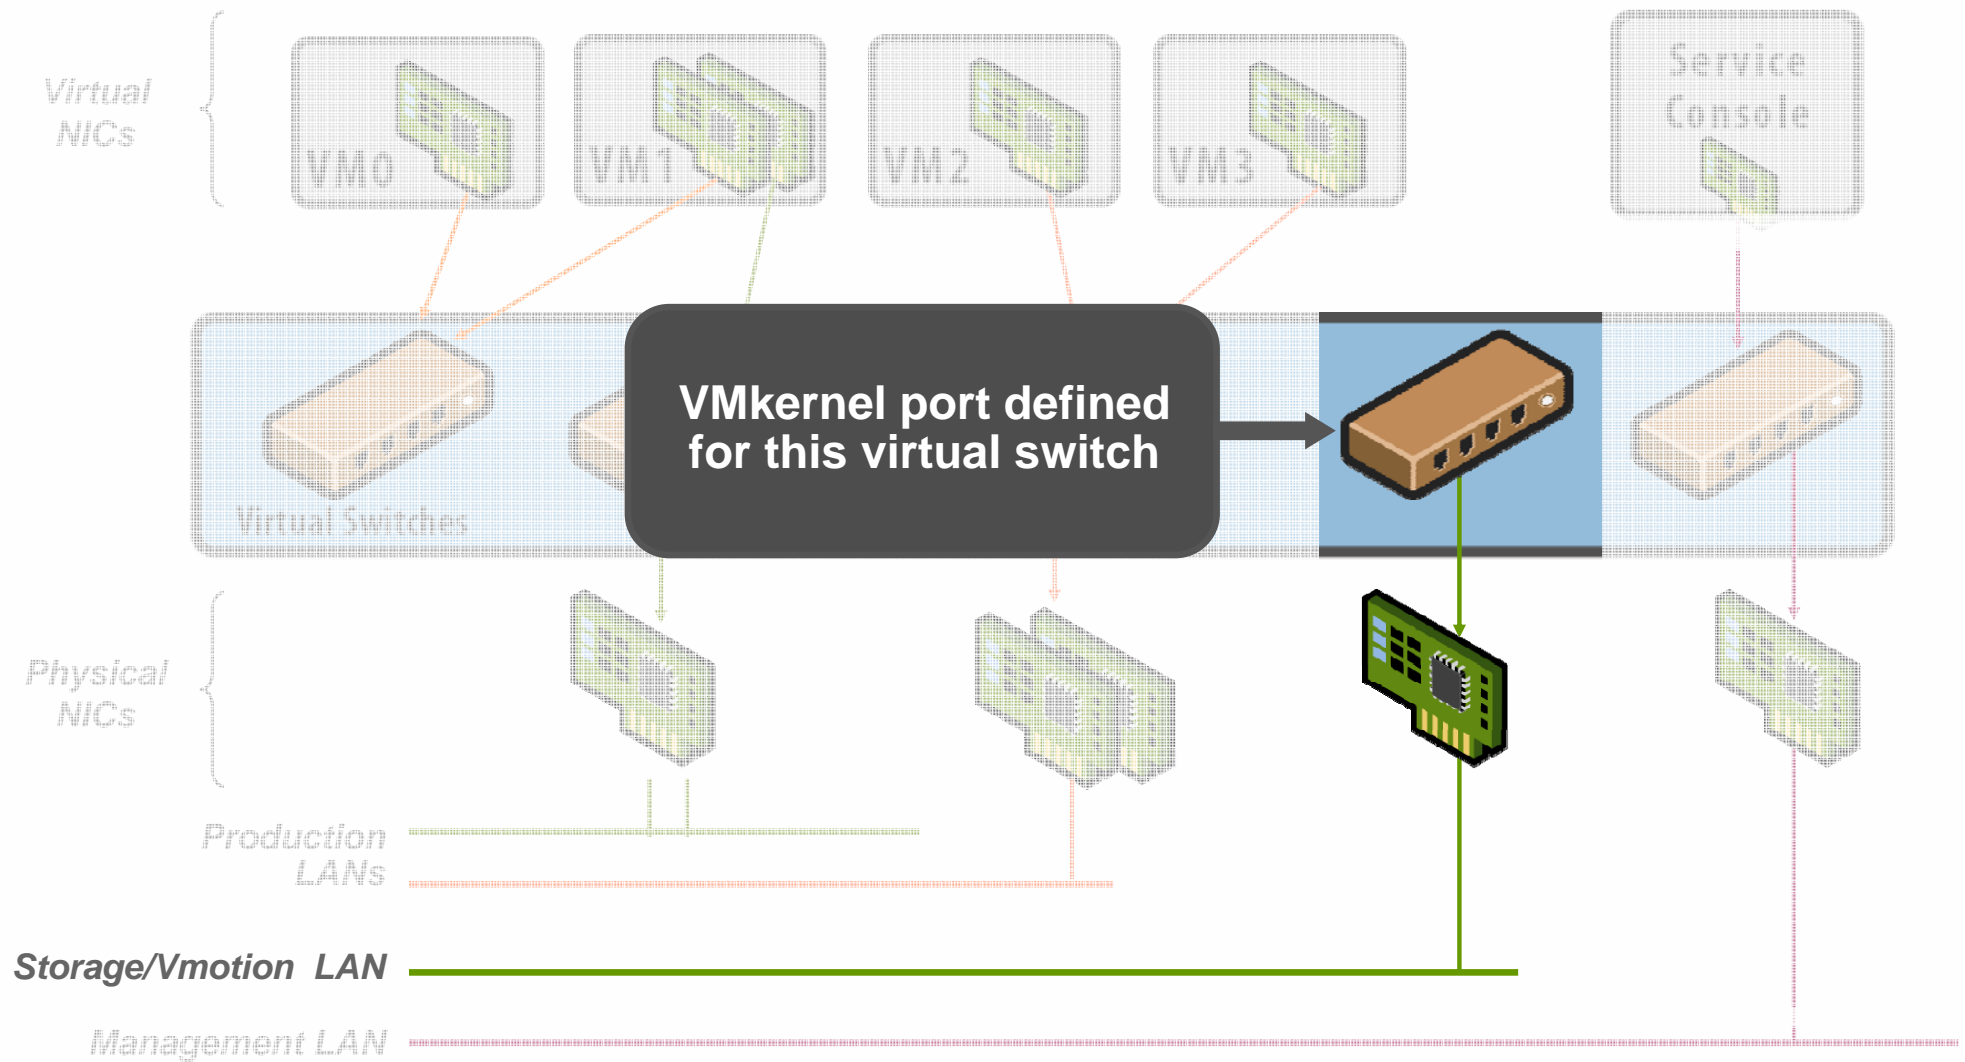
- There are three types of network connections:
  - Service console port – access to ESX Server management network
  - VMkernel port – access to VMotion, iSCSI and/or NFS/NAS networks
  - Virtual machine port group – access to VM networks
- More than one connection type can exist on a single virtual switch, or each connection type can exist on its own virtual switch




## Connection Type: Service Console Port



## Connection Type: VMkernel Port

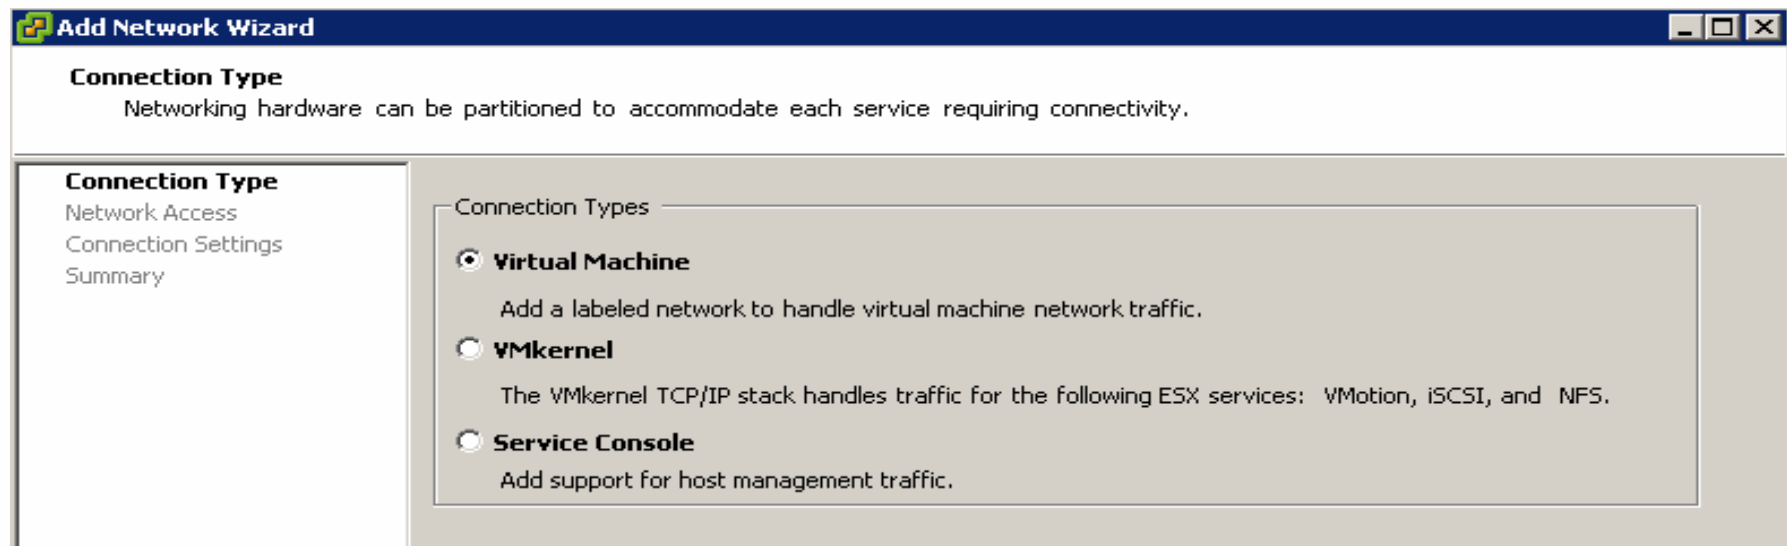


## Connection Type: Virtual Machine Port Group



# Defining Connections


- A connection type is specified when creating a new virtual switch
- Parameters for the connection are specified during setup
- More connections can be added later




# Naming Virtual Switches And Connections

- All virtual switches are known as vSwitch#
- Every port or port group has a network label
- Service console ports are known as vSwif#


Virtual Switch: vSwitch0



Virtual Switch: vSwitch1



Virtual Switch: vSwitch2



Policies



**VMWORLD 2006**



# Network Policies

- There are four network policies:
  - VLAN
  - Security
  - Traffic shaping
  - NIC teaming
- Policies are defined
  - At the virtual switch level
    - Default policies for all the ports on the virtual switch
  - At the port or port group level
    - Effective policies: Policies defined at this level override the default policies set at the virtual switch level

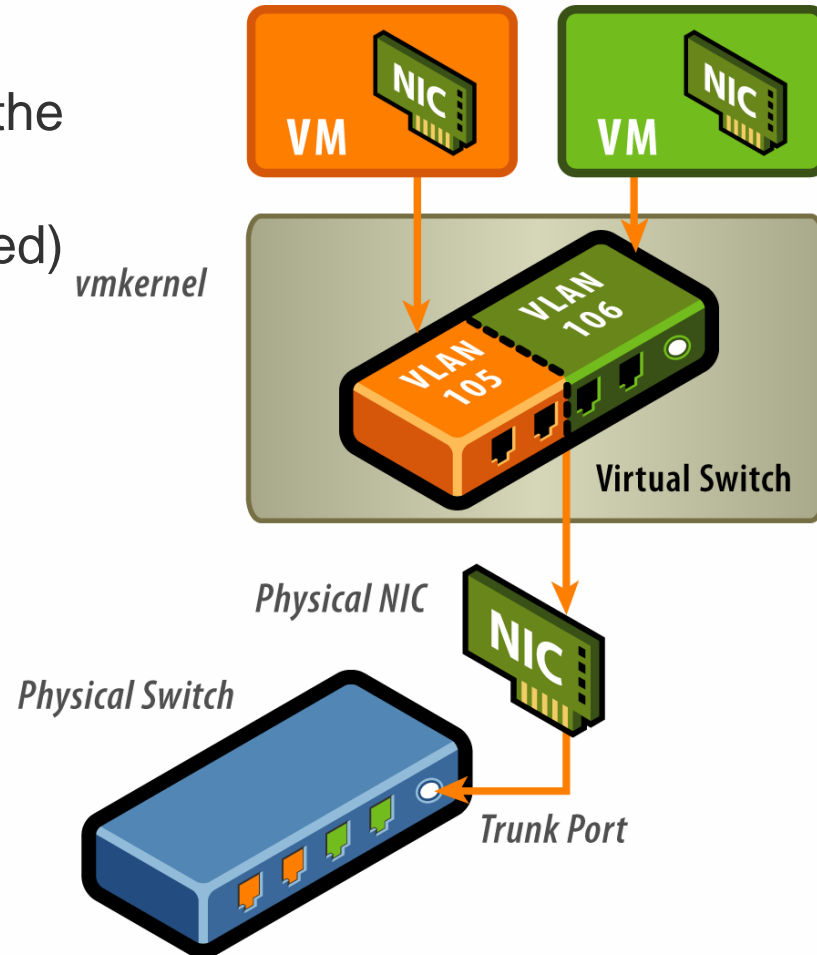
## Network Policy: VLANs

- Virtual LANs (VLANs) allow the creation of multiple logical LANs within or across physical network segments
- VLANs free network administrators from the limitations of physical network configuration
- VLANs provide several important benefits
  - Improved security: the switch only presents frames to those stations in the right VLANs
  - Improved performance: each VLAN is its own broadcast domain
  - Lower cost: less hardware required for multiple LANs
- ESX Server includes support for IEEE 802.1Q VLAN Tagging

## Network Policy: VLANs (2)

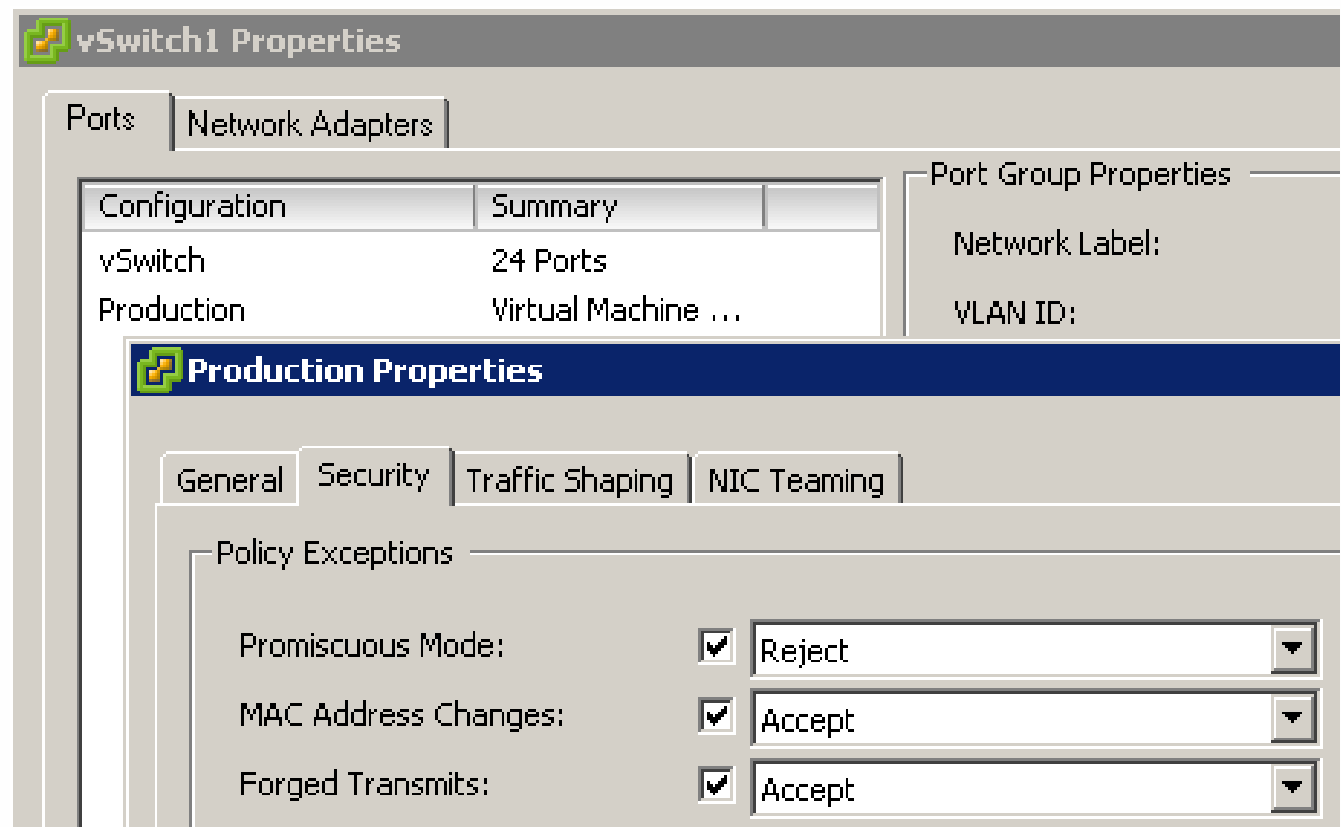
### ■ Virtual switch tagging

- Packets leaving a VM are tagged as they pass through the virtual switch
- Packets are cleared (untagged) as they return to the VM
- Little impact on performance




## Network Policy: Security

- Administrators can configure Layer 2 Ethernet security options at the virtual switch and at the port groups
- There are three security policy exceptions:
  - Promiscuous Mode
  - MAC Address Changes
  - Forged Transmits



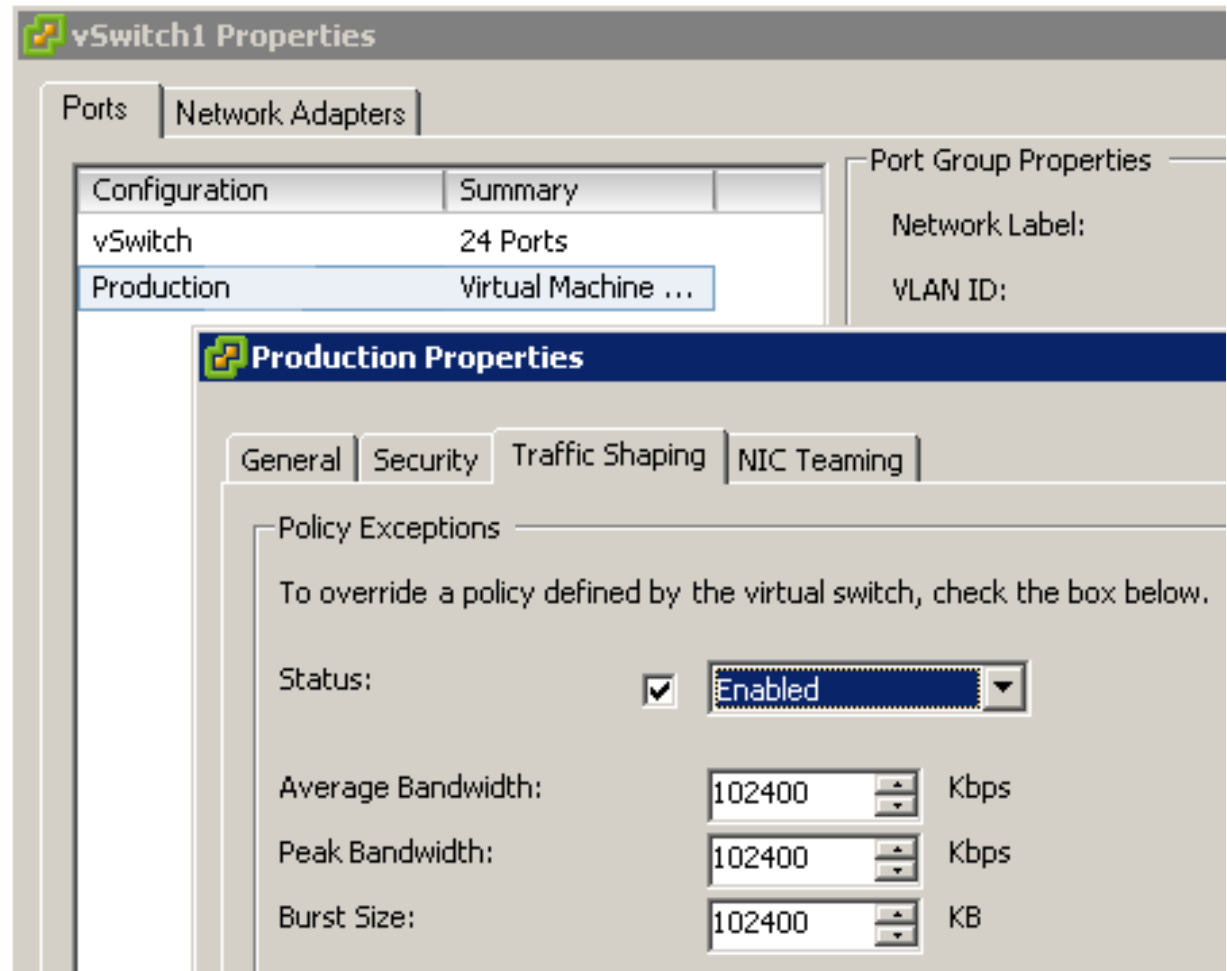
## Network Policy: Traffic Shaping

- Network traffic shaping is a mechanism for controlling a VM's outbound network bandwidth
- Average rate, peak rate, and burst size are configurable



## Network Policy: Traffic Shaping (2)

- Disabled by default
- Can be enabled for the entire virtual switch
  - Port group settings override the switch settings
- Shaping parameters apply to each virtual NIC in the virtual switch



# Network Policy: NIC Teaming

- NIC Teaming settings:
  - > Load Balancing
  - > Network Failure Detection
  - > Notify Switches
  - > Rolling Failover
  - > Failover Order
- Port group settings are similar to the virtual switch settings
  - > Except port group failover order can override vSwitch failover order

**Production Properties**

General | Security | Traffic Shaping | **NIC Teaming**

Policy Exceptions

Load Balancing: ☐ Route based on the originating virtual port ID

Network Failure Detection: ☒ Link Status only

Notify Switches: ☐ Yes

Rolling Failover: ☐ No

Failover Order:

☐ Override vSwitch failover order:

Select active and standby adapters for this port group. In a failover situation, standby adapters activate in the order specified below.

Name	Speed	Networks
<b>Active Adapters</b>		
vmnic1	100 Full	192.168.51.1-192.168.51.254
<b>Standby Adapters</b>		
<b>Unused Adapters</b>		

Move Up

Move Down

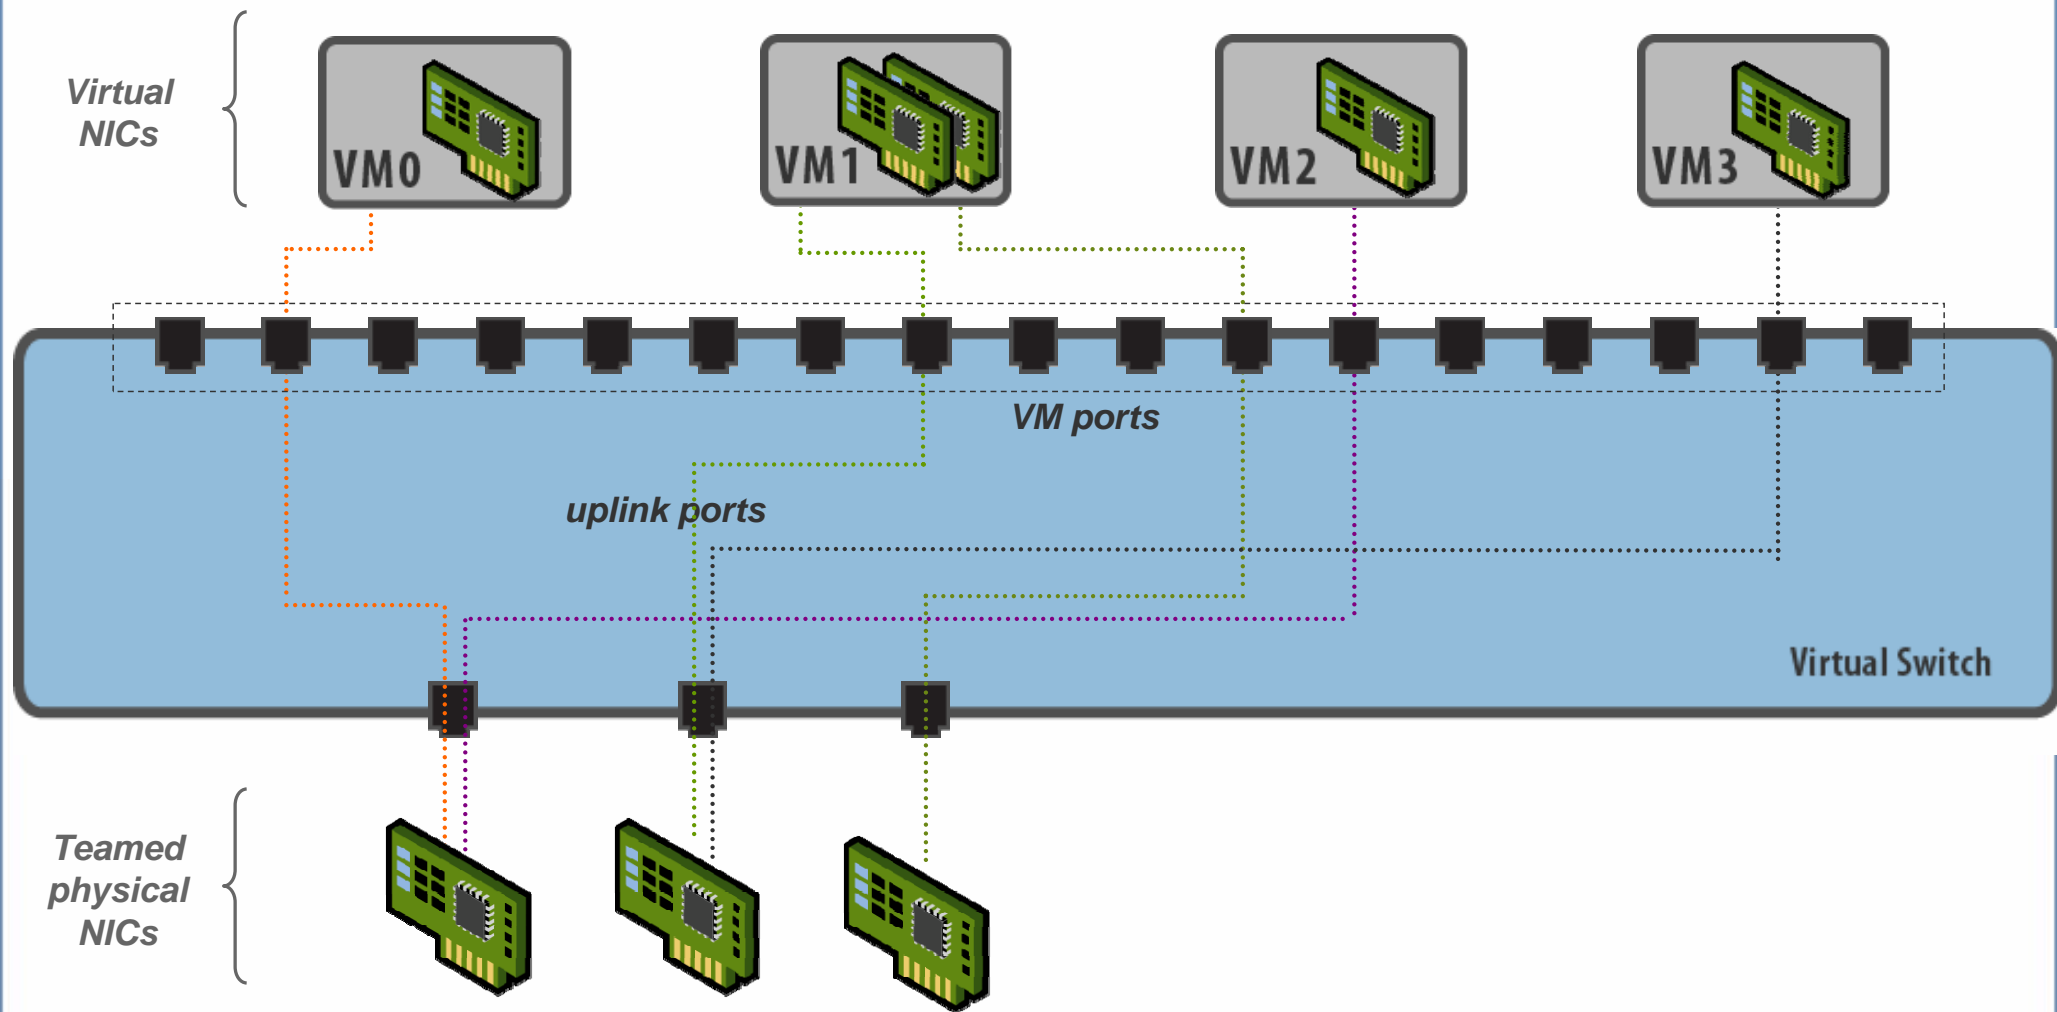
Adapter Details

No adapter selected

Driver:


Location:

## Load Balancing: vSwitch Port-based (Default)






## Load Balancing: Source MAC-based



## Load Balancing Method: IP-based

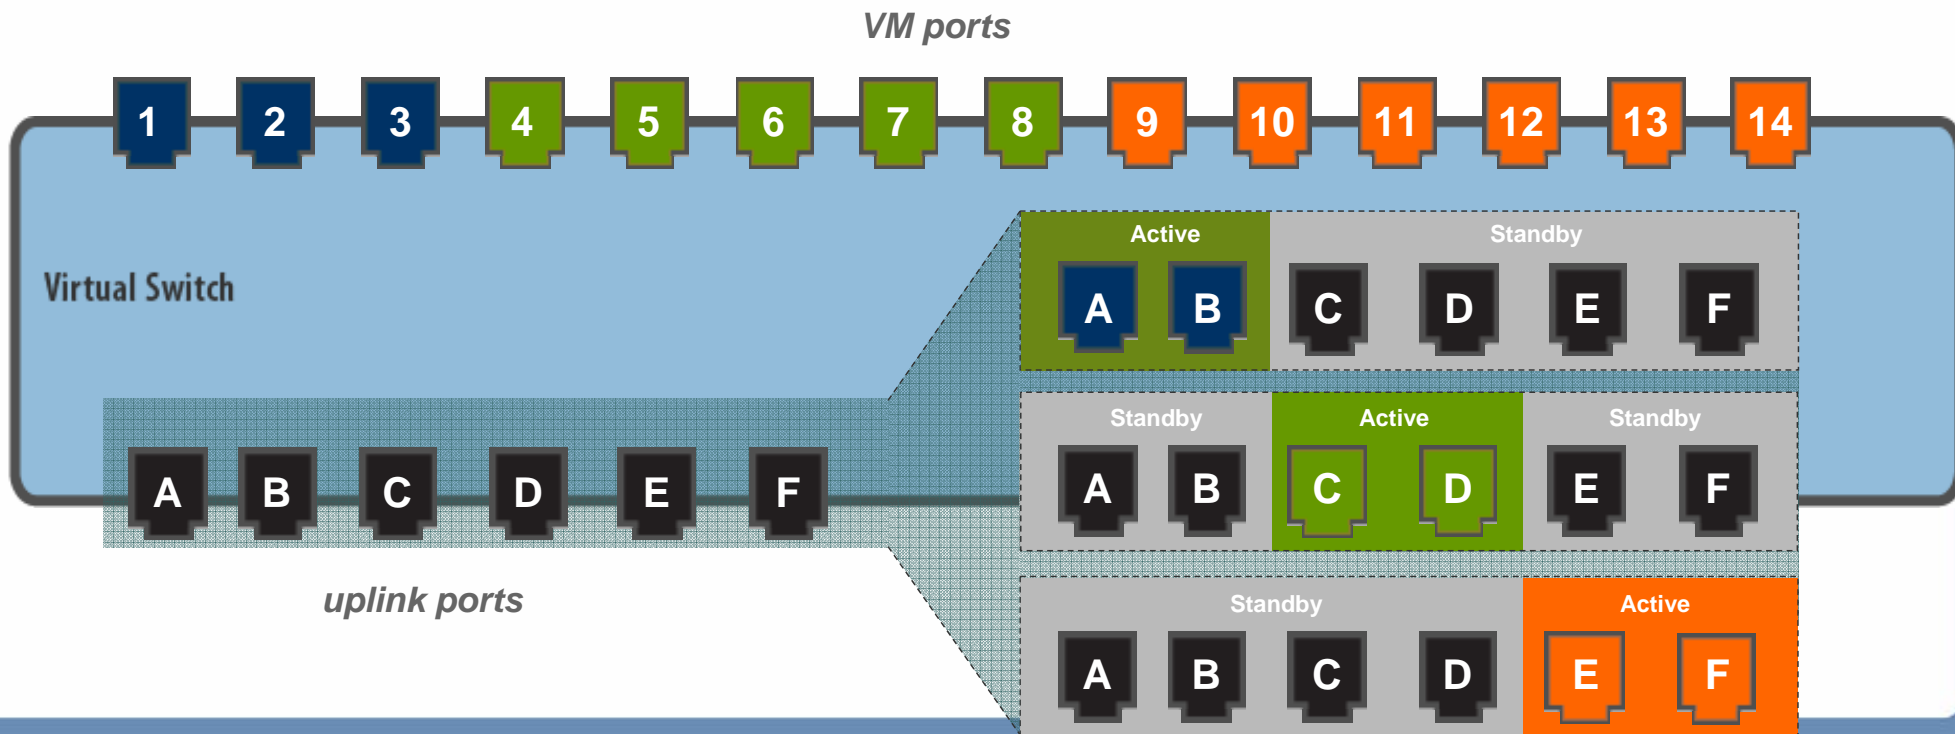


## Detecting And Handling Network Failure

- Network failure is detected by the VMkernel, which monitors the following:
  - Link state only
  - Link state + beaconing
- Switches can be notified whenever
  - There is a failover event
  - A new virtual NIC is connected to the virtual switch
  - Updates switch tables and minimizes failover latency
- Failover is implemented by the VMkernel based upon configurable parameters
  - Failover order: Explicit list of preferred links (uses highest-priority link which is up)
    - Maintains load balancing configuration
    - Good if using a lower bandwidth standby NIC
  - Rolling failover -- preferred uplink list sorted by uptime

## Multiple Policies Applied To A Single Team

- Different port groups within a vSwitch can implement different networking policies
  - This includes NIC teaming policies
- Example: different active/standby NICs for different port groups of a switch using NIC teaming




# IP Storage



**VMWORLD 2006**

# What is iSCSI?

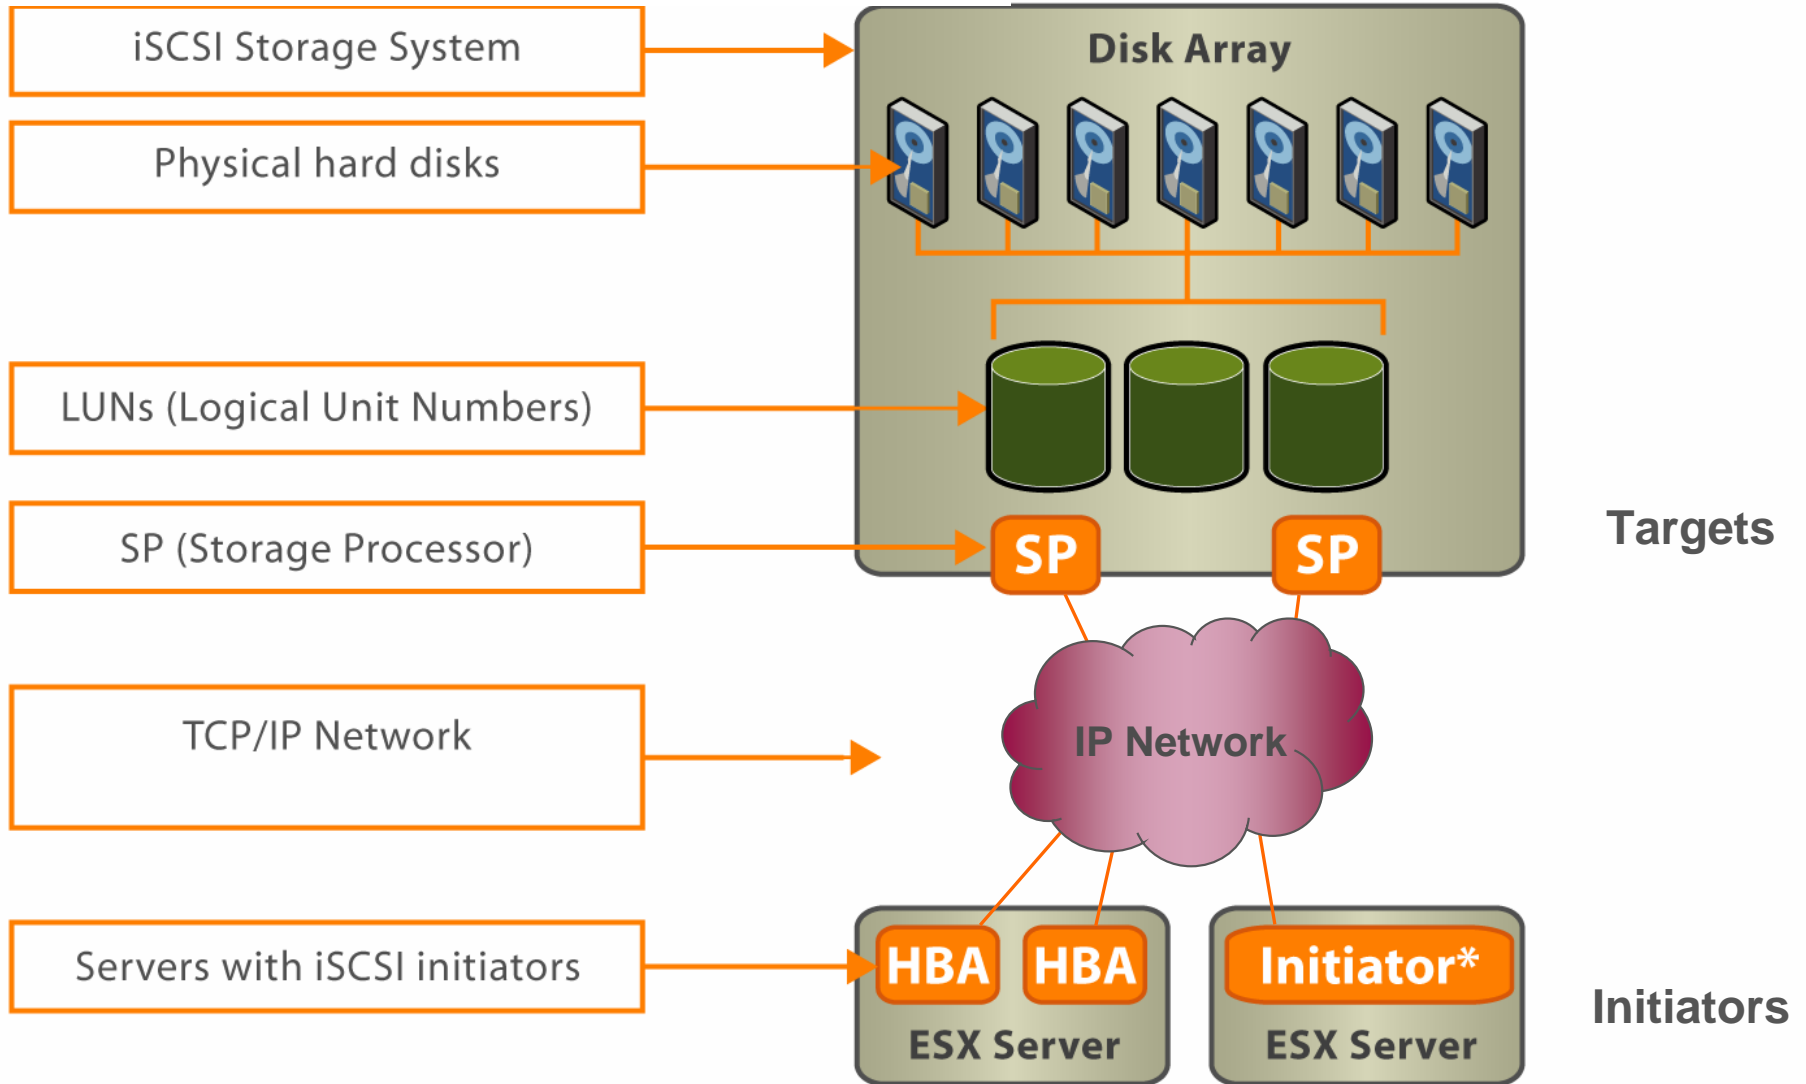
- A SCSI transport protocol, enabling access to storage devices over standard TCP/IP networks
  - Maps SCSI block-oriented storage over TCP/IP
  - Similar to mapping SCSI over Fibre Channel
- “Initiators”, such as an iSCSI HBA in an ESX Server, send SCSI commands to “targets”, located in iSCSI storage systems



## How is iSCSI Used With ESX Server?

- Boot ESX Server from iSCSI storage
  - Using hardware initiator only
- Create a VMFS on an iSCSI LUN
  - To hold VM State, ISO images, and templates
- Allows VM access to a raw iSCSI LUN
- Allows VMotion migration of a VM whose disk resides on an iSCSI LUN

# Components of an iSCSI SAN



*\* Software implementation*



# Addressing in an iSCSI SAN

*iSCSI target name*

`iqn.1992-08.com.netapp:stor1`

*iSCSI alias*

`stor1`

*IP address*

`192.168.36.101`

*iSCSI initiator name*

`iqn.1998-01.com.vmware:train1`

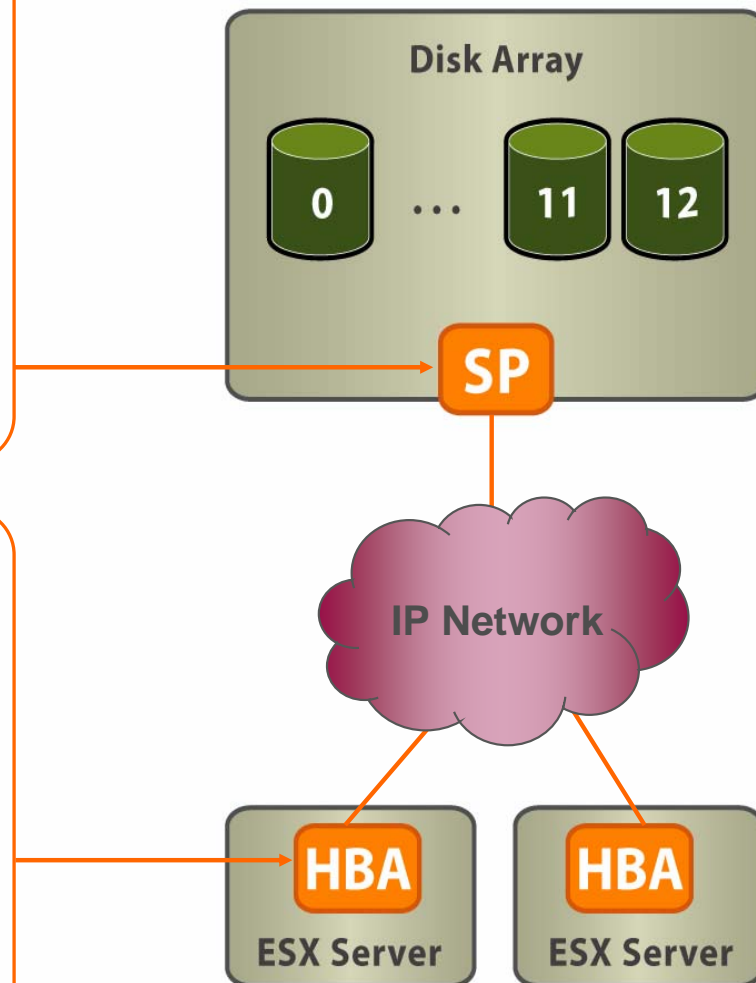
*iSCSI alias*

`train1`

*IP address*


`192.168.36.88`

iqn – iSCSI Qualified Name




## How iSCSI LUNs Are Discovered

- Two discovery methods are supported:
  - Static Configuration
  - SendTargets
- iSCSI device returns its target info as well as any additional target info that it knows about.



## Multipathing With iSCSI

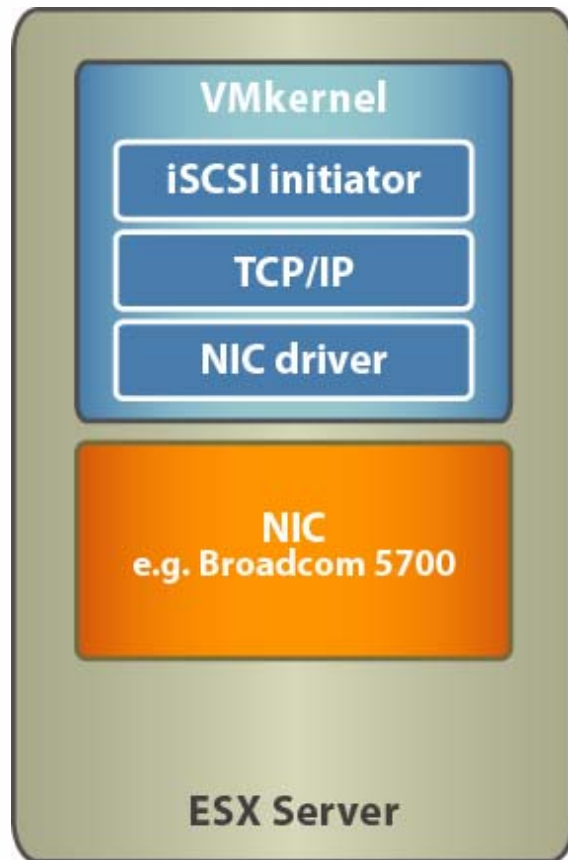
- SendTargets advertises multiple routes
  - It reports different IP addresses to allow different paths to the iSCSI LUNs
- Routing done via IP network
- For the software initiator
  - Counts as one network interface
  - NIC teaming and multiple SPs allow for multiple paths
- Currently supported via mru policy only



# iSCSI Software and Hardware Initiator

ESX Server 3 provides full support for software initiators

## *Software Initiator*




## *Hardware Initiator*



## Set Up Networking For iSCSI Software Initiator

- Both Service Console and VMkernel need to access the iSCSI storage (software initiator uses *vmkiscsid*, a daemon that runs in the service console)
- Two ways to do this:
  1. Have Service Console port and VMkernel port share a virtual switch and be in the same subnet

vSwitch2



2. Have routing in place so both the Service Console port and VMkernel port can access the storage

# Enable the Software iSCSI Client

## Security Profile

### Firewall

#### Incoming Connections

SSH Server  
CIM Server  
CIM Secure Serv  
EMC AAM Client  
CIM SLP

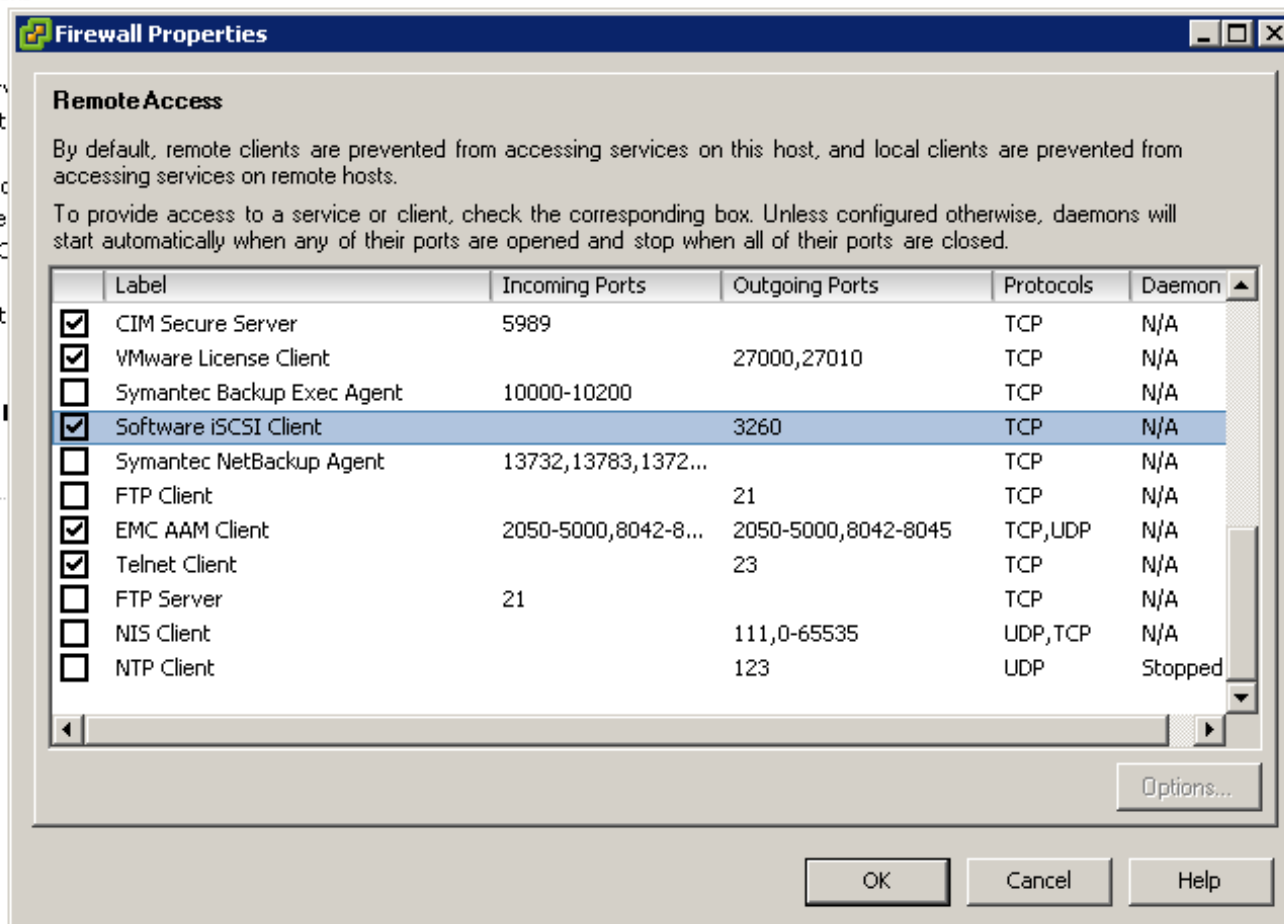
#### Outgoing Connections

VMware License  
VMware VirtualC  
Telnet Client  
EMC AAM Client  
CIM SLP

#### Virtual Machine I

Read and write to

User Name:



# Configure the iSCSI Software Adapter

## Storage Adapters

[Rescan...](#)

Device	Type	SAN Identifier	
<b>Smart Array 6i</b>			
vmhba1	Block SCSI		
<b>QLA2340/2340L</b>			
vmhba0	Fibre Channel SCSI	21:00:00:e0:8b:89:19:9c	
<b>iSCSI Software Adapter</b>			
iSCSI Software Adapter	iSCSI		

## Details

[Properties...](#)

Model:

iSCSI Name:

iSCSI Alias:

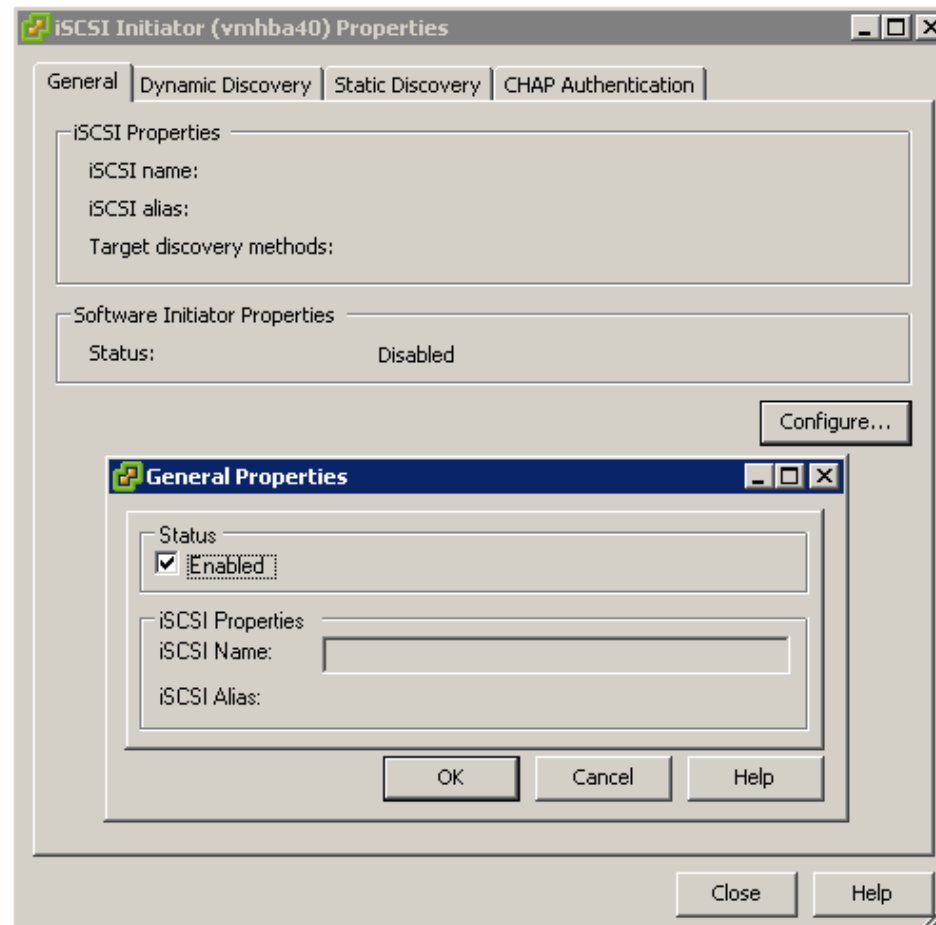
IP Address:

Discovery Methods:

Targets:

# Configure Software Initiator: General Properties

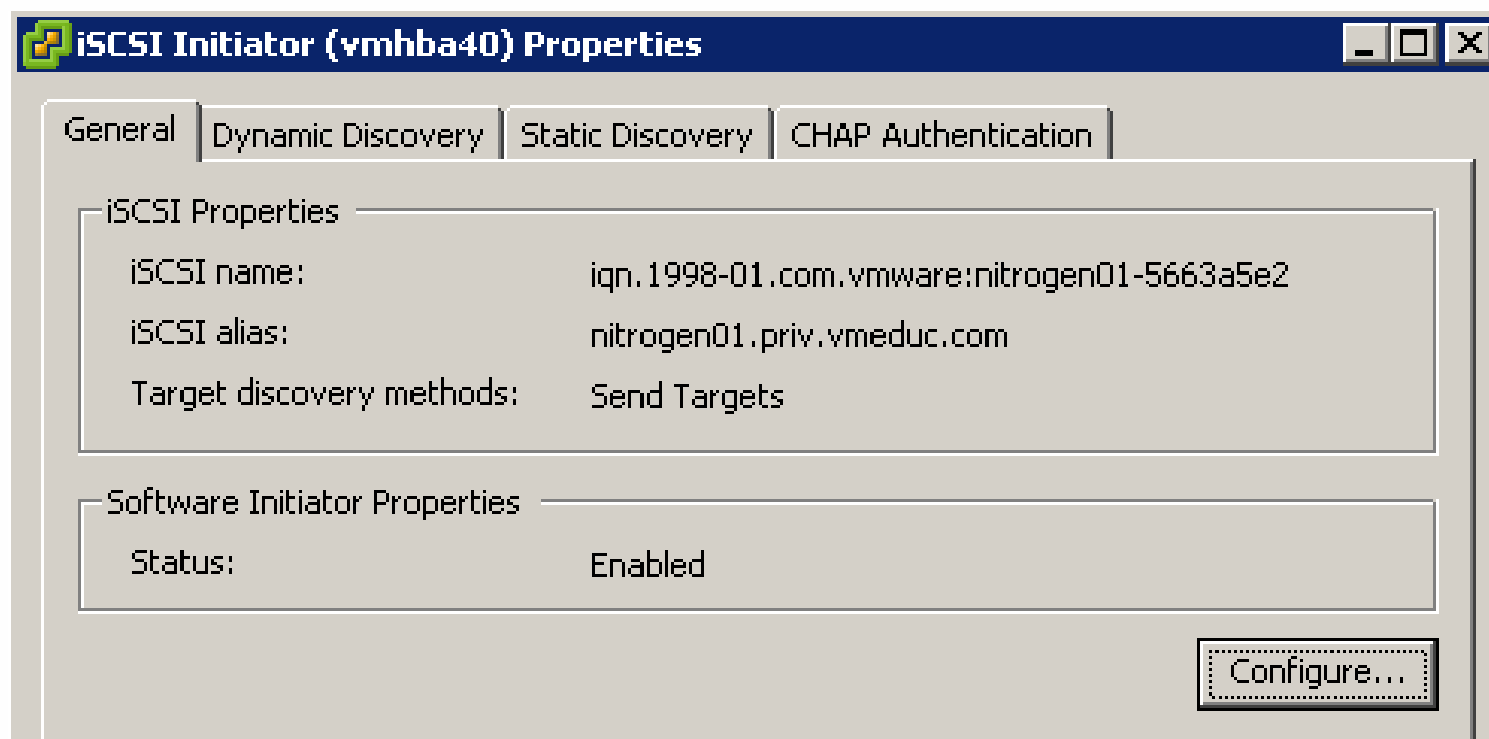
- Enable the iSCSI initiator





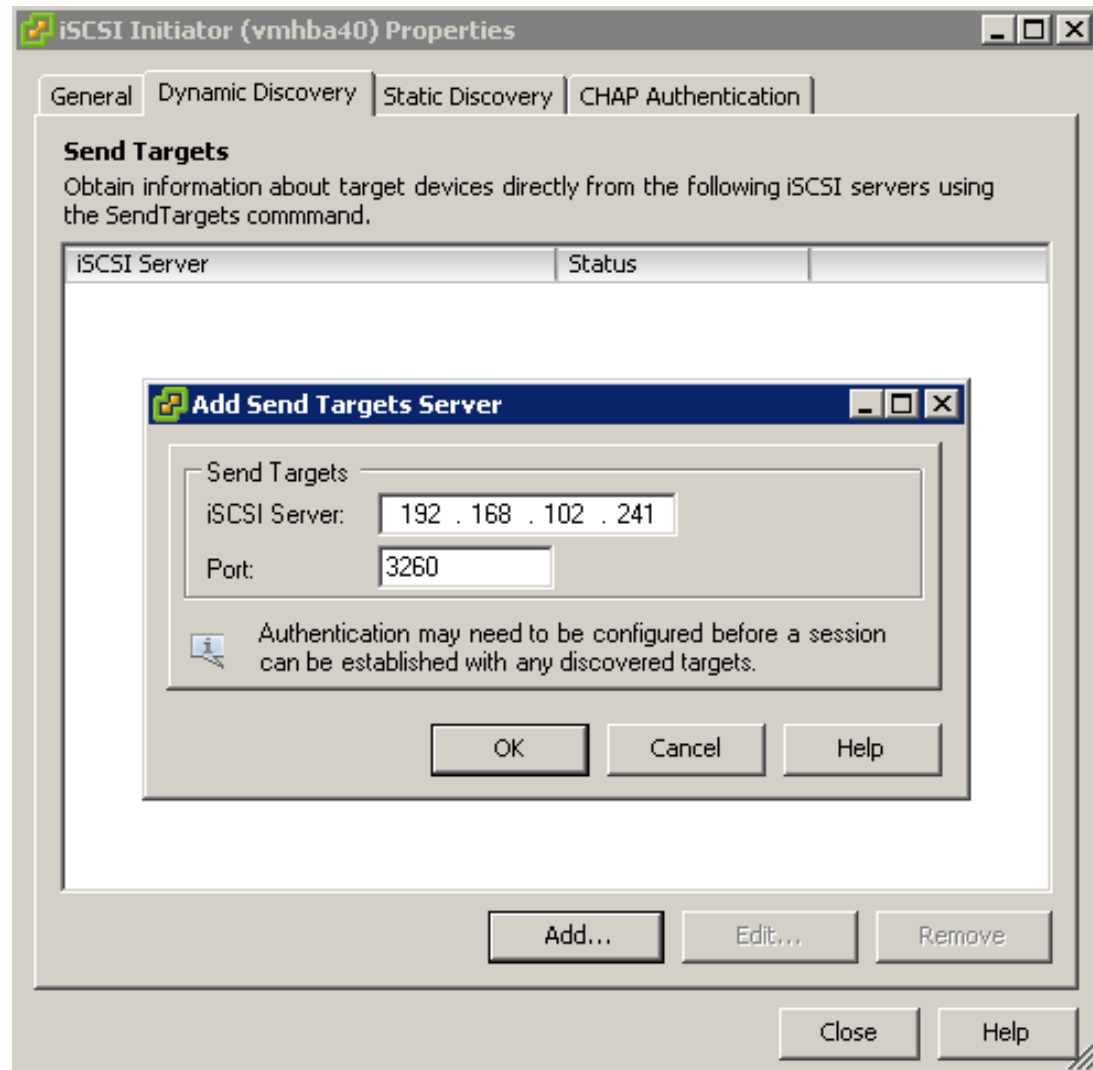
## Configure Software Initiator: General Properties (2)

- The iSCSI name and alias are automatically filled in after initiator is enabled



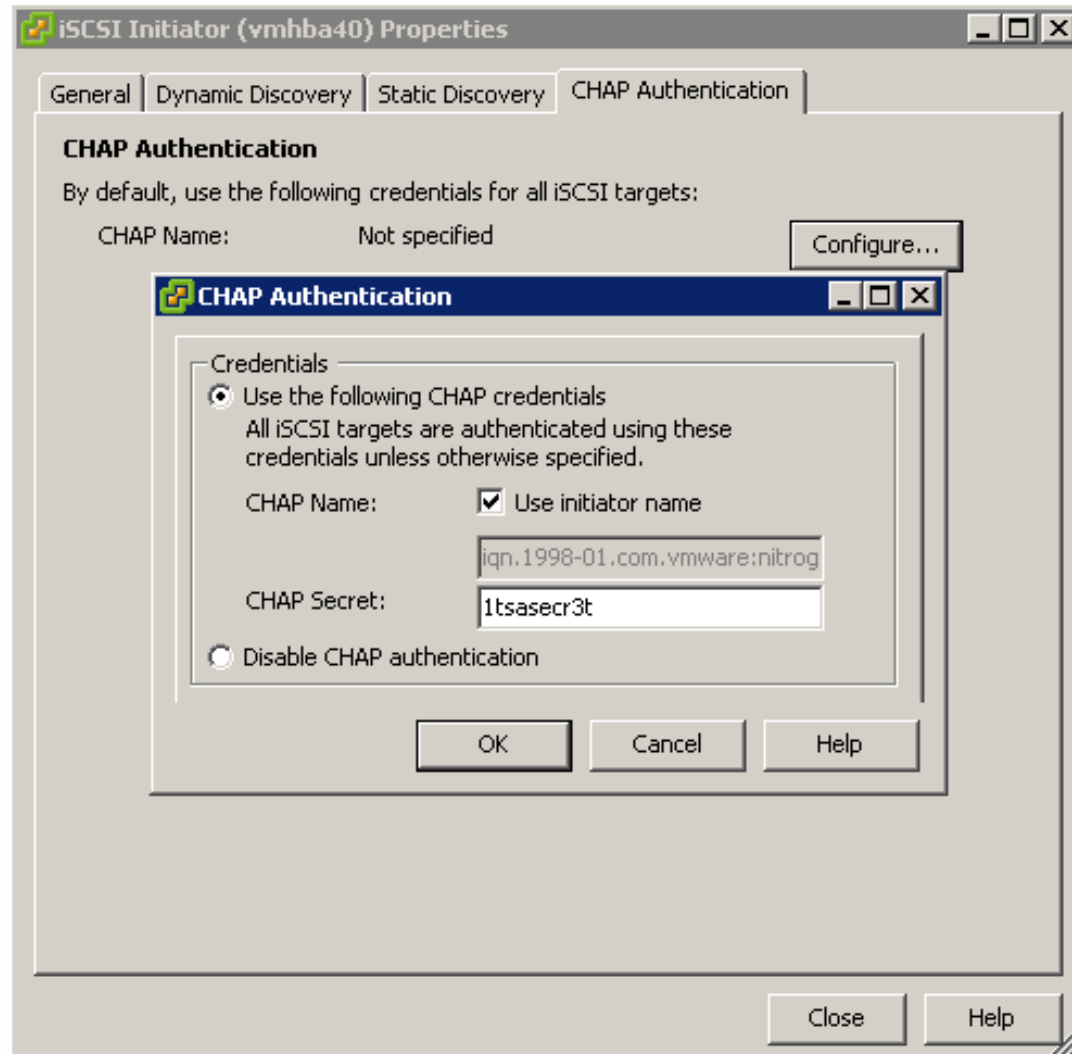
# Configure Software Initiator: Dynamic Discovery

- In the Dynamic Discovery tab, enter the IP address of each target server for initiator to establish a discovery session
- All available targets returned by the target server show up in the Static Discovery tab



# Configure Software Initiator: CHAP Authentication

- By default, CHAP is disabled
- Enable CHAP and enter CHAP name and secret



# Discover iSCSI LUNs

## ■ Rescan to find new LUNs

dhiltgen-dev1 VMware Host Agent, e.x.p, 19319

Summary

Virtual Machines

Resource Allocation

Performance

Configuration

Users & Groups

System Logs

Events

Permissions

### Hardware

[Storage \(SCSI, SAN, and NFS\)](#)

[Networking](#)

[Processors](#)

[Memory](#)

► [Storage Adapters](#)

[Network Adapters](#)

### Software

[Licensed Features](#)

[DNS and Routing](#)

[Virtual Machine Startup/Shutdown](#)

[SNMP Agents](#)

[Security Profile](#)

[Service Console Resources](#)

[Advanced Settings](#)

### Storage Adapters

[Rescan](#)

Device	Type	Target ID
<b>iSCSI Software Adapter</b>		
vmhba40	iSCSI	iqn.1998...
<b>53c1030 PCI-X Fusion-MPT Dual Ultra320 SCSI</b>		
vmhba0	Parallel SCSI	

### Details

#### vmhba40

[Properties...](#)

Model: iSCSI Software Adapter

iSCSI Name: iqn.1998-01.com.vmware.dhiltgen-dev1

iSCSI Alias: dhiltgen-dev1

IP Address:

Discovery Methods: Send Targets

Targets: 1

#### SCSI Target 0

iSCSI Name: iqn.1992-08.com.netapp:burton

iSCSI Alias:

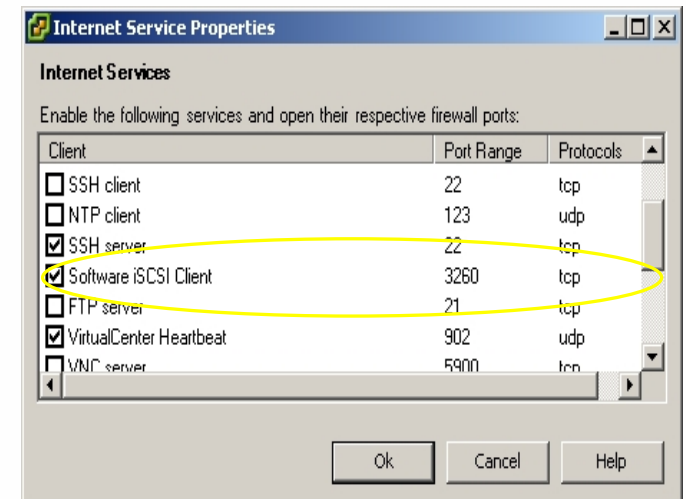
Target LUNs: 5

[Hide LUNs](#)

Path	Canonical Path	Capacity	LUN ID
vmhba40:0:0	vmhba40:0:0	TBD	vmhba40:0:0
vmhba40:0:1	vmhba40:0:1	TBD	vmhba40:0:1
vmhba40:0:2	vmhba40:0:2	TBD	vmhba40:0:2

# iSCSI Tips and Tricks

- Do not use software iSCSI initiators in virtual machines
- Set console OS firewall to allow iSCSI port traffic if using software initiator
- Default iqn names incompatible with some targets – use this format
  - iqn.yyyy-mm.<domain>.<hostname>:<user defined string>
  - For example: iqn.2006-03.esxtest.vmware.com:esx3a-0a97886a.
- Can use QLogic SANsurfer for QLA4010 setup
  - Install on COS with:
    - `sh ./iSCSI_SANsurfer_4_01_00_linux_x86.bin -i silent -D SILENT_INSTALL_SET="QMSJ_LA"`
  - Start iqlremote in COS, connect from remote UI application



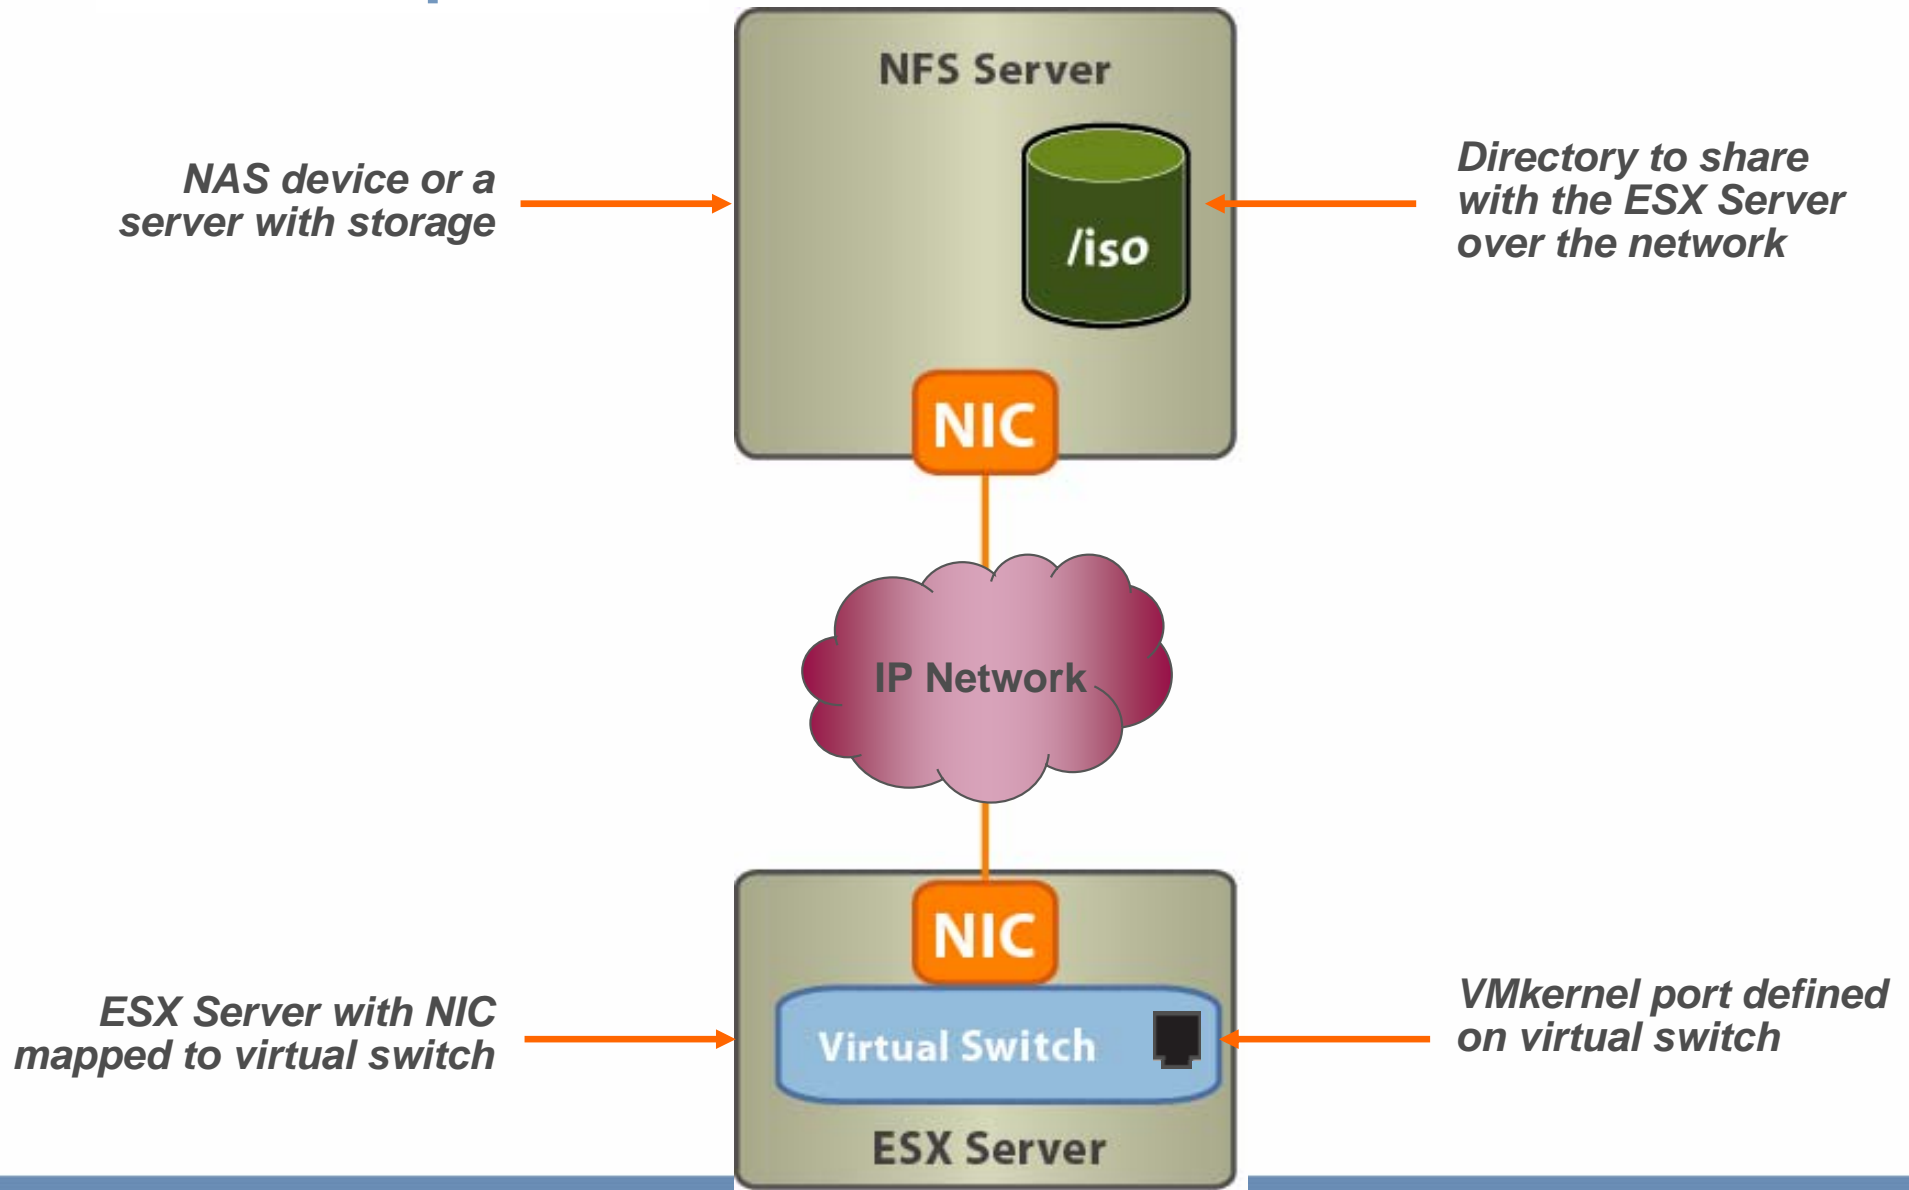
## What is NAS and NFS?

- What is NAS?
  - Network-Attached Storage
  - Storage shared over the network at a filesystem level
- Why use NAS?
  - A low-cost, moderate-performance option
  - Less infrastructure investment required than with Fibre Channel
- There are two key NAS protocols:
  - NFS (the “Network File System”)
  - SMB (Windows networking, also known as “CIFS”)
- Major NAS appliances support both NFS and SMB
  - Notably those from Network Appliance and EMC
- Server operating systems also support both

## How is NAS Used With ESX Server?

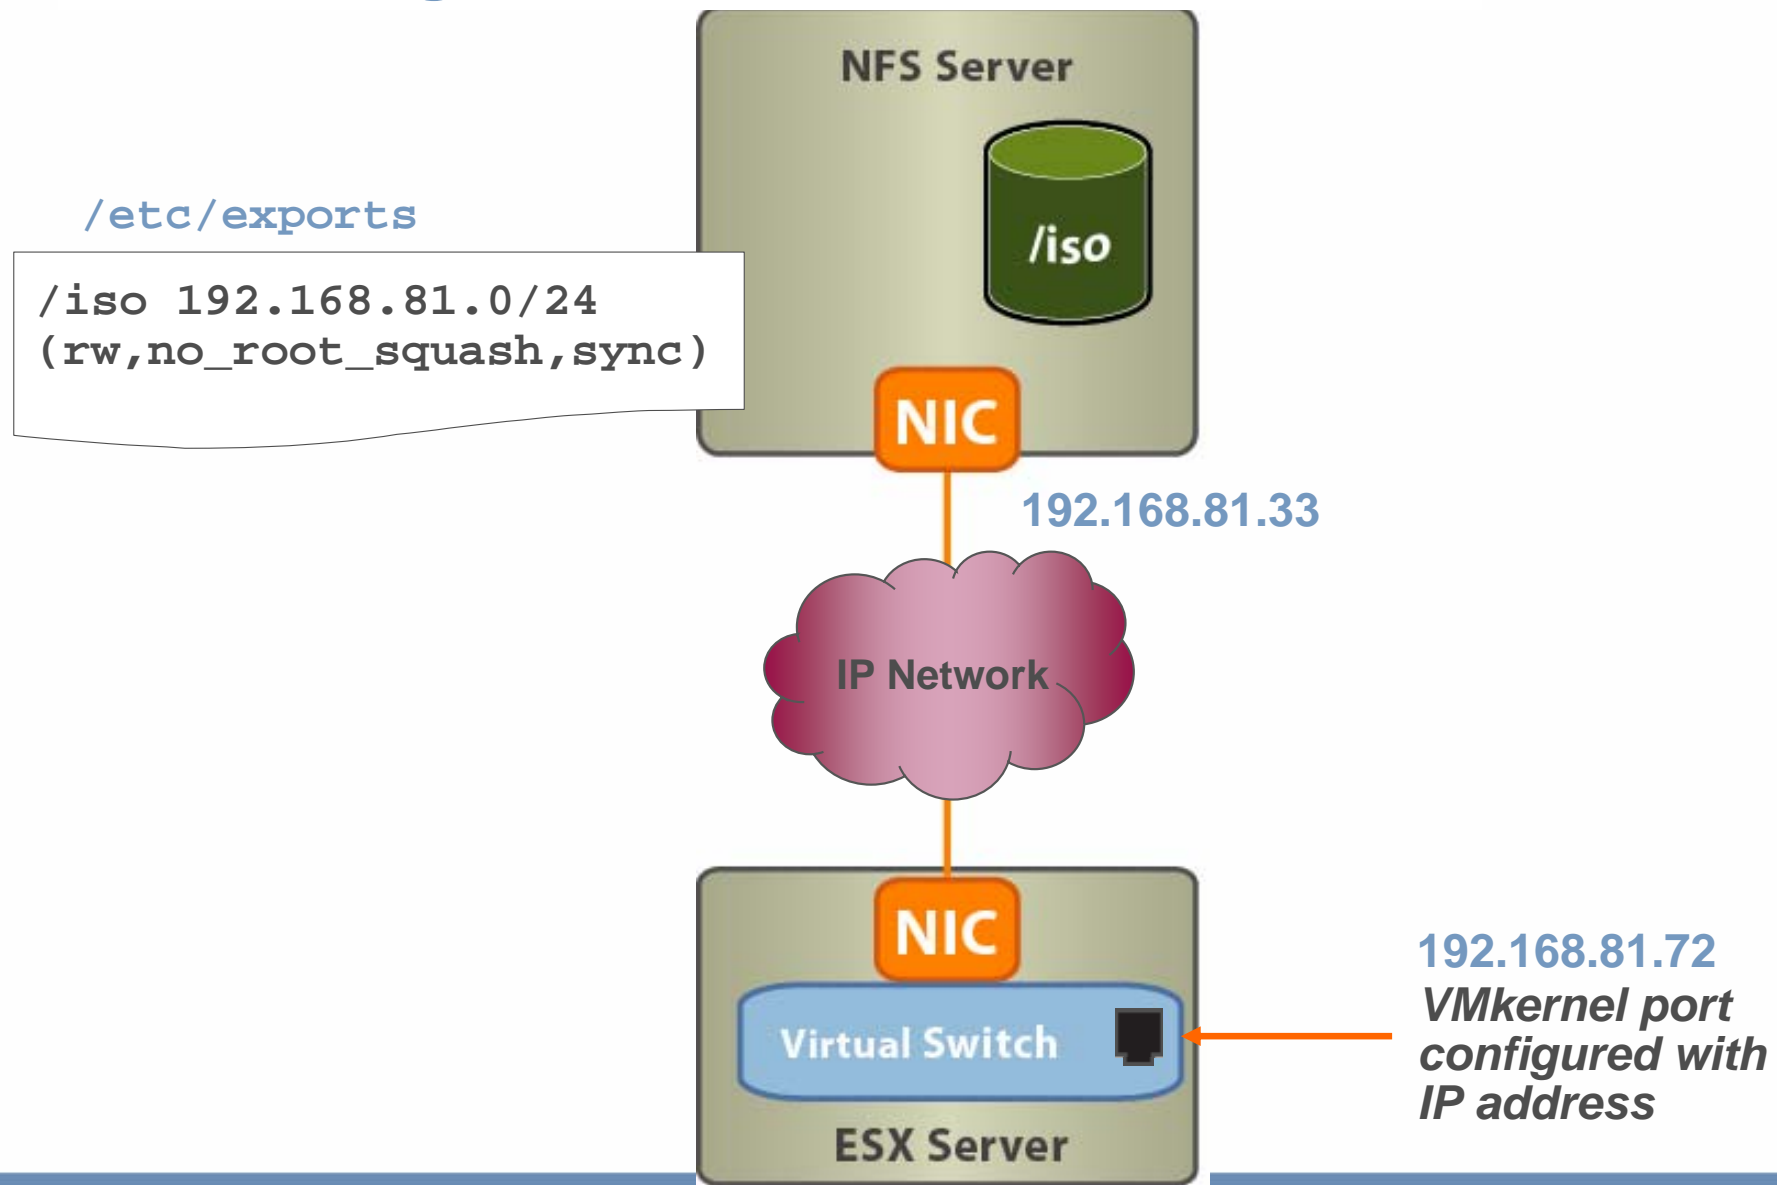
- The VMkernel only supports NFS
  - More specifically NFS version 3, carried over TCP
- NFS volumes are treated just like VMFS volumes in Fibre Channel or iSCSI storage
  - Any can hold VMs' running virtual disks
  - Any can hold ISO images
  - Any can hold VM templates
- Virtual machines with virtual disks on NAS storage can be VMotioned, subject to the usual constraints
  - Compatible CPUs
  - All needed networks and storage must be visible at destination

## NFS Components





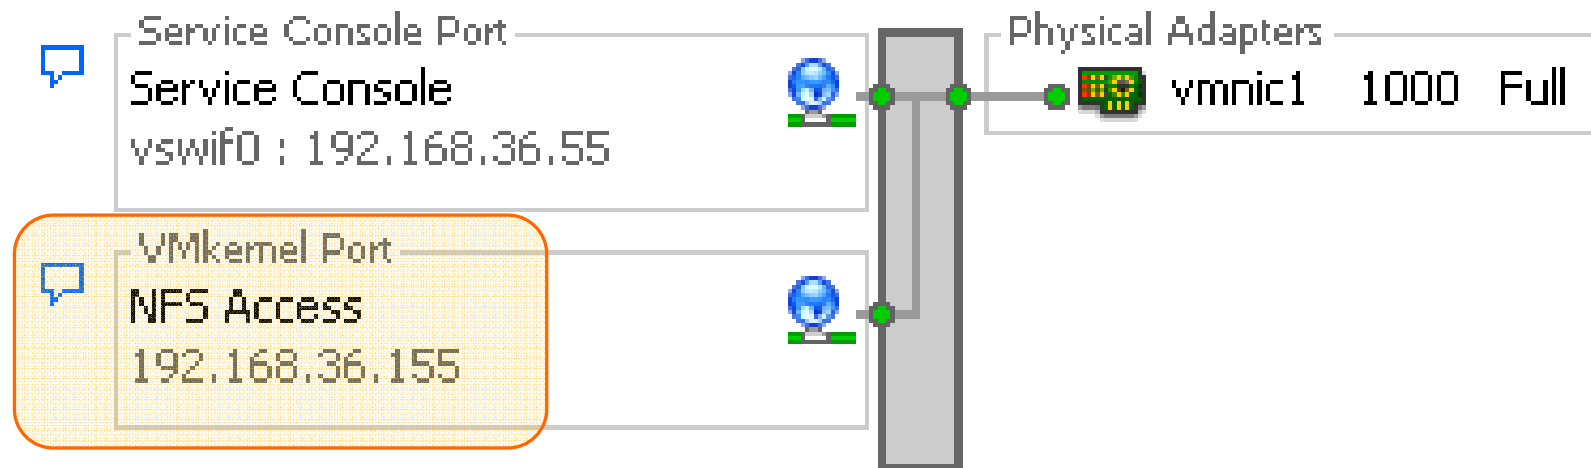
# Addressing and Access Control With NFS



## Before You Begin Using NAS/NFS

- Create a VMkernel port on a virtual switch

Virtual Switch: vSwitch0



***You must define a new IP address for NAS use, different from the Service Console's IP address***

# Configure an NFS Datastore

- Describe the NFS share

**Add Storage**

**Locate Network File System**  
Which shared folder will be used as a VMware datastore?

☒ **NAS**  
**Network File System**  
Ready to Complete

**Properties**

Server:   
Examples: nas, nas.it.com or 192.168.0.1

Folder:   
Example: /vols/vol0/datastore-001




☐ Mount NFS read only

**Datastore Name**

## Configure an NFS Datastore (cont.)

- Verify that the NFS datastore has been added

### Storage

Identification	Device	Capacity	Free	Type
 NFS01	192.168.56.131:/...	7.10 GB	5.02 GB	nfs
 storage1	vmhba0:0:0:3	60.25 GB	59.64 GB	vmfs3
 SharedVMs	vmhba1:0:25:1	99.75 GB	99.14 GB	vmfs3

### Details

#### NFS01


Server: 192.168.56.131

Folder: /iso

7.10 GB Capacity

2.08 GB  Used

5.02 GB  Free



Questions



**VMWORLD 2006**

Some or all of the features in this document may be representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.

# VMWORLD 2006

