VMware vSphere® 5.0 Evaluation Guide

Volume Three – Advanced Networking Features

TECHNICAL WHITE PAPER

Table of Contents

About This Guide

The purpose of the *VMware vSphere 5.0 Evaluation Guide, Volume Three – Advanced Networking Features,* is to support a self-guided, hands-on evaluation of VMware vSphere® 5.0 ("vSphere") advanced networking features, such as NetFlow, port mirroring, and user-defined resource pools.

This guide covers evaluation cases that are suitable for IT professionals who have an existing VMware virtualization environment and who want to evaluate features in vSphere that enable greater consolidation while maintaining service levels.

System Requirements

To ensure the best experience when using this guide, the user will need to configure hardware and software as detailed in the Hardware Requirements section.

Hardware Requirements

This guide makes the following assumptions about users' existing physical infrastructure:

Servers

Users must have at least three dedicated servers capable of running VMware ESXi™ 5.0 to provide resources for this evaluation.[1]

Storage

Users must have shared storage with enough space available to allow for three 100GB dedicated datastores. Shared storage can be SAN or network-attached storage (NAS). This document assumes users have SAN-based storage.

Networking

Users must also have at least three virtual networks configured to separate virtual machine, VMware vSphere® vMotion® (vMotion), and vSphere management networks. These networks can be set up on a single virtual switch with multiple port groups, or across multiple virtual switches. For the purpose of this evaluation guide, the configuration starts with a single vSphere standard switch and migrates port groups to a VMware vSphere Distributed Switch (VDS) to support evaluation of advanced networking features.

For more detailed requirements, see the following table.

[1] These servers must be on the *VMware vSphere 5.0 Hardware Compatibility List.*

HARDWARE	MINIMUM	WHAT'S USED IN THIS GUIDE
ESXi	3 ESXi/ESX servers CPU – 2 processors of 2GHz Memory – 6GB Network – 2 x 1GB network adaptor	3 ESXi servers (Cisco UCS 1.3.1) CPU – 2 quad-core Intel eon Nehalem processors at 2.6GHz Memory – 48GB Network – 4 x 10GB network adaptor
Storage	1 datastore (100GB)	3 datastores (Fibre Channel—100GB each)
Network	1 VLAN for carrying virtual machine traffic, 1 VLAN for carrying management traffic	Separate VLANs for ESXi management, vMotion, and virtual machine traffic

Software and Licensing Requirements

This guide makes the following assumptions about users' existing software infrastructure:

vSphere

This volume of the *VMware vSphere 5.0 Evaluation Guide* requires vSphere 5.0 and licensing for Enterprise Plus. The vSphere 5.0 evaluation license available from the VMware evaluation portal provides Enterprise Plus functionality for 60 days and is the best choice for performing the vSphere 5.0 evaluations.

Guest Operating Systems

This volume of the *VMware vSphere 5.0 Evaluation Guide* will require five or six virtual machines running Windows 2003 or Windows 2008.

Evaluation Guide Environment Setup

The VMware technical marketing lab was built using a combination of Cisco UCS server hardware and EMC Clariion CX4 Fibre Channel (FC) storage. The environment consisted of eight identical four-node pods with most pods configured as a three-node ESXi cluster and a fourth node for management. In many cases, additional resources have been configured in the technical marketing test-bed configuration to support other evaluation projects and are present in the diagrams. The user can configure only what is called for in Figure 1 and can safely ignore additional resources in screen shots and topology diagrams. The following picture shows the technical marketing test rack.

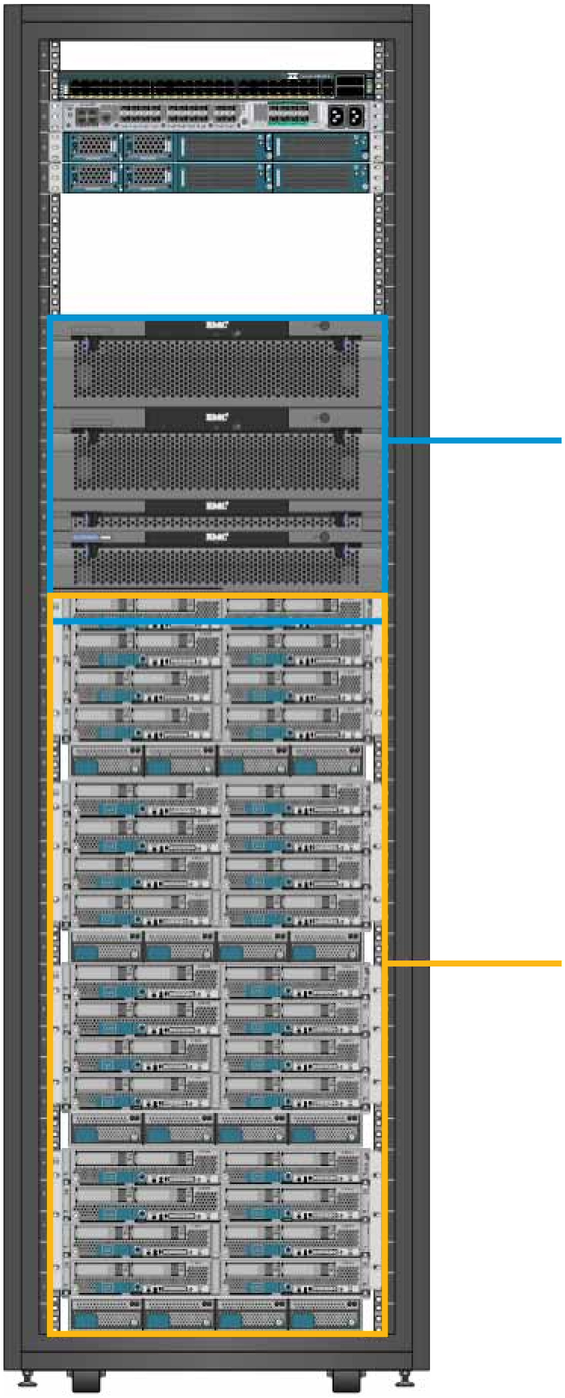

EMC CLARiiON CX$-120
Flare version 4.30.000.5.509
15x 600GB 15K SCSI
5x 200GB SSD
10x open slots

Cisco UCS 1.1.1(p)
32x dual quad-core Intel Xeon Nehalem Processors
48GB Memory
Dual-port Palo cards
1x 146GB local HDD
Configure with:
4x 10GB network adapters
2x 4/8GB HBAs
Slot 1 of each chassis reserved for core management

Server Configuration

The *VMware vSphere 5.0 Evaluation Guide* expects three modern server-class systems with adequate processors and memory to host 6–8 minimally configured virtual machines used for testing. The servers used for this evaluation do not need to be overly powerful, just reliable, and they must be on the *VMware vSphere 5.0 Hardware Compatibility List* (HCL).

Each server must have at least 2 x 1GB or 2 x 10GB network adaptors and a proper connection to shared storage. The following diagram summarizes the Evaluation Guide test-bed configuration.

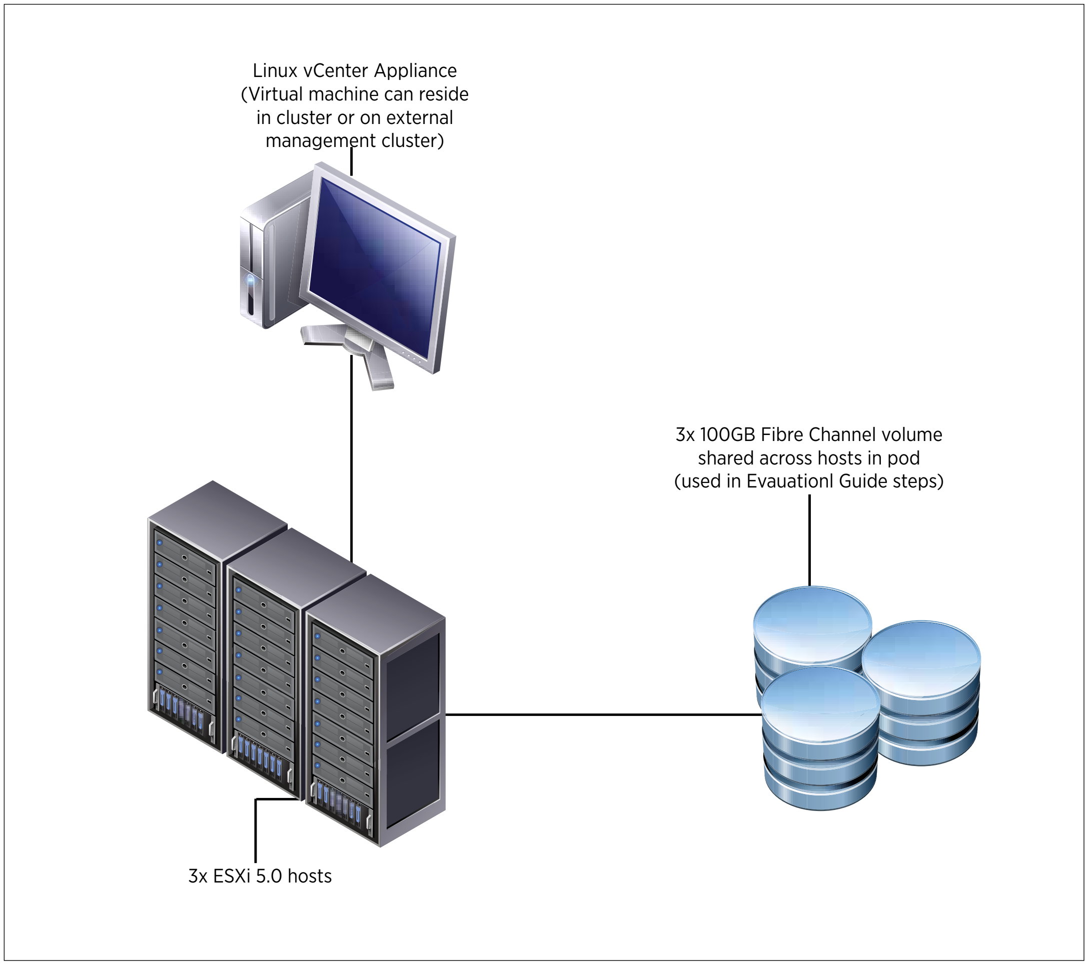
Linux vCenter Appliance
(Virtual machine can reside
in cluster or on external
management cluster)
3x 100GB Fibre Channel volume
shared across hosts in pod
(used in Evauationl Guide steps)
3x ESXi 5.0 hosts

Logical Network Setup

The *VMware vSphere 5.0 Evaluation Guide, Volume Two,* uses a simple network configuration consisting of three logical networks. The first is for vSphere management traffic, including vSphere High Availability (HA). The second is for vMotion and the third is for virtual machine traffic. Each logical network is configured as a port group on a standard switch, with a corresponding VLAN configured to provide physical isolation of the network traffic.

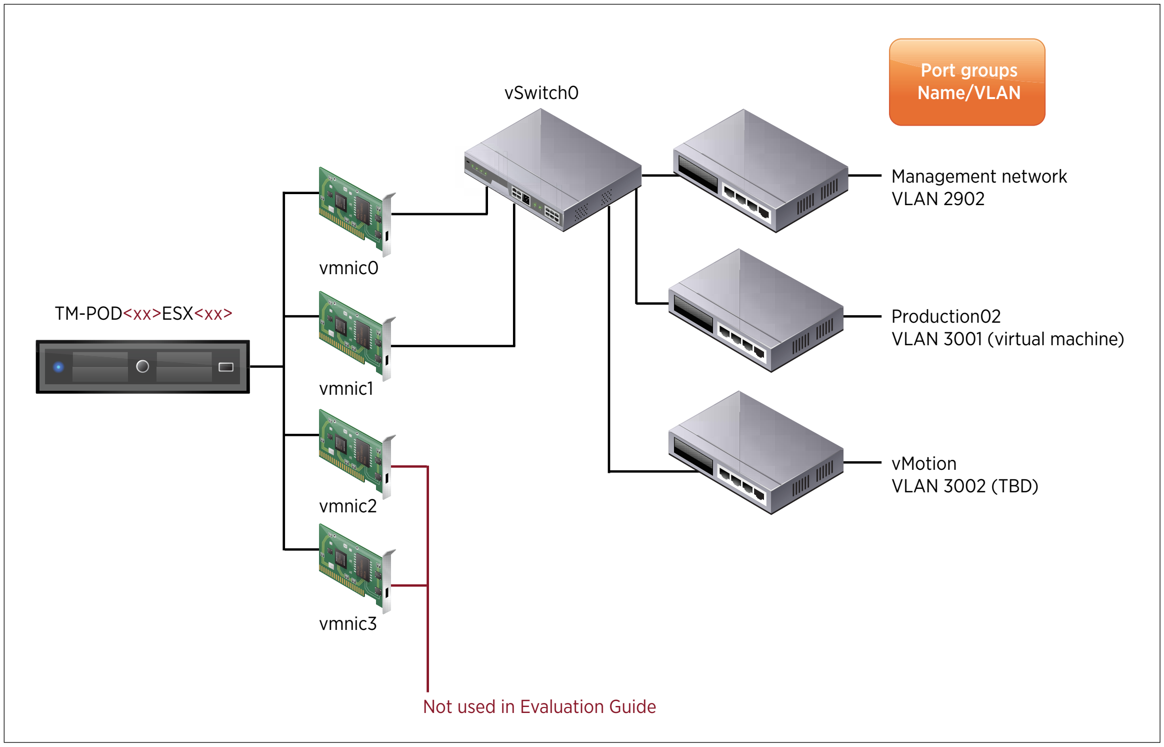

On the vSphere side, the network configuration looks like that in the following diagram.

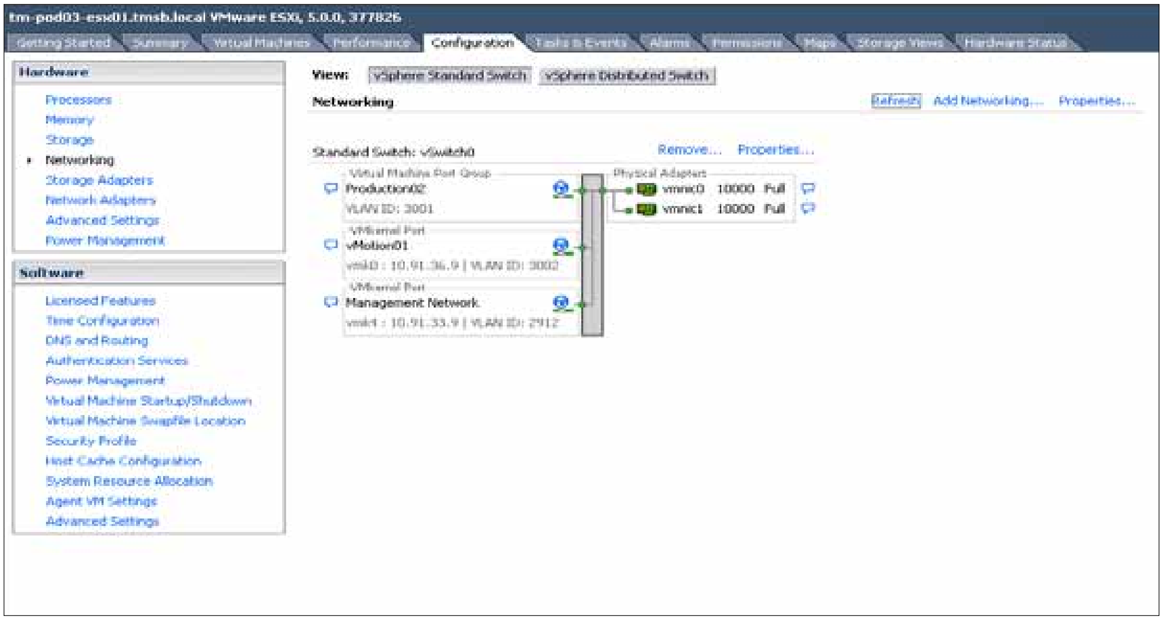

tm-pod03-esx01.tmsb.local VMware ESXi, 5.0.0, 377826

Networking

Standard Switch: vSwitch0

Virtual Machine Port Group
Production02
VLAN ID: 3001

VMkernel Port
vMotion01
vmk0 : 10.91.36.9 | VLAN ID: 3002

VMkernel Port
Management Network
vmk4 : 10.91.33.9 | VLAN ID: 2912

Physical Adapters
vmnic0 10000 Full
vmnic1 10000 Full

Storage Setup

The *VMware vSphere 5.0 Evaluation Guide, Volume Two,* uses a storage configuration consisting of three 100GB FC LUNs presented to each host, allowing the creation of three datastores.

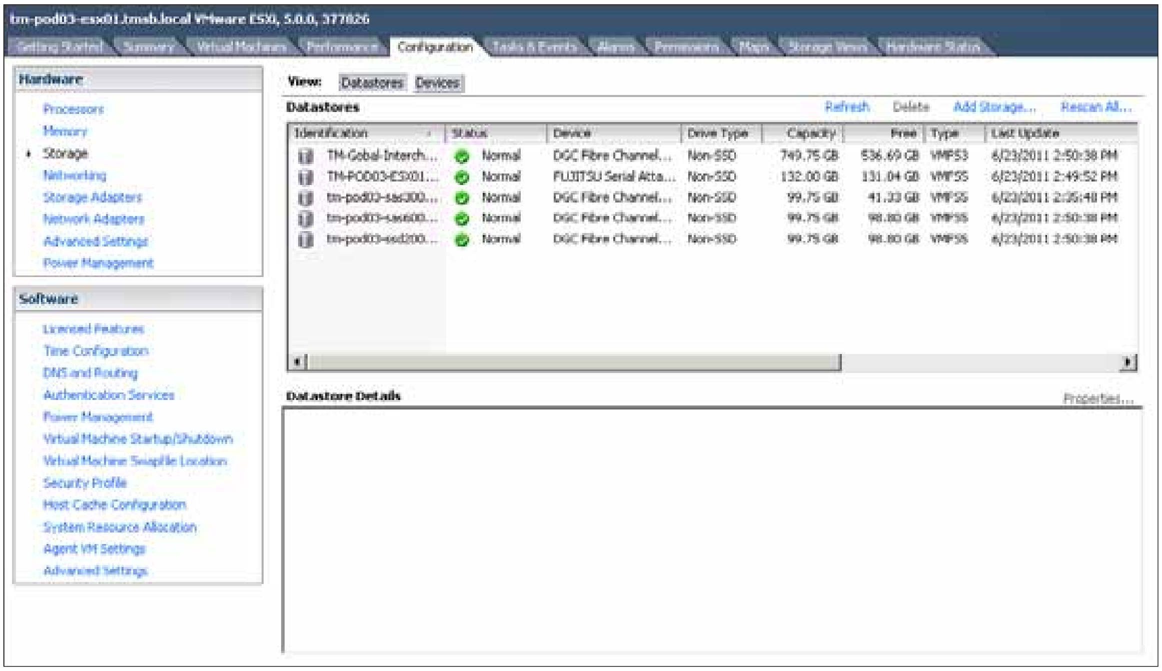

Datastores

Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update
TM-Gobal-Interch...	Normal	DGC Fibre Channel...	Non-SSD	749.75 GB	536.69 GB	VMFS3	6/23/2011 2:50:38 PM
TM-POD03-ESX01...	Normal	FUJITSU Serial Atta...	Non-SSD	132.00 GB	131.04 GB	VMFS5	6/23/2011 2:49:52 PM
tm-pod03-sas300...	Normal	DGC Fibre Channel...	Non-SSD	99.75 GB	41.33 GB	VMFS5	6/23/2011 2:35:48 PM
tm-pod03-sas600...	Normal	DGC Fibre Channel...	Non-SSD	99.75 GB	98.80 GB	VMFS5	6/23/2011 2:50:38 PM
tm-pod03-ssd200...	Normal	DGC Fibre Channel...	Non-SSD	99.75 GB	98.80 GB	VMFS5	6/23/2011 2:50:38 PM

Virtual Machine Setup

The *VMware vSphere 5.0 Evaluation Guide, Volume Two,* uses a total of seven virtual machines for testing. This volume will require Windows 2003 or Windows 2008 guest operating systems. It is up to the user to configure virtual machines that can be brought up to a running state for testing. The following diagram shows VM_01 through VM_06 configured in the technical marketing test lab.

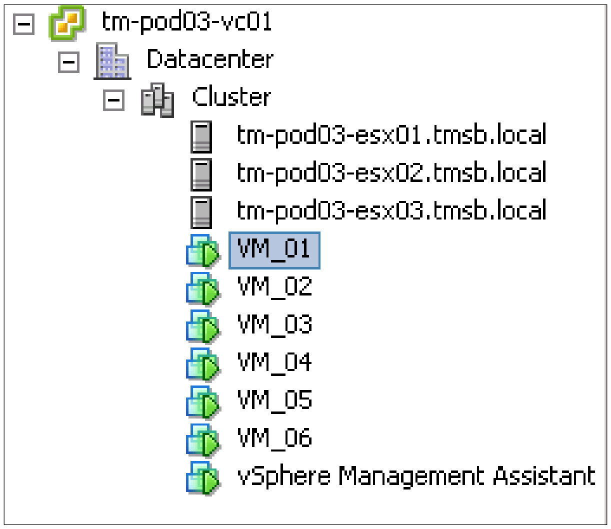
tm-pod03-vc01
Datacenter
Cluster
tm-pod03-esx01.tmsb.local
tm-pod03-esx02.tmsb.local
tm-pod03-esx03.tmsb.local
VM_01
VM_02
VM_03
VM_04
VM_05
VM_06
vSphere Management Assistant

VMware vSphere 5.0 Evaluation Guide, Volume Three, Worksheet

Use the following worksheet to organize your evaluation process.

HARDWARE CHECKLIST	
All hardware has been validated against the *VMware vSphere 5.0 Hardware Compatibility List* (HCL)	
Each server has 2 x 1GB or 2 x 10GB network adaptors connected to a common switch (they will be configured as a network adaptor team)	
Each server has the required HBA/network adaptor to access shared storage	

SOFTWARE CHECKLIST	
VMware ESXi™ installation media available	
VMware vCenter Server Appliance downloaded	
VMware vSphere Client installed	
ESXi host 1 name	
ESXi host 2 name	
ESXi host 3 name	
Subnet, netmask, and default gateway for management network	
Subnet, netmask, and default gateway for virtual machine network	
Subnet, netmask, and default gateway for vMotion network	

STORAGE CHECKLIST	
All servers can see at least three common 100GB LUNs (or NFS exports)	
Datastore 1 name	
Datastore 2 name	
Datastore 3 name	

vSphere Advanced Network Features

With the release of vSphere 5.0, VMware brings a number of powerful new features and enhancements to the networking capabilities of the vSphere platform. There are two broad categories of enhancements:

- **Enhanced network I/O control:** vSphere 5.0 builds on network I/O control (NIOC) to allow user-defined network resource pools, enabling multitenancy deployment, and to bridge virtual and physical infrastructure quality of service with per–resource pool 802.1 tagging.
- **VDS improvements:** vSphere 5.0 provides improved visibility into virtual machine traffic through NetFlow and enhances monitoring and troubleshooting capabilities through the Switch Port Analyzer (SPAN) and LLDP.

In this Evaluation Guide volume, you will learn how to configure and test the following new networking features that are available on the VDS:

- NetFlow
- Port mirroring
- NIOC – user-defined resource pools

First, you will configure the VDS and then enable and test each of the new features one at a time.

vSphere Distributed Switch Configuration

Overview

The VDS simplifies virtual machine networking by enabling you to set up virtual machine networking for your entire datacenter from a centralized interface. A single VDS spans many VMware ESX®/ESXi hosts and aggregates networking to a centralized datacenter level.

Configuring the VDS

In this exercise, you will configure a VDS that spans across three hosts and provides a single network management interface to configure network parameters. The VDS configuration can be accomplished in either of two ways:

1. Using only the VDS user interface
2. Using a combination of the VDS and host profiles

This use case describes the configuration through the user interface. For more details on host profile–based migration, refer to *VMware vSphere Distributed Switch: Migration and Configuration,* available on vmware.com.

Evaluation Environment for the VDS Configuration

The evaluation environment consists of the following components, as shown in Figure 1:

1. A single vSphere datacenter (Datacenter)
2. Three ESXi 5.0 servers (tm-pod03-esx01.tmsb.local, tm-pod03-esx02.tmsb.local, tm-pod03-esx03.tmsb.local) in a cluster (Cluster)
3. A virtual network environment supporting the following different traffic types:
    a. Production02 (virtual machine traffic)
    b. vMotion traffic
    c. Management network traffic
    d. User-defined tenant traffic
4. Six virtual machines with a single vNIC attachment to the VDS
5. A vSphere Management Assistant, providing remote console access

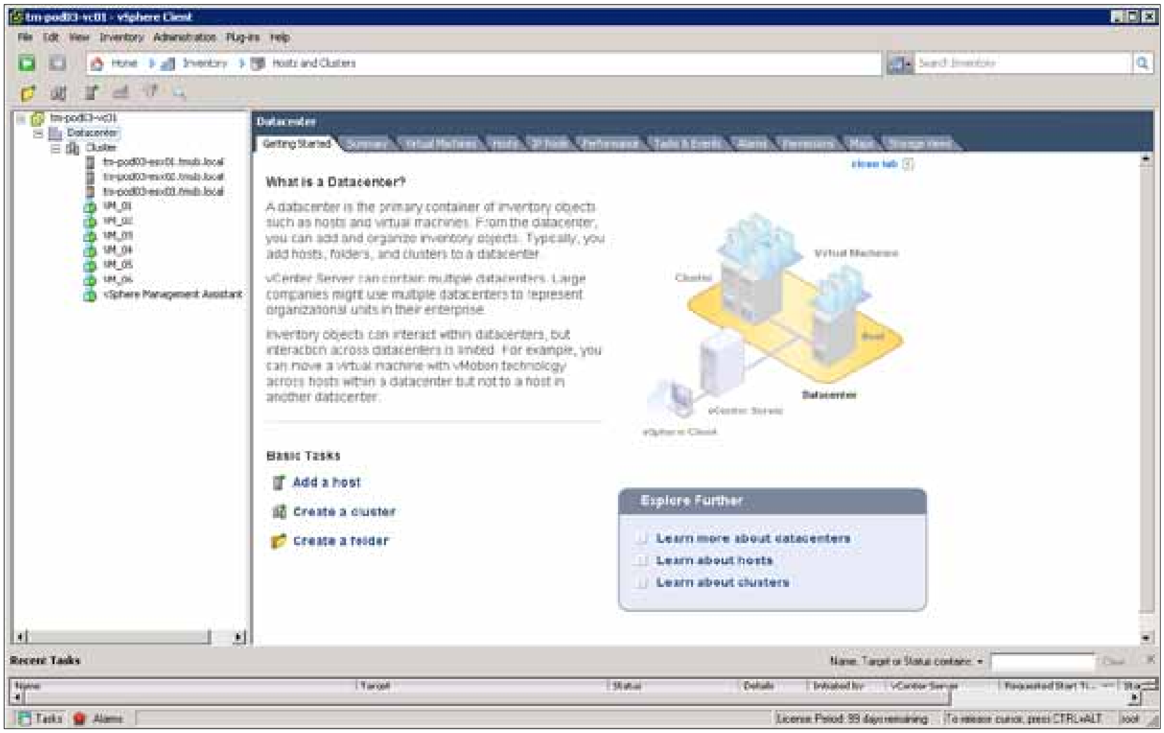

What is a Datacenter?

A datacenter is the primary container of inventory objects such as hosts and virtual machines. From the datacenter, you can add and organize inventory objects. Typically, you add hosts, folders, and clusters to a datacenter.

vCenter Server can contain multiple datacenters. Large companies might use multiple datacenters to represent organizational units in their enterprise.

Inventory objects can interact within datacenters, but interaction across datacenters is limited. For example, you can move a virtual machine with vMotion technology across hosts within a datacenter but not to a host in another datacenter.

Basic Tasks

Figure 1. VMware vSphere Client View with Hosts and Clusters View

Creating a VDS

The VDS is created at the datacenter level in the vSphere environment. As shown in Figure 2 under the **Networking** view, in the **Datacenter** level, you can click **Create VDS button** to configure a new VDS.

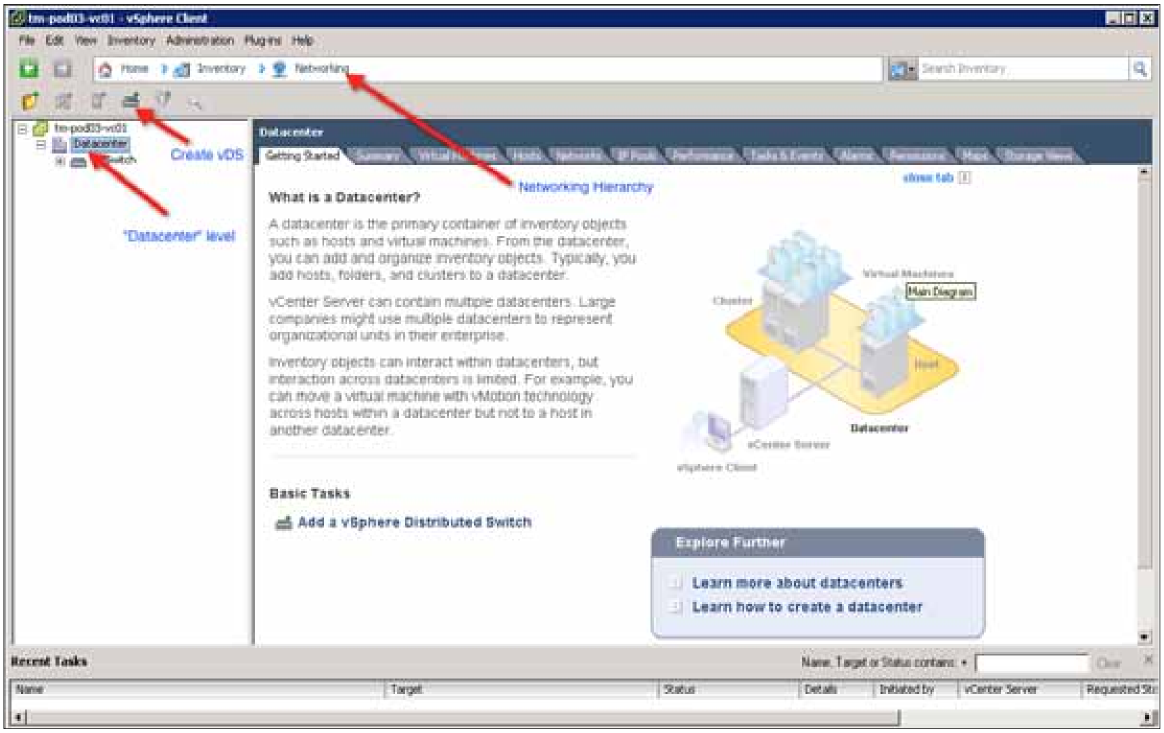

What is a Datacenter?

A datacenter is the primary container of inventory objects such as hosts and virtual machines. From the datacenter, you can add and organize inventory objects. Typically, you add hosts, folders, and clusters to a datacenter.

vCenter Server can contain multiple datacenters. Large companies might use multiple datacenters to represent organizational units in their enterprise.

Inventory objects can interact within datacenters, but interaction across datacenters is limited. For example, you can move a virtual machine with vMotion technology across hosts within a datacenter but not to a host in another datacenter.

Basic Tasks

Add a vSphere Distributed Switch

Explore Further

Learn more about datacenters

Learn how to create a datacenter

Figure 2. Starting Point in Creating a VDS

After the VDS has been created, the Networking panel will show a dvSwitch (the default name), and an uplink group for the uplinks (in this example, it is named dvSwitch-DVUplinks-26).

An uplink group provides a policy template for the uplinks on that VDS. Security policies, VLAN trunk ranges, traffic shaping, and teaming/failover settings can be set at the uplink group level for the entire VDS. In this example, the environment shown in Figure 3, the uplink group consists of four uplinks (dvUplink1 to dvUplink4) and the first two uplinks are connected to vmnic0 and vmnic1 on the host. Depending on the number of vmnics available on a host, you can decide the number of dvUplink ports.

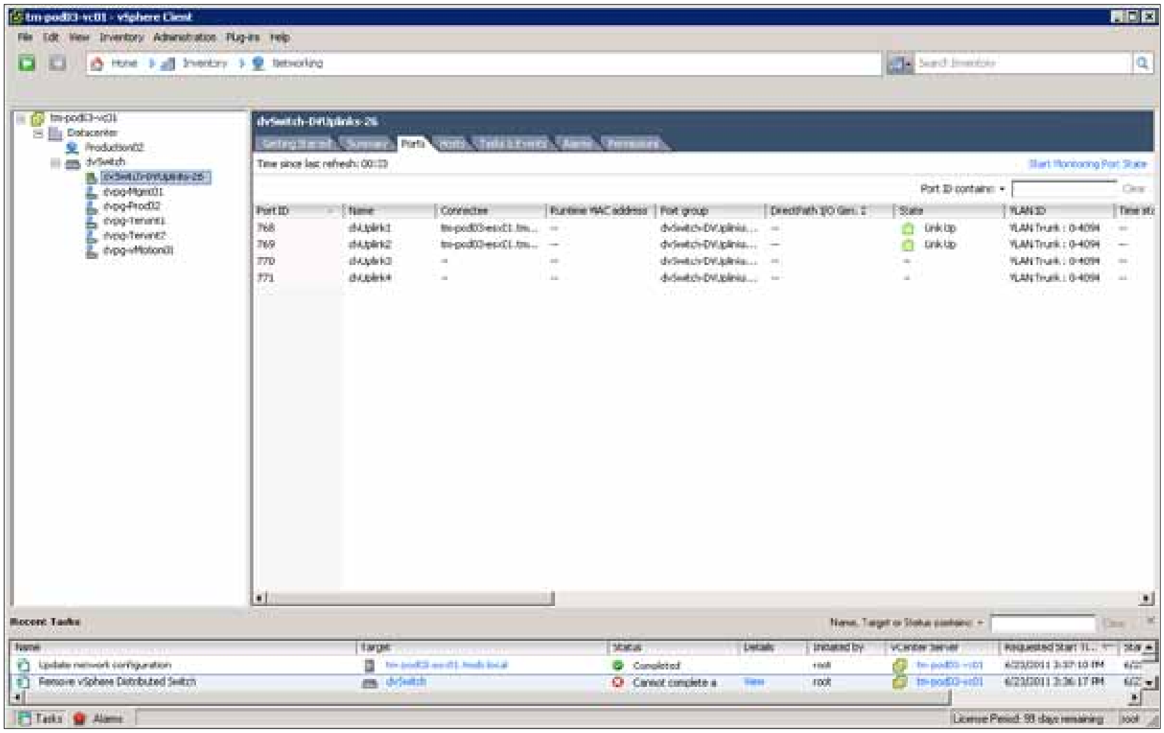

Time since last refresh: 00:33

Port ID	Name	Connectee	Runtime MAC address	Port group	DirectPath I/O Gen. 2	State	VLAN ID	Time st
768	dvUplink1	tm-pod03-esx01.tm...	--	dvSwitch-DVUplinks...	--	Link Up	VLAN Trunk : 0-4094	--
769	dvUplink2	tm-pod03-esx01.tm...	--	dvSwitch-DVUplinks...	--	Link Up	VLAN Trunk : 0-4094	--
770	dvUplink3	--	--	dvSwitch-DVUplinks...	--	--	VLAN Trunk : 0-4094	--
771	dvUplink4	--	--	dvSwitch-DVUplinks...	--	--	VLAN Trunk : 0-4094	--

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Star
Update network configuration	tm-pod03-esx01.tmdb.local	Completed		root	tm-pod03-vc01	6/23/2011 3:37:10 PM	6/2
Remove vSphere Distributed Switch	dvSwitch	Cannot complete a	View	root	tm-pod03-vc01	6/23/2011 3:36:17 PM	6/2

Figure 3. Uplink Details

Configuring Distributed Virtual Port Groups

In this step, you will create distributed virtual port groups (DV port groups) on the VDS according to the evaluation environment requirements. The evaluation environment needs support for different types of system and user network traffic types. These traffic types are isolated from each other using different VLANs. Table 1 shows the different traffic types and corresponding port group, VLAN, and IP subnet information. Check the VLAN policies with your network administrators when configuring different port groups.

TRAFFIC TYPE	PORT GROUP NAME	VLAN	IP NETWORK
Virtual Machine Traffic	Prod02	3001	10.91.35.0 (DHCP allocation)
ESXi Host Management Traffic	Mgmt01	2912	10.91.33.0
vMotion Traffic	vMotion01	3002	10.91.36.9
Tenant1 Traffic	Tenant1	3001	10.91.35.0
Tenant2 Traffic	Tenant2	3001	10.91.35.0

Table 1. Traffic Types and VLAN Assignments in the Example Environment

Once you have the table of port groups and associated VLAN and IP mapping, you can start creating the individual DV port groups.

1. From the **Home > Inventory > Networking** view, select the VDS. In this example environment, the VDS is labeled **dvSwitch.**
2. Click **New Port Group.** Figure 4 shows the first panel in creating the dvpg-Mgmt01 DV port group. Note the number of ports. This defaults to 128 and is the number of ports that this DV port group will allow once created. This also means that up to 128 virtual machines can use this DV port group. You can modify this to a higher number based on the number of virtual machines you want to support within a single DV port group.

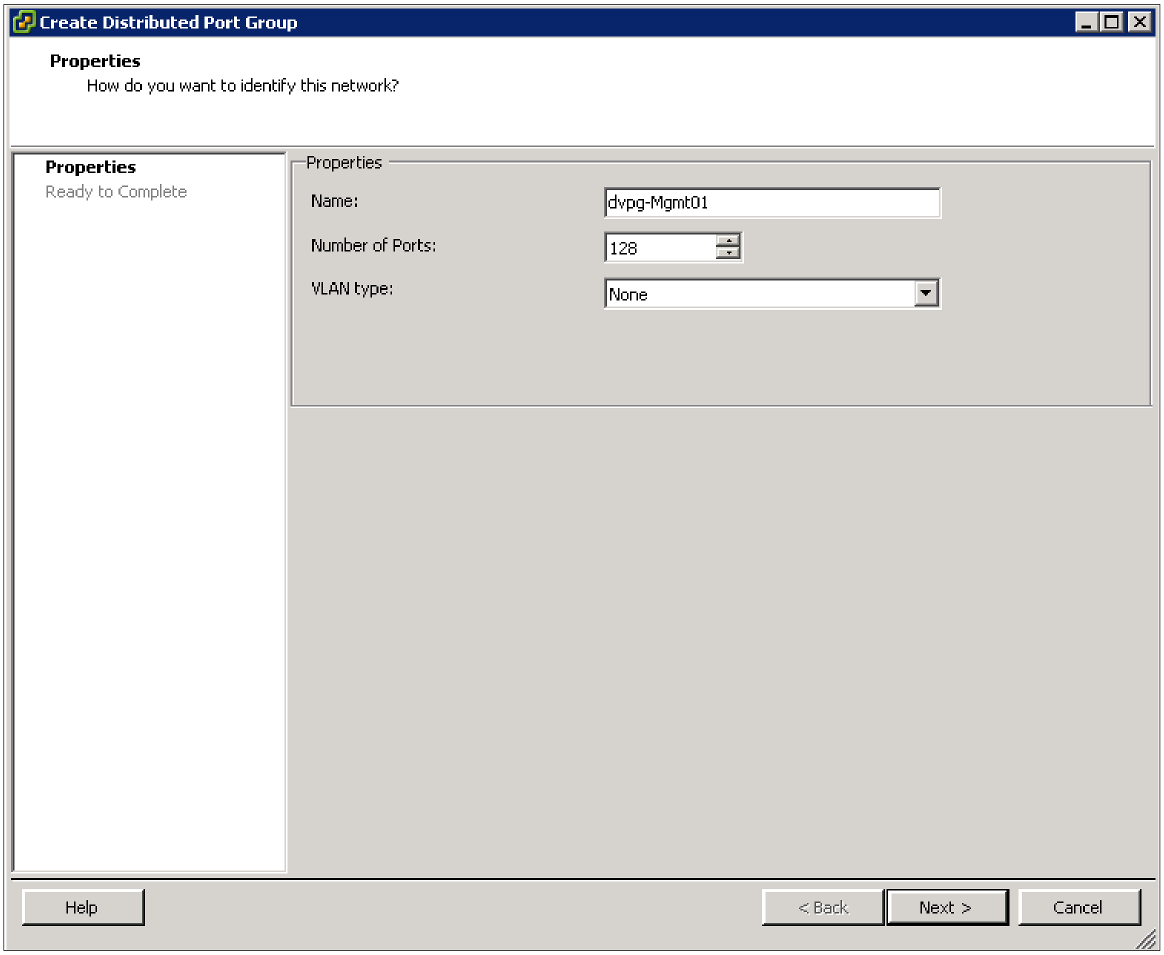
Create Distributed Port Group
Properties
How do you want to identify this network?
Properties
Ready to Complete
Properties
Help
< Back
Next >
Cancel

Name: dvpg-Mgmt01

Number of Ports: 128

VLAN type: None

Figure 4. Creating a Port Group

3. Continue creating the DV port groups according to the table and enter the VLAN and IP information during the configuration. The Tenant1 and Tenant2 port groups are used during the testing of the enhanced Network I/O Control (NIOC) feature. After creating the DV port groups, the VDS panel should look like it does in Figure 5.

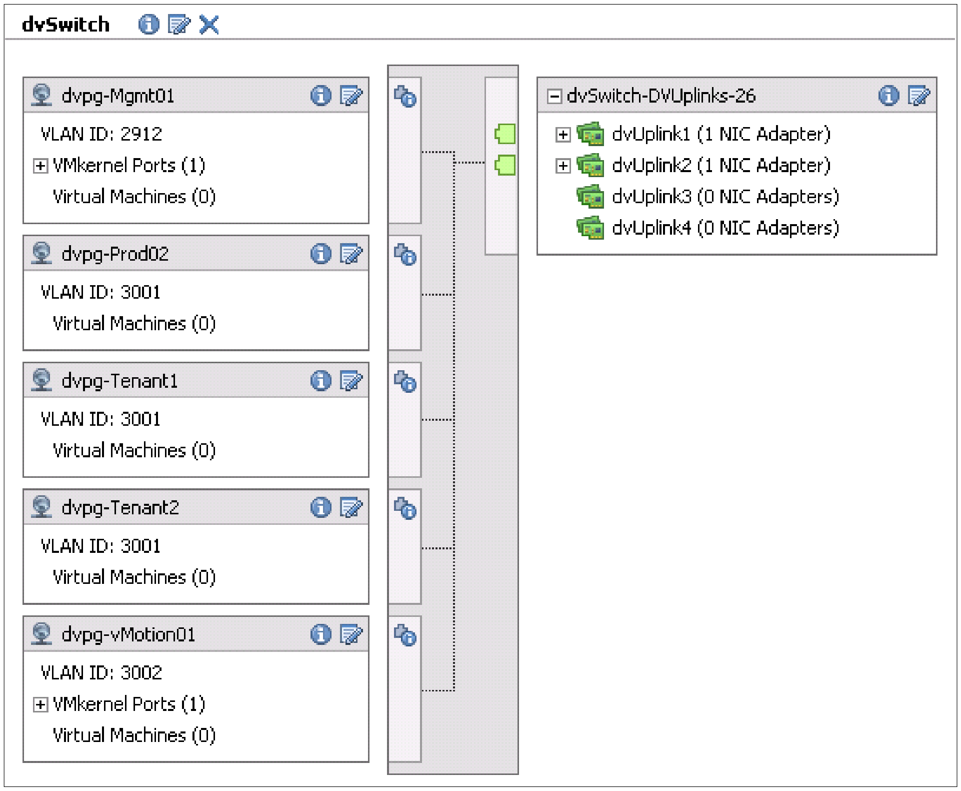

dvSwitch

dvpg-Mgmt01
VLAN ID: 2912
⊞ VMkernel Ports (1)
Virtual Machines (0)

dvpg-Prod02
VLAN ID: 3001
Virtual Machines (0)

dvpg-Tenant1
VLAN ID: 3001
Virtual Machines (0)

dvpg-Tenant2
VLAN ID: 3001
Virtual Machines (0)

dvpg-vMotion01
VLAN ID: 3002
⊞ VMkernel Ports (1)
Virtual Machines (0)

⊟ dvSwitch-DVUplinks-26
⊞ dvUplink1 (1 NIC Adapter)
⊞ dvUplink2 (1 NIC Adapter)
dvUplink3 (0 NIC Adapters)
dvUplink4 (0 NIC Adapters)

Figure 5. View of the VDS Panel After Creating DV Port Groups

Adding a Host to a VDS

After creating a VDS, you can migrate hosts and physical adapters to this VDS. In this step, you will migrate the standard switch environment of the host to the VDS and DV port groups created in steps 1 and 2.

Figure 6 shows the standard switch configuration of a host that is going to be migrated to the VDS. The switch configuration shows the current port groups and physical adapters. Migration of these port groups and physical adapters is carried out through the following steps:

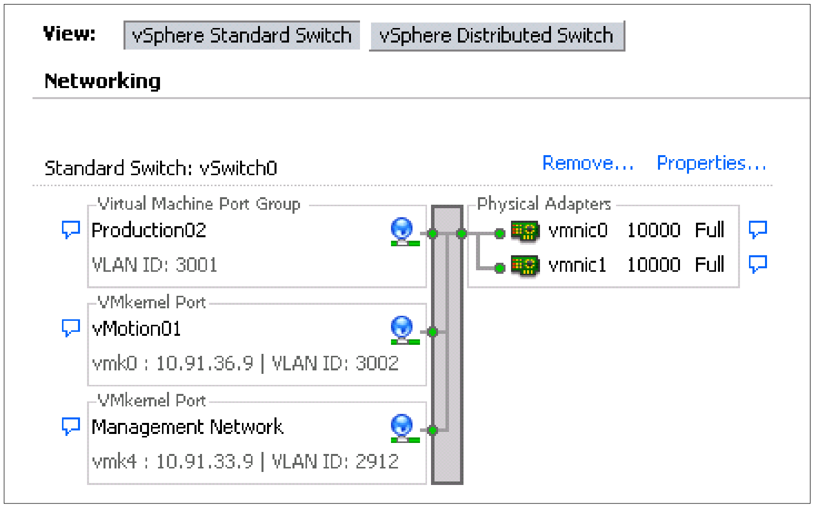

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0 — Remove... Properties...

Virtual Machine Port Group
Production02
VLAN ID: 3001

VMkernel Port
vMotion01
vmk0 : 10.91.36.9 | VLAN ID: 3002

VMkernel Port
Management Network
vmk4 : 10.91.33.9 | VLAN ID: 2912

Physical Adapters
vmnic0 10000 Full
vmnic1 10000 Full

Figure 6. Standard Switch Configuration of a Host to Be Migrated to the VDS

1. Switch to the **Home > Inventory > Networking** view.
2. Right-click the dvSwitch and select **Add Host..** See Figure 7.

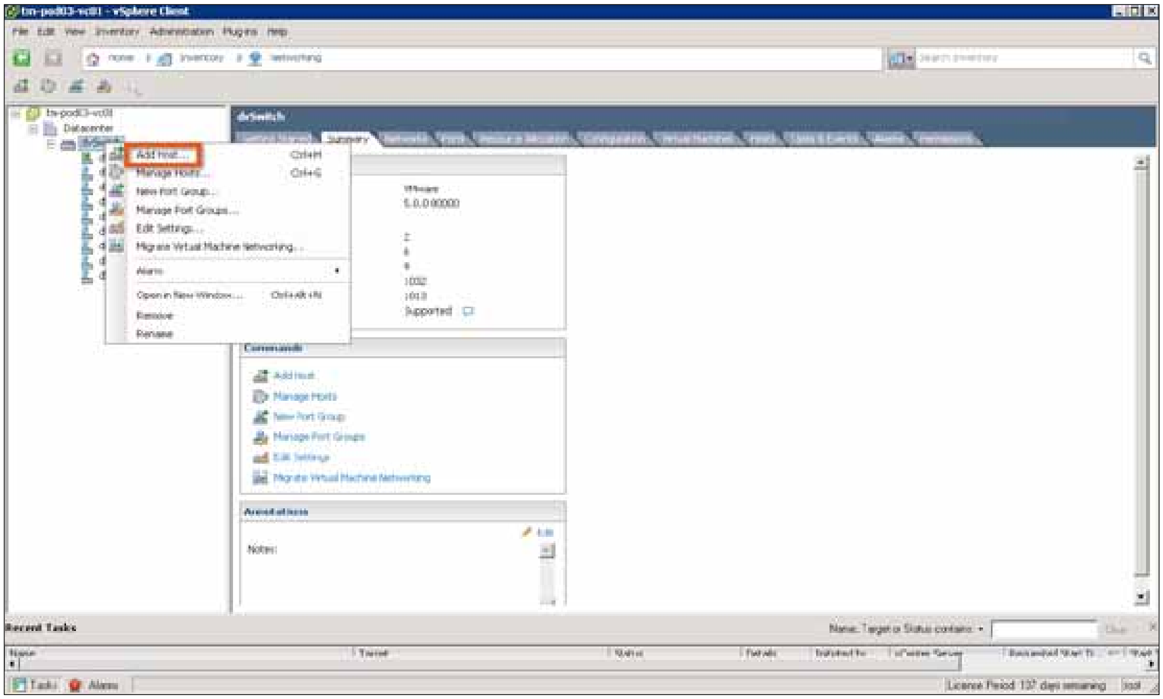

VMware
5.0.0 00000
Supported

Figure 7. Preparing to Add a Host

3. Select the hosts to migrate to the VDS as shown in Figure 8. In your environment, if you have three hosts in a cluster you will see those hosts along with their physical adapter listed as shown in the following panel. In this example, choose to migrate vmnic0 and vmnic1 from the standard switch, vSwitch0, to the uplink port group of the VDS dvSwitch-DVUplinks-26. Click **Next.**

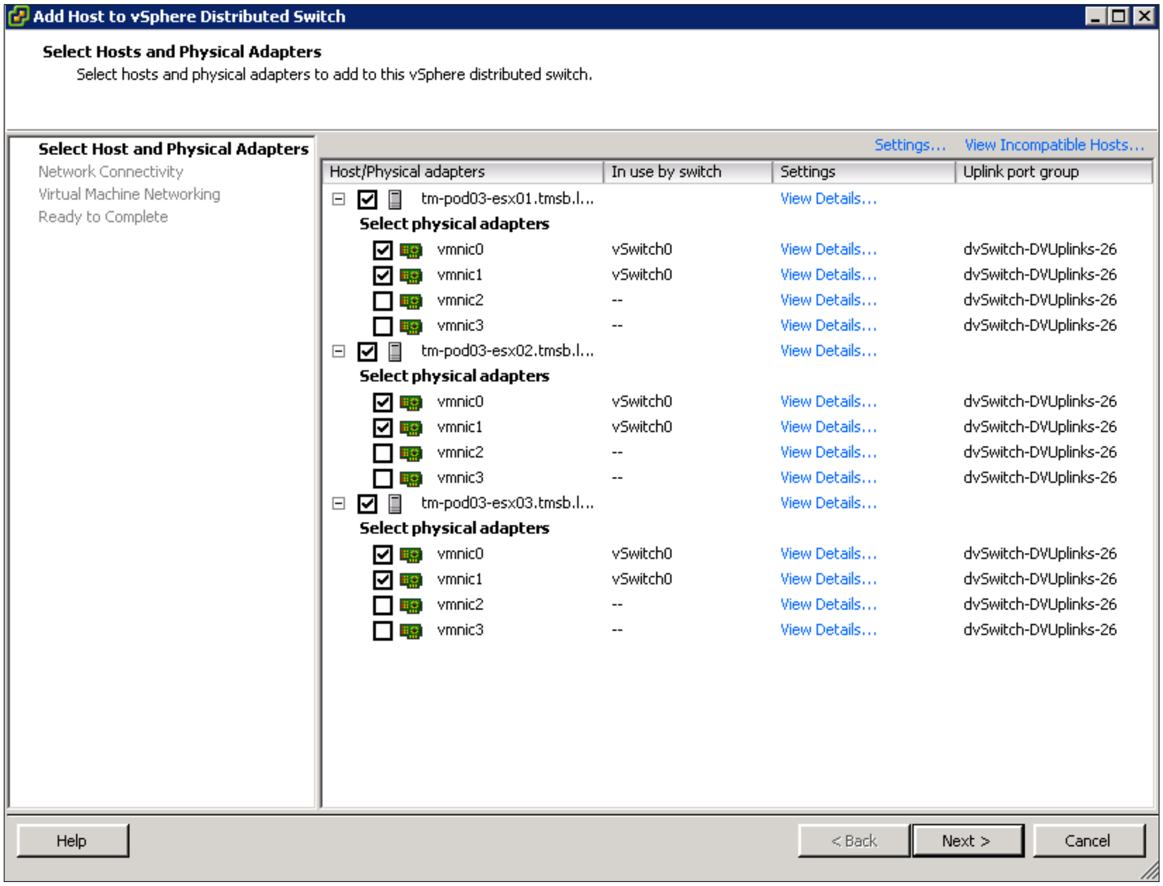
Select Host and Physical Adapters
Network Connectivity
Virtual Machine Networking
Ready to Complete
Settings...
View Incompatible Hosts...
Help
< Back
Next >
Cancel

Add Host to vSphere Distributed Switch

Select Hosts and Physical Adapters

Select hosts and physical adapters to add to this vSphere distributed switch.

Host/Physical adapters	In use by switch	Settings	Uplink port group
⊟ ☑ tm-pod03-esx01.tmsb.l...		View Details...	
Select physical adapters			
☑ vmnic0	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☑ vmnic1	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☐ vmnic2	--	View Details...	dvSwitch-DVUplinks-26
☐ vmnic3	--	View Details...	dvSwitch-DVUplinks-26
⊟ ☑ tm-pod03-esx02.tmsb.l...		View Details...	
Select physical adapters			
☑ vmnic0	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☑ vmnic1	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☐ vmnic2	--	View Details...	dvSwitch-DVUplinks-26
☐ vmnic3	--	View Details...	dvSwitch-DVUplinks-26
⊟ ☑ tm-pod03-esx03.tmsb.l...		View Details...	
Select physical adapters			
☑ vmnic0	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☑ vmnic1	vSwitch0	View Details...	dvSwitch-DVUplinks-26
☐ vmnic2	--	View Details...	dvSwitch-DVUplinks-26
☐ vmnic3	--	View Details...	dvSwitch-DVUplinks-26

Figure 8. Selecting Hosts to Migrate

4. Match up the port groups on the standard switch with the DV port groups of the VDS. In Figure 9, the Management Network port group on the standard switch is matched with the dvpg-Mgmt01 DV port group on the VDS.

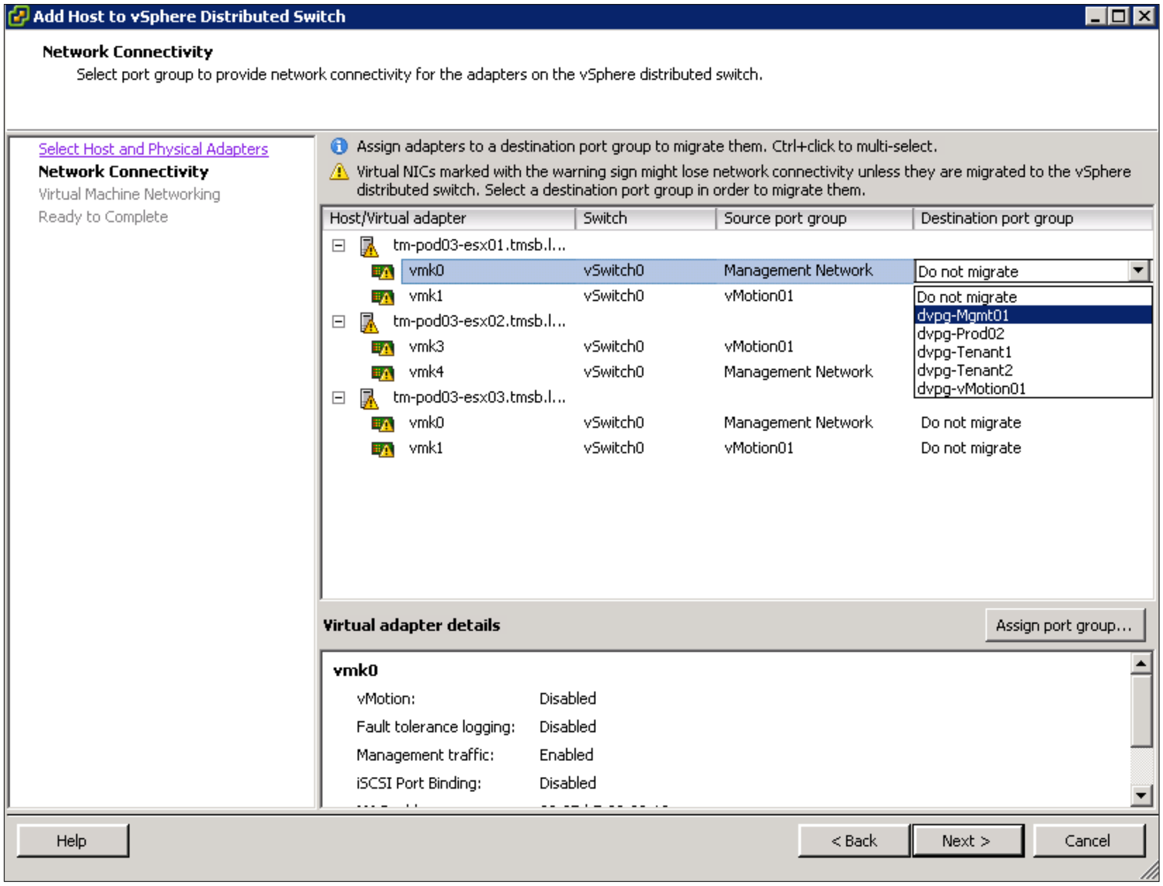

Add Host to vSphere Distributed Switch

Network Connectivity

Select port group to provide network connectivity for the adapters on the vSphere distributed switch.

Assign adapters to a destination port group to migrate them. Ctrl+click to multi-select.

⚠ Virtual NICs marked with the warning sign might lose network connectivity unless they are migrated to the vSphere distributed switch. Select a destination port group in order to migrate them.

Host/Virtual adapter	Switch	Source port group	Destination port group
⊟ tm-pod03-esx01.tmsb.l...			
vmk0	vSwitch0	Management Network	Do not migrate Do not migrate dvpg-Mgmt01 dvpg-Prod02 dvpg-Tenant1 dvpg-Tenant2 dvpg-vMotion01
vmk1	vSwitch0	vMotion01	
⊟ tm-pod03-esx02.tmsb.l...			
vmk3	vSwitch0	vMotion01	
vmk4	vSwitch0	Management Network	
⊟ tm-pod03-esx03.tmsb.l...			
vmk0	vSwitch0	Management Network	Do not migrate
vmk1	vSwitch0	vMotion01	Do not migrate

Virtual adapter details

vmk0

vMotion:	Disabled
Fault tolerance logging:	Disabled
Management traffic:	Enabled
iSCSI Port Binding:	Disabled

Figure 9. Migration of the Port Group

5. Repeat the matching process of port groups from the standard switch to DV port groups of the VDS. Double-check the matchups before starting the migration action by clicking **Next** in the panel shown in Figure 10.

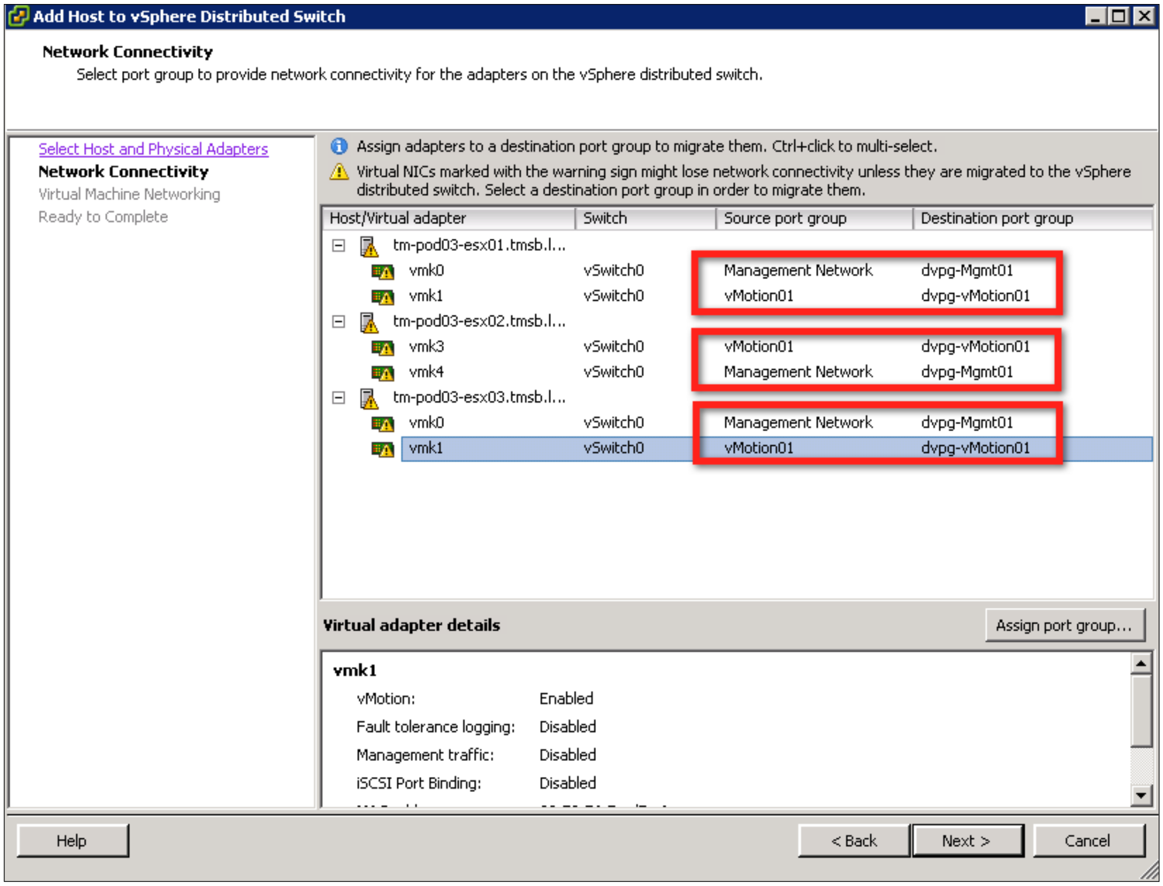
Add Host to vSphere Distributed Switch
Select Host and Physical Adapters
Network Connectivity
Virtual Machine Networking
Ready to Complete
Assign port group...
Help
< Back
Next >
Cancel

Network Connectivity

Select port group to provide network connectivity for the adapters on the vSphere distributed switch.

Assign adapters to a destination port group to migrate them. Ctrl+click to multi-select.

Virtual NICs marked with the warning sign might lose network connectivity unless they are migrated to the vSphere distributed switch. Select a destination port group in order to migrate them.

Host/Virtual adapter	Switch	Source port group	Destination port group
tm-pod03-esx01.tmsb.l...			
vmk0	vSwitch0	Management Network	dvpg-Mgmt01
vmk1	vSwitch0	vMotion01	dvpg-vMotion01
tm-pod03-esx02.tmsb.l...			
vmk3	vSwitch0	vMotion01	dvpg-vMotion01
vmk4	vSwitch0	Management Network	dvpg-Mgmt01
tm-pod03-esx03.tmsb.l...			
vmk0	vSwitch0	Management Network	dvpg-Mgmt01
vmk1	vSwitch0	vMotion01	dvpg-vMotion01

Virtual adapter details

vmk1

vMotion:	Enabled
Fault tolerance logging:	Disabled
Management traffic:	Disabled
iSCSI Port Binding:	Disabled

Figure 10. Association of Standard Switch Port Groups to VDS DV Port Groups

This step does not provide the option of transferring the port groups for the virtual machines (Prod02).

After you click **Next** in the step 5, you will be presented with the panel shown in Figure 11. In this panel, you have an option to migrate your viritual machines from the standard switch port groups to VDS DV port groups. In this example environment, there are six virtual machines that are running on three hosts. In your environment, if you have virtual machines running on the hosts, you can use this screen to migrate those virtual machines to the appropriate VDS DV port groups. Select **Migrate virtual machine networking** as in Figure 11.

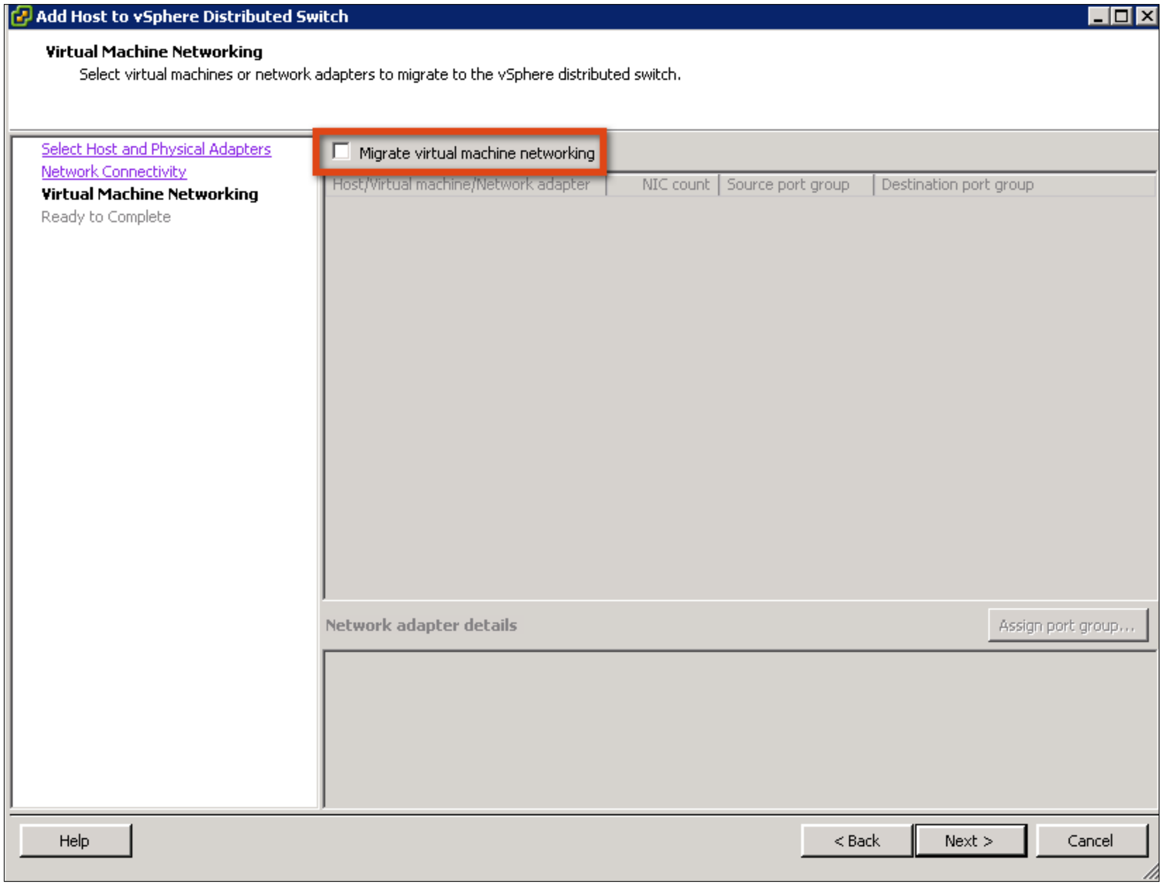

Add Host to vSphere Distributed Switch
Virtual Machine Networking
Select virtual machines or network adapters to migrate to the vSphere distributed switch.
Select Host and Physical Adapters
Network Connectivity
Virtual Machine Networking
Ready to Complete
□ Migrate virtual machine networking
Host/Virtual machine/Network adapter
NIC count
Source port group
Destination port group
Network adapter details
Assign port group...
Help
< Back
Next >
Cancel

Figure 11. Starting the Migration of Virtual Machines to the VDS

6. You will be prompted with the following screen, shown in Figure 12, with all the available virtual machines listed, along with the option to migrate them to a destination port group.

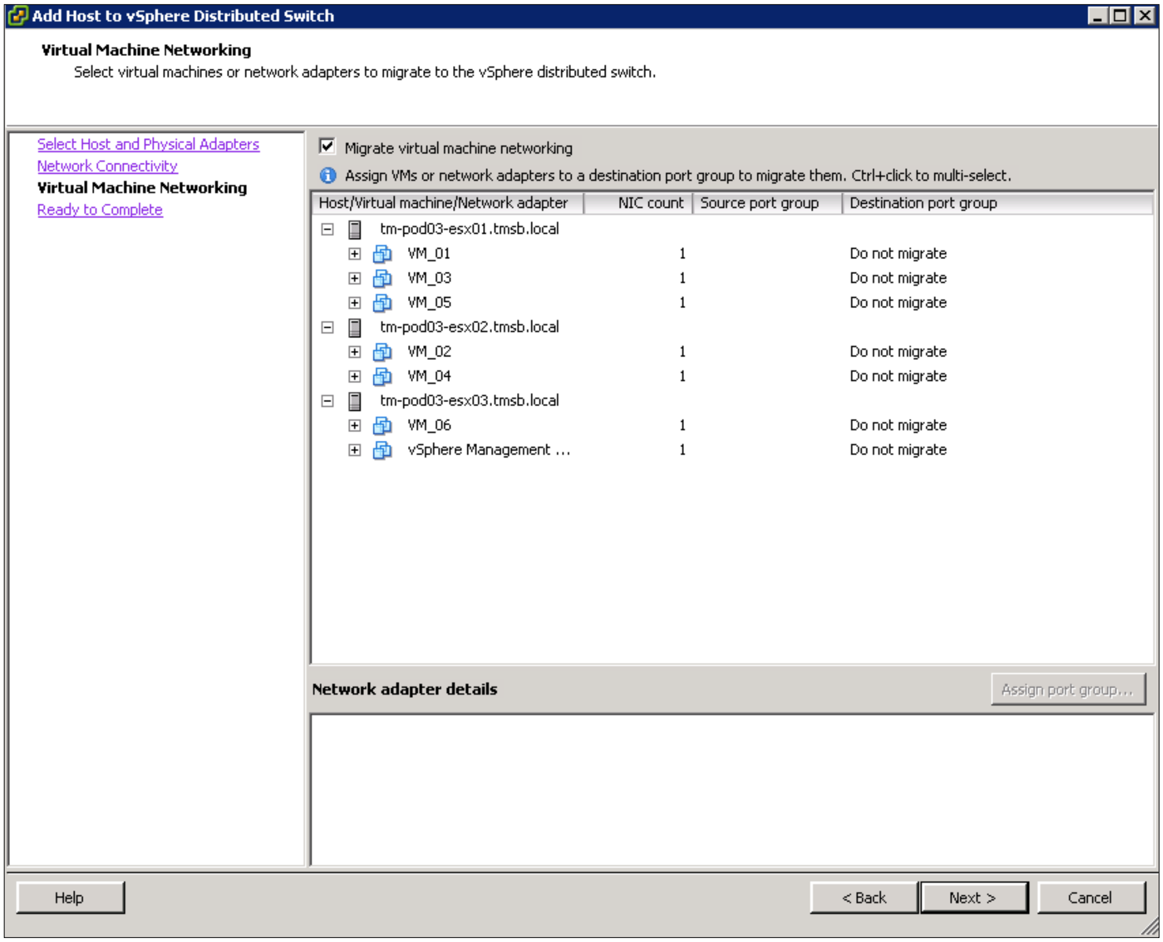

Add Host to vSphere Distributed Switch

Virtual Machine Networking

Select virtual machines or network adapters to migrate to the vSphere distributed switch.

☑ Migrate virtual machine networking

Assign VMs or network adapters to a destination port group to migrate them. Ctrl+click to multi-select.

Host/Virtual machine/Network adapter	NIC count	Source port group	Destination port group
⊟ tm-pod03-esx01.tmsb.local			
⊞ VM_01	1		Do not migrate
⊞ VM_03	1		Do not migrate
⊞ VM_05	1		Do not migrate
⊟ tm-pod03-esx02.tmsb.local			
⊞ VM_02	1		Do not migrate
⊞ VM_04	1		Do not migrate
⊟ tm-pod03-esx03.tmsb.local			
⊞ VM_06	1		Do not migrate
⊞ vSphere Management ...	1		Do not migrate

Network adapter details

Figure 12. List of Virtual Machines to Migrate

7. Using the drop-down menu, select the DV port group to which you want the virtual machine connected. In this example environment, all the virtual machines are migrated to the **dvpg-Prod02** port group, as shown in Figure 13 and Figure 14.

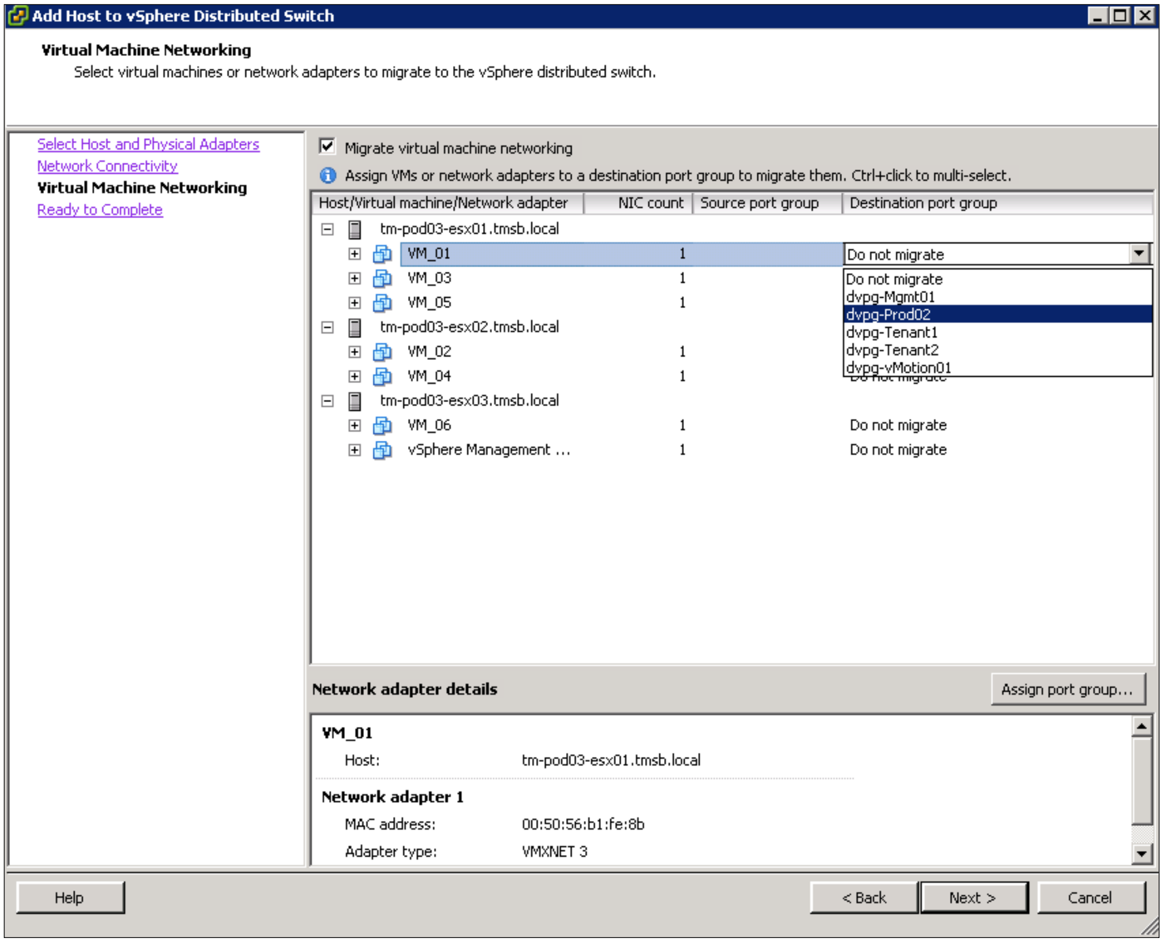

Add Host to vSphere Distributed Switch

Virtual Machine Networking

Select virtual machines or network adapters to migrate to the vSphere distributed switch.

Select Host and Physical Adapters
Network Connectivity
Virtual Machine Networking
Ready to Complete

☑ Migrate virtual machine networking

Assign VMs or network adapters to a destination port group to migrate them. Ctrl+click to multi-select.

Host/Virtual machine/Network adapter	NIC count	Source port group	Destination port group
tm-pod03-esx01.tmsb.local			
VM_01	1		Do not migrate
VM_03	1		
VM_05	1		
tm-pod03-esx02.tmsb.local			
VM_02	1		
VM_04	1		Do not migrate
tm-pod03-esx03.tmsb.local			
VM_06	1		Do not migrate
vSphere Management …	1		Do not migrate

Network adapter details

VM_01

Host: tm-pod03-esx01.tmsb.local

Network adapter 1

MAC address: 00:50:56:b1:fe:8b

Adapter type: VMXNET 3

Figure 13. Choosing a Destination DV Port Group for Each Virtual Machine

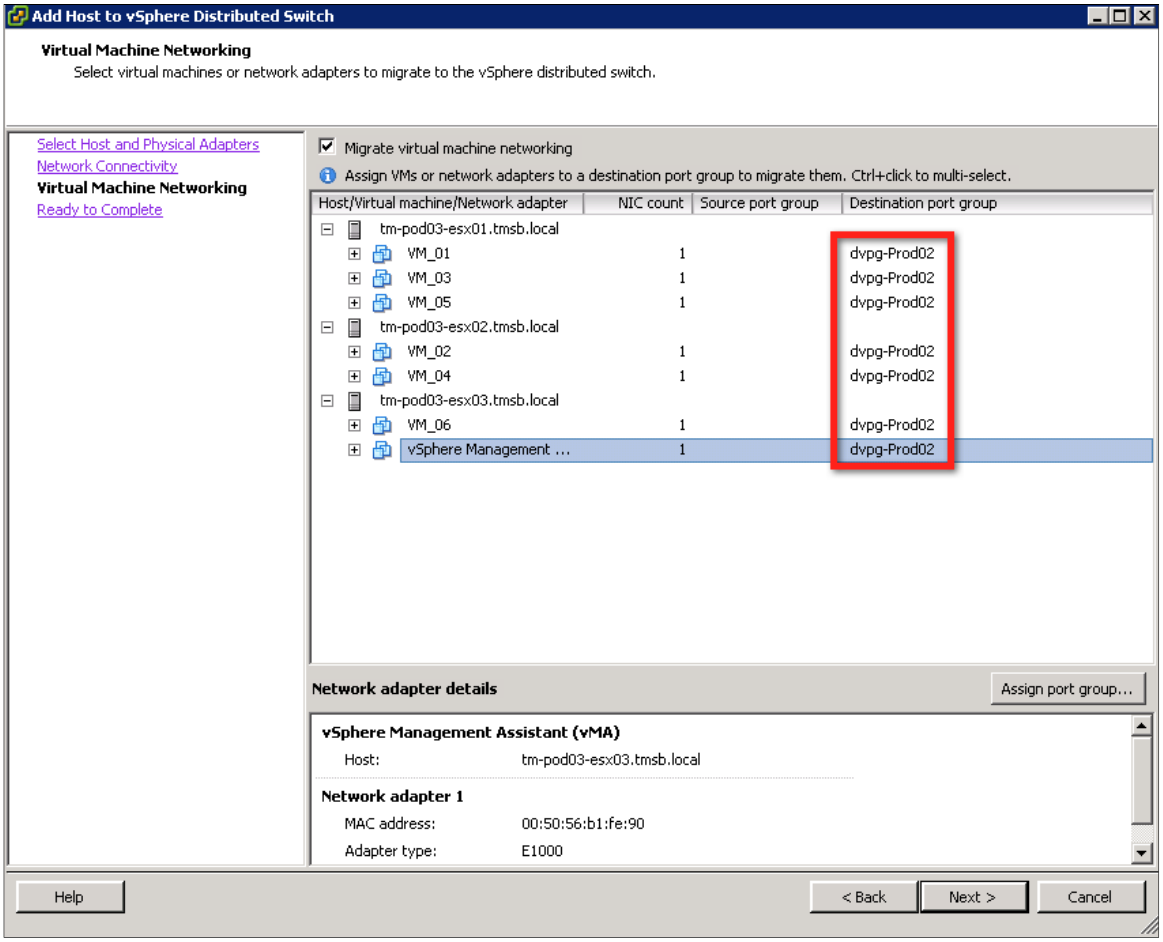
Add Host to vSphere Distributed Switch
Select Host and Physical Adapters
Network Connectivity
Virtual Machine Networking
Ready to Complete
Assign port group...
Help
< Back
Next >
Cancel

Virtual Machine Networking

Select virtual machines or network adapters to migrate to the vSphere distributed switch.

☑ Migrate virtual machine networking

Assign VMs or network adapters to a destination port group to migrate them. Ctrl+click to multi-select.

Host/Virtual machine/Network adapter	NIC count	Source port group	Destination port group
tm-pod03-esx01.tmsb.local			
VM_01	1		dvpg-Prod02
VM_03	1		dvpg-Prod02
VM_05	1		dvpg-Prod02
tm-pod03-esx02.tmsb.local			
VM_02	1		dvpg-Prod02
VM_04	1		dvpg-Prod02
tm-pod03-esx03.tmsb.local			
VM_06	1		dvpg-Prod02
vSphere Management ...	1		dvpg-Prod02

Network adapter details

vSphere Management Assistant (vMA)

Host: tm-pod03-esx03.tmsb.local

Network adapter 1

MAC address: 00:50:56:b1:fe:90

Adapter type: E1000

Figure 14. Migration of Virtual Machines to the DV Port Group

8. Click **Next** and **Finish,** and wait for the operation to complete. Track the status in the Recent Tasks panel at the bottom of the vSphere Client panel. The VDS should now appear as shown in Figure 15.

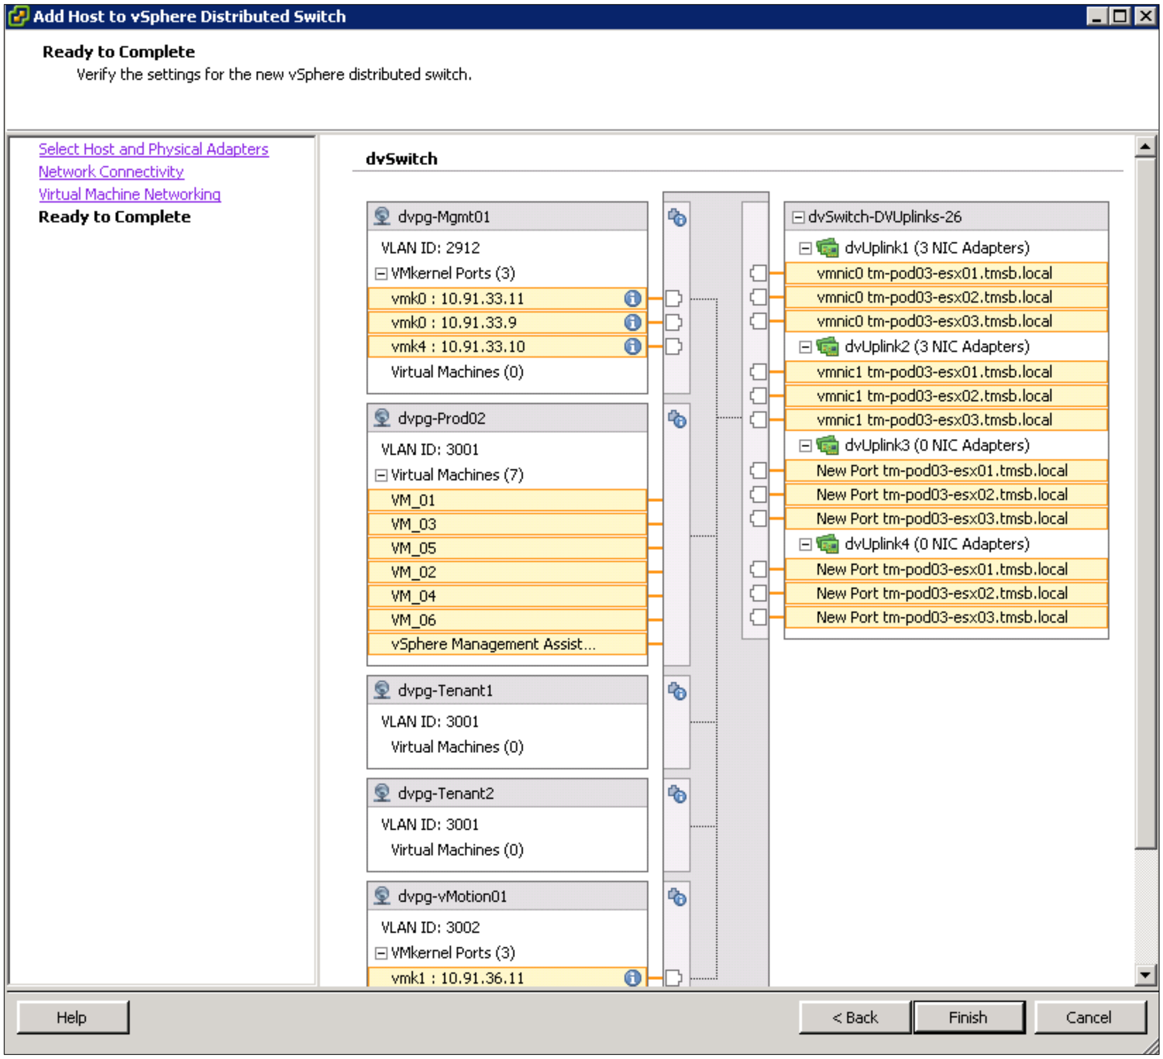

Add Host to vSphere Distributed Switch

Ready to Complete
Verify the settings for the new vSphere distributed switch.

dvSwitch

dvpg-Mgmt01
VLAN ID: 2912
VMkernel Ports (3)
vmk0 : 10.91.33.11
vmk0 : 10.91.33.9
vmk4 : 10.91.33.10
Virtual Machines (0)

dvpg-Prod02
VLAN ID: 3001
Virtual Machines (7)
VM_01
VM_03
VM_05
VM_02
VM_04
VM_06
vSphere Management Assist...

dvpg-Tenant1
VLAN ID: 3001
Virtual Machines (0)

dvpg-Tenant2
VLAN ID: 3001
Virtual Machines (0)

dvpg-vMotion01
VLAN ID: 3002
VMkernel Ports (3)
vmk1 : 10.91.36.11

dvSwitch-DVUplinks-26
dvUplink1 (3 NIC Adapters)
vmnic0 tm-pod03-esx01.tmsb.local
vmnic0 tm-pod03-esx02.tmsb.local
vmnic0 tm-pod03-esx03.tmsb.local
dvUplink2 (3 NIC Adapters)
vmnic1 tm-pod03-esx01.tmsb.local
vmnic1 tm-pod03-esx02.tmsb.local
vmnic1 tm-pod03-esx03.tmsb.local
dvUplink3 (0 NIC Adapters)
New Port tm-pod03-esx01.tmsb.local
New Port tm-pod03-esx02.tmsb.local
New Port tm-pod03-esx03.tmsb.local
dvUplink4 (0 NIC Adapters)
New Port tm-pod03-esx01.tmsb.local
New Port tm-pod03-esx02.tmsb.local
New Port tm-pod03-esx03.tmsb.local

Figure 15. VDS After Migration Is Complete

Some evaluators might not have the same port group configurations on the standard switch as shown in this example environment. However, the migration steps remain the same as in this example.

Deleting a Standard Switch from a Host

Deleting the standard switch from the host is not mandatory, but preferred as a way of cleaning up after the migration to the VDS.

To delete the standard switch, follow these steps:

1. Go to the **Home > Inventory > Hosts and Clusters** view and select the **Configuration** tab, and then select **Networking** from the Hardware box.
2. Select **Remove..** from the panel above the vSwitch0 graphic.

NetFlow

Overview

NetFlow is a networking protocol that collects IP traffic information as records and sends them to a collector such as CA NetQoS for traffic flow analysis. VDS now supports NetFlow version 5 and helps in monitoring virtual infrastructure traffic.

In this exercise, you will configure the NetFlow session that sends the flow information to the collector. To demonstrate how different flows in a virtual infrastructure are collected and sent to the collector, evaluators are expected to create traffic among different virtual machines using traffic generators.

Evaluation Environment for NetFlow

The evaluation environment consists of the following components as shown in Figure 16 and Figure 17:

1. Three virtual machines running Windows OS on Host1
2. Three virtual machines running Windows OS on Host2
3. Each virtual machine has following software tool installed:
    a. JPerf tool (You can download this tool at http://sourceforge.net/projects/iperf/files/)
4. On one of the virtual machines (VM_02) on Host2 ManageEngine NetFlow Analyzer is installed as a collector tool.

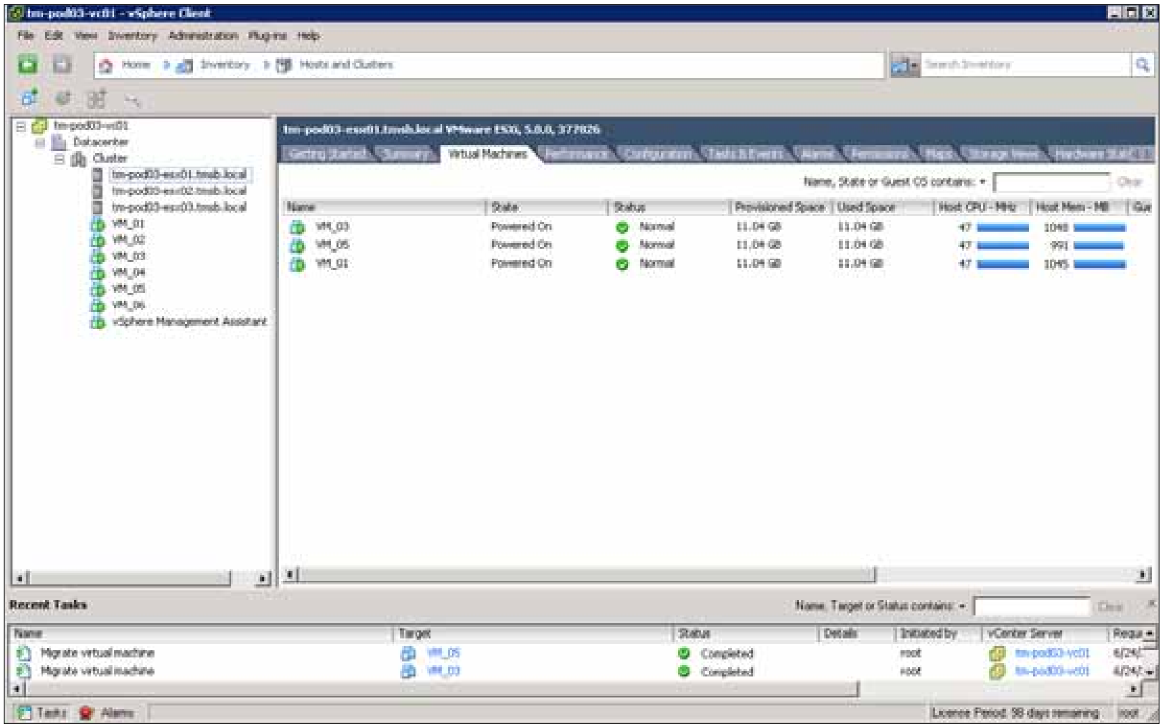
tm-pod03-vc01 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home
Inventory
Hosts and Clusters
Search Inventory
tm-pod03-vc01
Datacenter
Cluster
tm-pod03-esx01.tmsb.local
tm-pod03-esx02.tmsb.local
tm-pod03-esx03.tmsb.local
VM_01
VM_02
VM_03
VM_04
VM_05
VM_06
vSphere Management Assistant
tm-pod03-esx01.tmsb.local VMware ESXi, 5.0.0, 377826
Getting Started
Summary
Virtual Machines
Performance
Configuration
Tasks & Events
Alarms
Permissions
Maps
Storage Views
Hardware Status
Name, State or Guest OS contains:
Clear
Recent Tasks
Name, Target or Status contains:
Clear
Tasks
Alarms
License Period: 98 days remaining
root

Name	State	Status	Provisioned Space	Used Space	Host CPU - MHz	Host Mem - MB	Gue
VM_03	Powered On	Normal	11.04 GB	11.04 GB	47	1048	
VM_05	Powered On	Normal	11.04 GB	11.04 GB	47	991	
VM_01	Powered On	Normal	11.04 GB	11.04 GB	47	1045	

Name	Target	Status	Details	Initiated by	vCenter Server	Requ
Migrate virtual machine	VM_05	Completed		root	tm-pod03-vc01	6/24/
Migrate virtual machine	VM_03	Completed		root	tm-pod03-vc01	6/24/

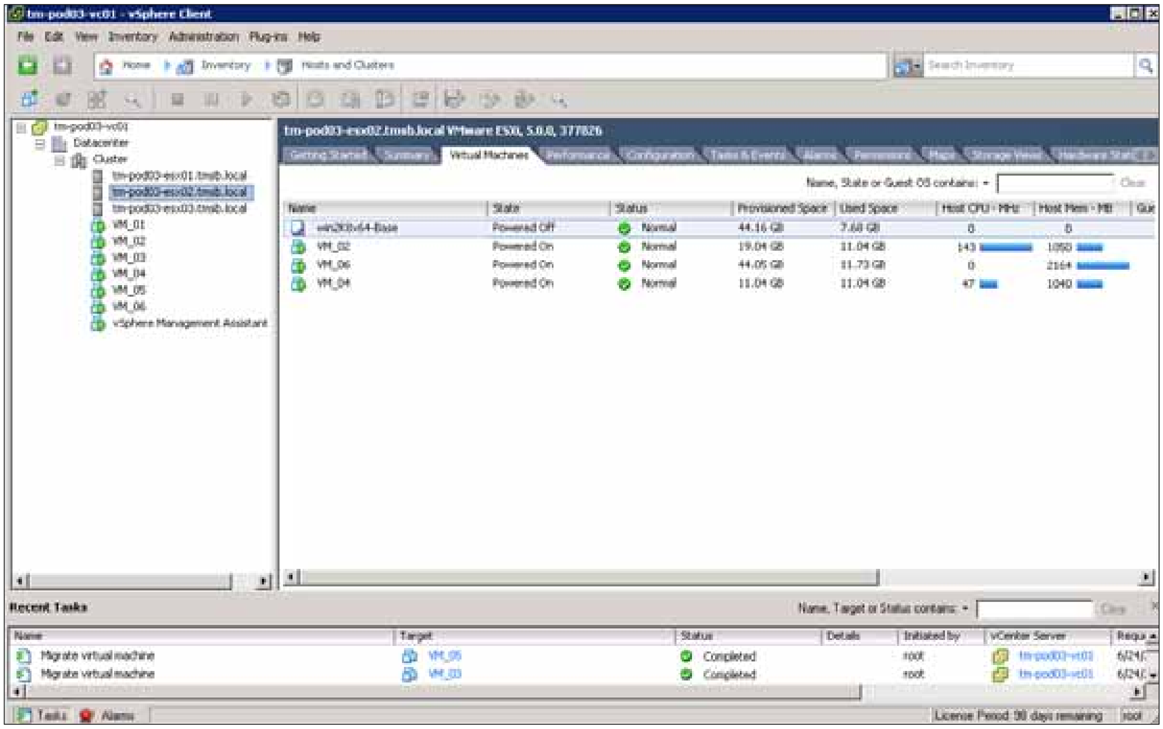
tm-pod03-vc01 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home
Inventory
Hosts and Clusters
Search Inventory
tm-pod03-vc01
Datacenter
Cluster
tm-pod03-esx01.tmsb.local
tm-pod03-esx02.tmsb.local
tm-pod03-esx03.tmsb.local
VM_01
VM_02
VM_03
VM_04
VM_05
VM_06
vSphere Management Assistant
tm-pod03-esx02.tmsb.local VMware ESXi, 5.0.0, 377826
Getting Started
Summary
Virtual Machines
Performance
Configuration
Tasks & Events
Alarms
Permissions
Maps
Storage Views
Hardware Status
Name, State or Guest OS contains:
Clear
Recent Tasks
Name, Target or Status contains:
Tasks
Alarms
License Period: 90 days remaining
root

Name	State	Status	Provisioned Space	Used Space	Host CPU - MHz	Host Mem - MB
win2k8r64-Base	Powered Off	Normal	44.16 GB	7.68 GB	0	0
VM_02	Powered On	Normal	19.04 GB	11.04 GB	143	1050
VM_06	Powered On	Normal	44.05 GB	11.73 GB	0	2164
VM_04	Powered On	Normal	11.04 GB	11.04 GB	47	1040

Name	Target	Status	Details	Initiated by	vCenter Server	Requ
Migrate virtual machine	VM_05	Completed		root	tm-pod03-vc01	6/24/
Migrate virtual machine	VM_03	Completed		root	tm-pod03-vc01	6/24/

Figure 17. Virtual Machines on Host2

Configuring NetFlow

The NetFlow session can be configured at the VDS level. You should collect the following information regarding the collector before the configuration process starts:

1. The collector's IP address: In this evaluation the virtual machine VM_02 has the collector tool installed, and its IP address is 10.91.35.72.
2. NetFlow's Listener Port: As shown in Figure 18, the listener port number is 9996.

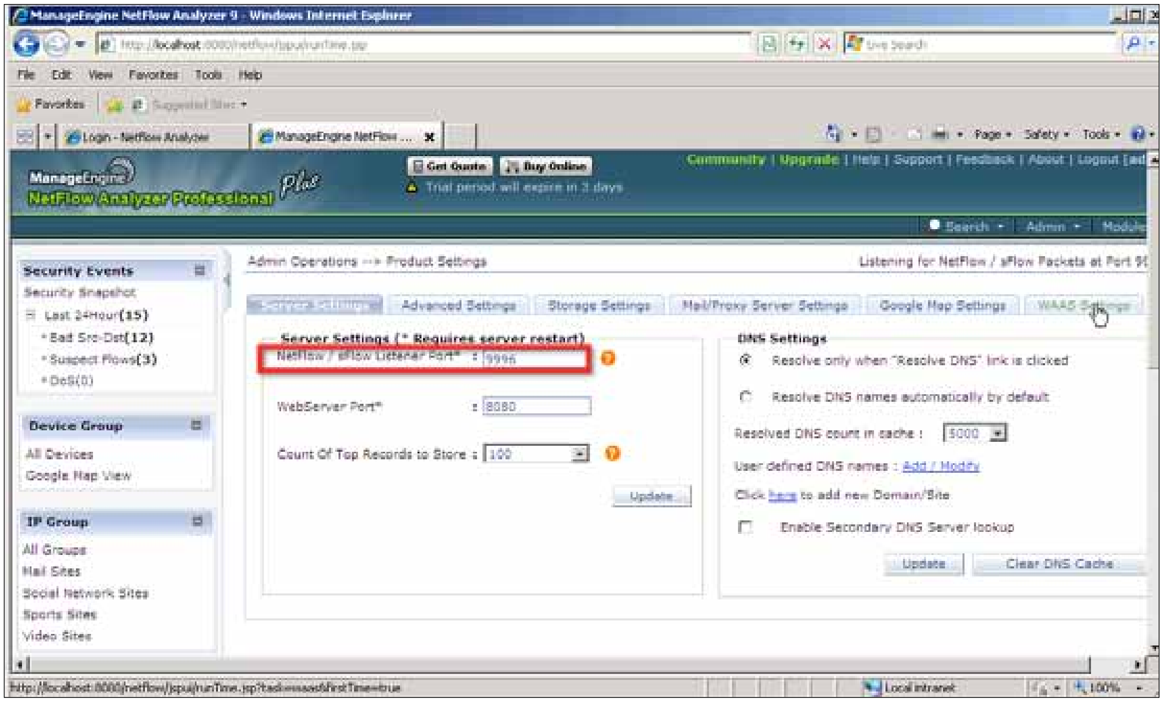

Trial period will expire in 3 days

Admin Operations --> Product Settings

Listening for NetFlow / sFlow Packets at Port 9

Server Settings (* Requires server restart)

NetFlow / sFlow Listener Port* : 9996

WebServer Port* : 8080

Count Of Top Records to Store : 100

DNS Settings

● Resolve only when "Resolve DNS" link is clicked

○ Resolve DNS names automatically by default

Resolved DNS count in cache : 5000

User defined DNS names : Add / Modify

Click here to add new Domain/Site

□ Enable Secondary DNS Server lookup

http://localhost:8080/netflow/jspui/runTime.jsp?tad=waas&firstTime=true

Figure 18. Collector Information

Once you have the required information about the collector, you can now create a NetFlow session on the VDS.

Start the NetFlow configuration process by editing the VDS settings and selecting the NetFlow tab. Enter the following parameters, as shown in Figure 19, to set up the session.

1. The Collector Settings of IP address and Port is configured according to the information collected about the collector tool installed on VM_02.
2. The other NetFlow parameters remain default parameters, but you can modify them. To change the amount of information that is collected, you can change the sampling rate. For example, a sampling rate of 2 indicates that the VDS will collect data from every other packet. You can also modify the Idle flow export timeout values.
3. The VDS IP address configuration is useful when you want to see all flow information in the collector tool as part of one VDS IP address and not as a separate host management network IP address. In this example, because the VDS IP address is not entered, the collector tool will provide flow details under each host's management network IP address.

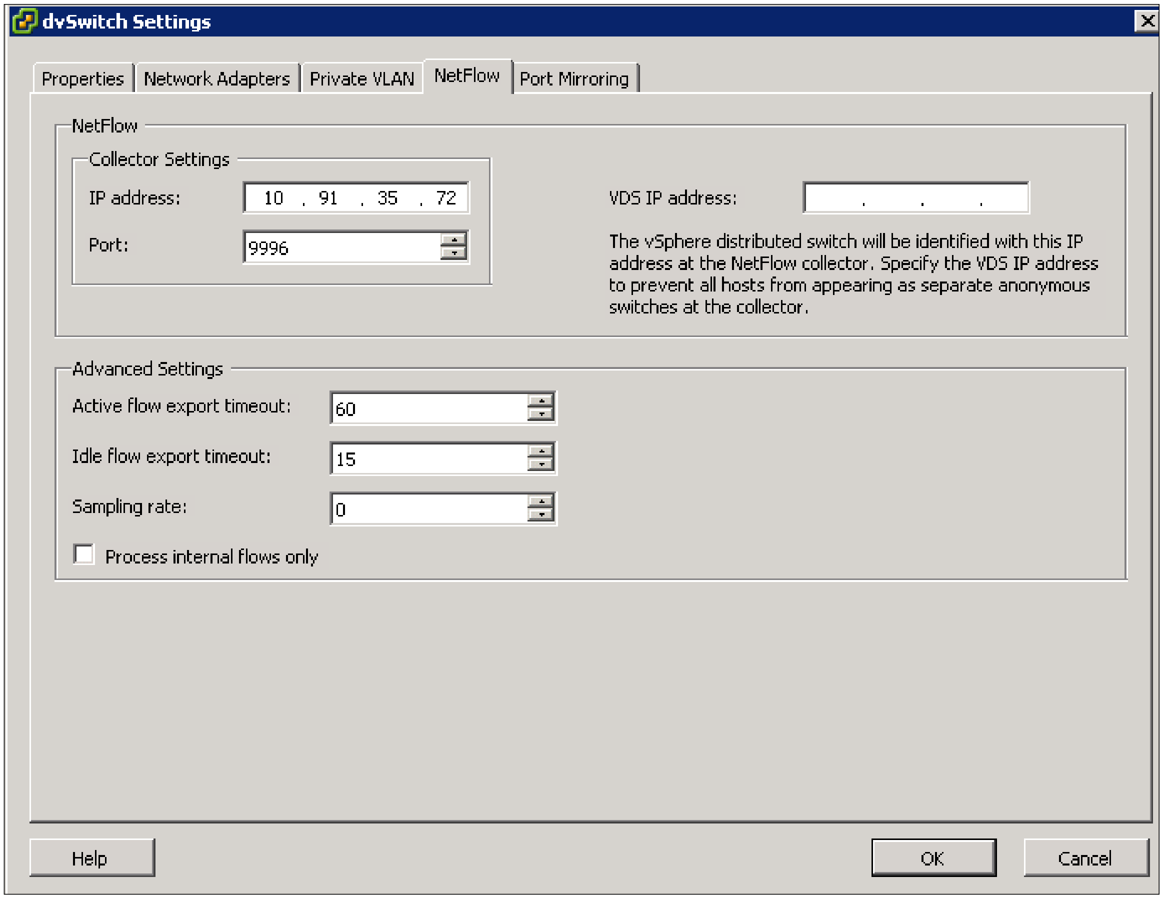
dvSwitch Settings

NetFlow

Collector Settings

IP address: 10 . 91 . 35 . 72

Port: 9996

VDS IP address: . . .

The vSphere distributed switch will be identified with this IP address at the NetFlow collector. Specify the VDS IP address to prevent all hosts from appearing as separate anonymous switches at the collector.

Advanced Settings

Active flow export timeout: 60

Idle flow export timeout: 15

Sampling rate: 0

□ Process internal flows only

Figure 19. NetFlow Configuration

Generating Traffic

After the NetFlow session configuration, you can test the way in which the VDS collects and sends flow data to the collector by generating some traffic using a standard traffic generator. Also, to demonstrate how the new monitoring capability provides the visibility into the virtual machine to virtual machine traffic, you can create traffic flows between two virtual machines on Host1.

The following are the different flows that are created using the JPerf tool:

1. The VM_01 to VM_03 TCP session running on Host1
2. The VM_05 on Host1 to VM_02 on Host2 TCP session

More details on how to configure the JPerf tool is provided below.

JPerf is a tool that helps in measuring network bandwidth. The tool requires a client- and server-side setup. Figure 20 shows the server side configuration running on VM_01. The server listens on port 5001 for any traffic from the client. Under **Choose iPerf Mode,** when you select “Server,” the tool automatically fills the **Iperf command** field with appropriate command.

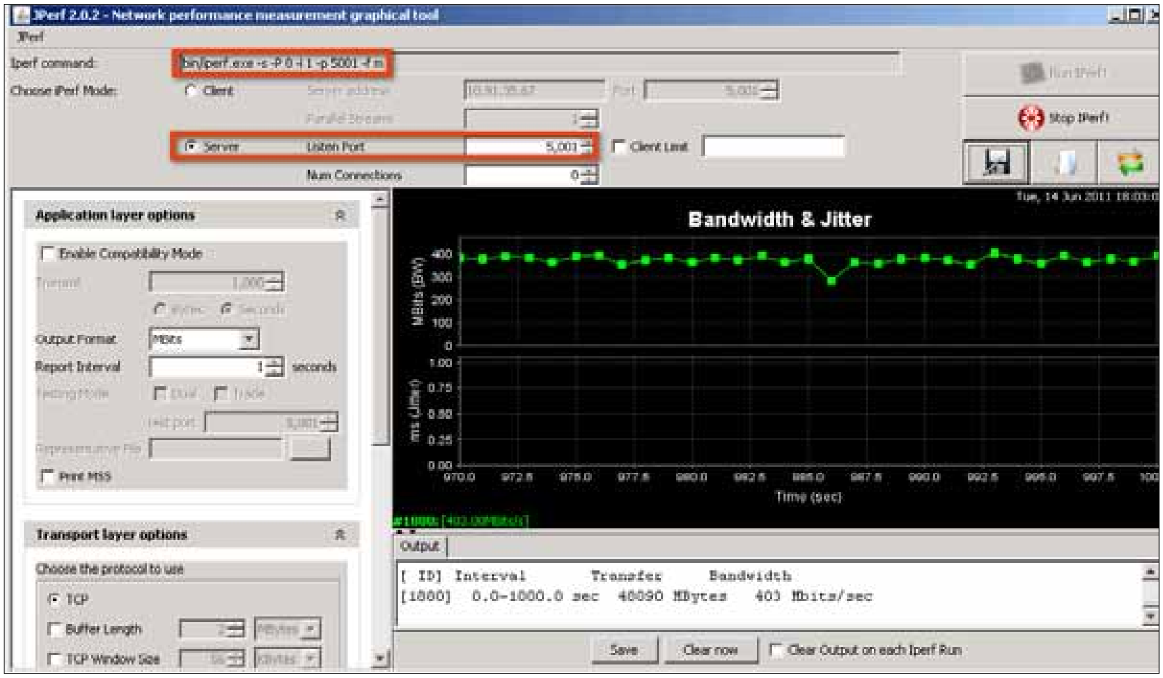

JPerf 2.0.2 - Network performance measurement graphical tool

Iperf command: bin/iperf.exe -s -P 0 -i 1 -p 5001 -f m

Choose iPerf Mode: ○ Client ● Server Listen Port 5,001 Num Connections 0 □ Client Limit

Application layer options: □ Enable Compatibility Mode; Output Format MBits; Report Interval 1 seconds; □ Print MSS

Transport layer options: Choose the protocol to use ● TCP; □ Buffer Length; □ TCP Window Size

```
[ ID] Interval       Transfer     Bandwidth
[1800]  0.0-1000.0 sec  48090 MBytes   403 Mbits/sec
```

Figure 20. The JPerf Server Settings

Figure 21 shows the client-side configuration running on VM_03. The JPerf client running on VM_03 should send traffic to server IP address 10.91.35.73 on port 5,001. To configure the client side, under **Choose iPerf Mode,** select "Client" and enter the Server address and Port fields. To control the amount of traffic that the client will send, you can configure the **Application layer options.**

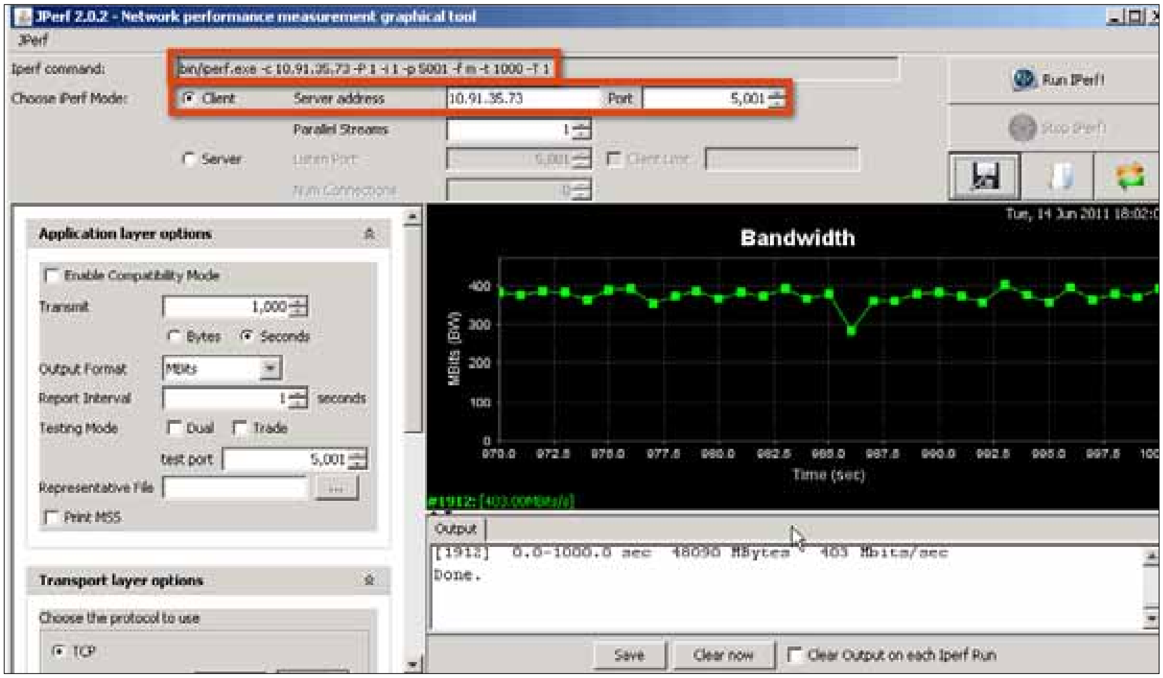

JPerf 2.0.2 - Network performance measurement graphical tool

Iperf command: bin/iperf.exe -c 10.91.35.73 -P 1 -i 1 -p 5001 -f m -t 1000 -T 1

Choose iPerf Mode: ● Client Server address 10.91.35.73 Port 5,001; Parallel Streams 1; ○ Server

Application layer options: □ Enable Compatibility Mode; Transmit 1,000 ○ Bytes ● Seconds; Output Format MBits; Report Interval 1 seconds; Testing Mode □ Dual □ Trade; test port 5,001; Representative File; □ Print MSS

Transport layer options: Choose the protocol to use ● TCP

```
[1912]  0.0-1000.0 sec  48090 MBytes   403 Mbits/sec
Done.
```

Figure 21. The JPerf Client Settings

You can generate additional traffic flows between other virtual machines by launching the client and server side of the JPerf tool.

Checking Collector Results

After the traffic generation, you can check the collector tool interface for the flow data that was processed and sent by the VDS. In this example environment, two TCP sessions are created. One TCP session runs between two virtual machines on the same host and other between two virtual machines on different hosts. A TCP session consists of two flows. Traffic flowing from the client to the server is one flow and traffic flowing from the server to the client is another flow.

VDS captures the flows that are flowing through the virtual switch on a host, and then sends the flow data over a UDP session to the collector. Along with the flow data, the VDS sends the VDS IP address configured in the NetFlow Session. If the VDS IP address is not configured, as is the case in this example environment, the VDS sends the Host Management IP address along with flow data to identify the flows monitored on a particular host. The following are the Management IP addresses used for the two hosts in this environment:

1. Host1 – tm-pod03-esx01.tmsb.local: 10.91.33.9
2. Host2 – tm-pod03-esx02.tmsb.local: 10.91.33.10

Figure 22 shows the collector screen shot that provides the information on the data collected. The highlighted (red rectangle) application is NetFlow and shows the Source IP address as the Management IP addresses of the two hosts.

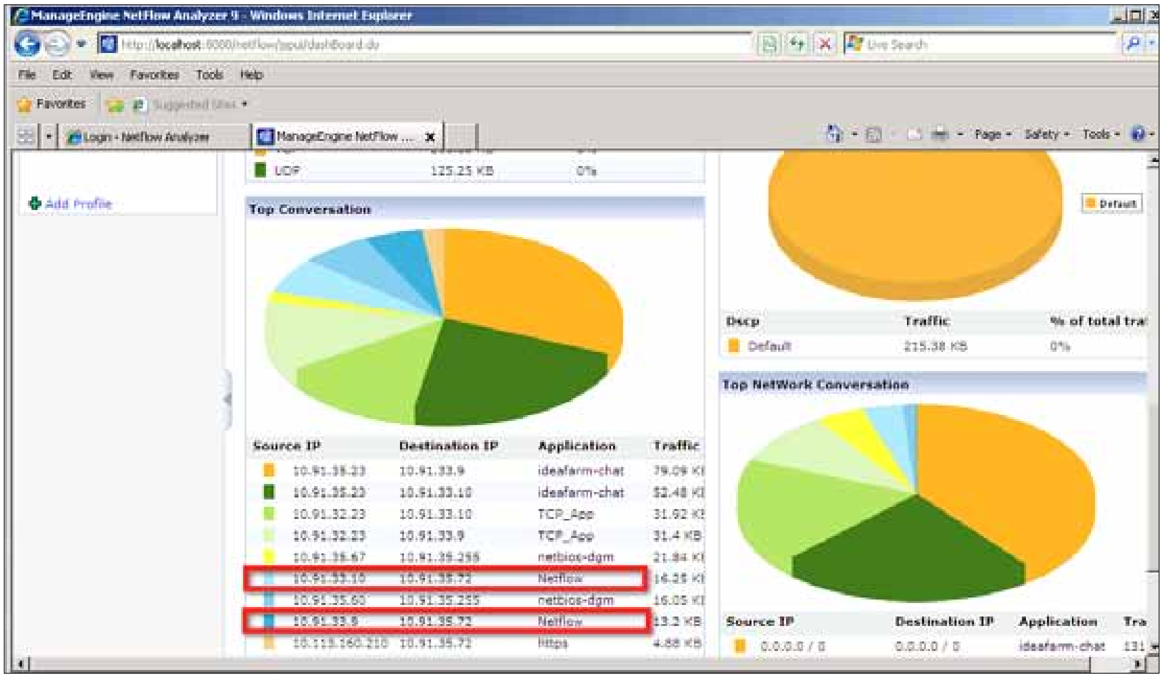

🟩 UDP	125.25 KB	0%

Top Conversation

Source IP	Destination IP	Application	Traffic
10.91.35.23	10.91.33.9	ideafarm-chat	79.09 K
10.91.35.23	10.91.33.10	ideafarm-chat	52.48 K
10.91.32.23	10.91.33.10	TCP_App	31.92 K
10.91.32.23	10.91.33.9	TCP_App	31.4 KB
10.91.35.67	10.91.35.255	netbios-dgm	21.84 K
10.91.33.10	10.91.35.72	Netflow	16.25 K
10.91.35.60	10.91.35.255	netbios-dgm	16.05 K
10.91.33.9	10.91.35.72	Netflow	13.2 KB
10.113.160.210	10.91.35.72	https	4.88 KB

Dscp	Traffic	% of total tra
Default	215.38 KB	0%

Top NetWork Conversation

Source IP	Destination IP	Application	Tra
0.0.0.0 / 0	0.0.0.0 / 0	ideafarm-chat	131

Figure 22. Collector Screen Shot

The collector tool also provides historical data as shown in the screen shot in Figure 23. It provides Top Application traffic pie chart information as well as statistics on the Top Devices sending traffic. You can find many more stats and reports in this Collector tool that will help you measure the performance of the application traffic as well as detect any security breaches.

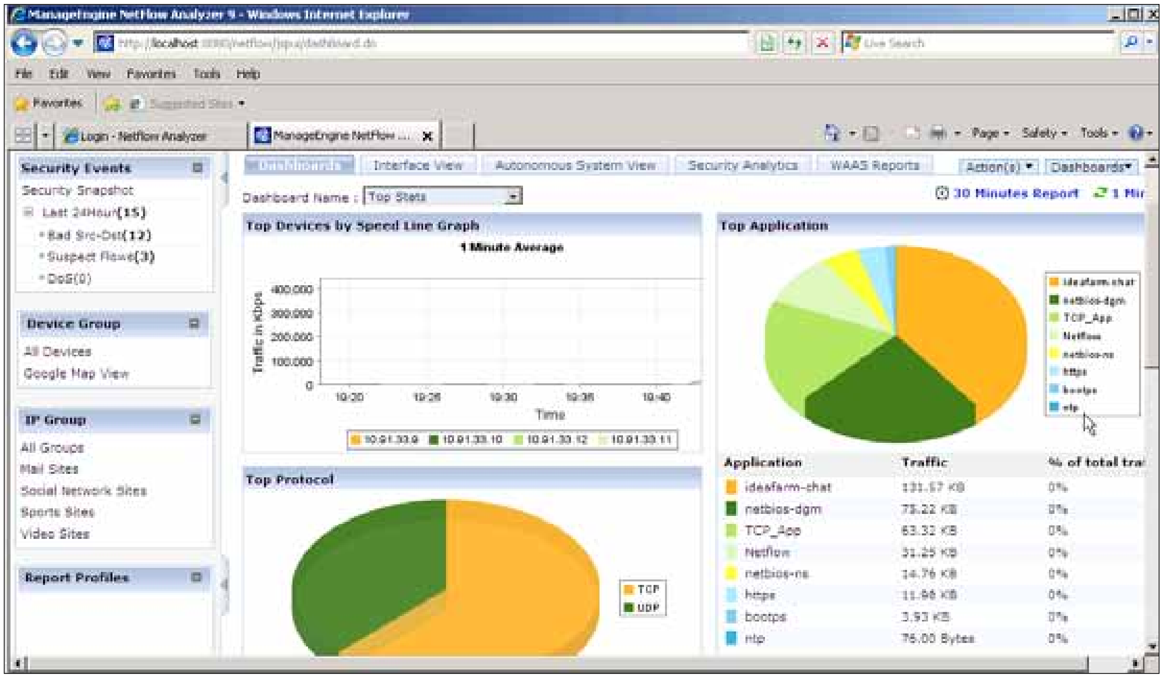

Application	Traffic	% of total tra
🟧 ideafarm-chat	131.57 KB	0%
🟩 netbios-dgm	75.22 KB	0%
🟩 TCP_App	63.32 KB	0%
🟩 Netflow	31.25 KB	0%
🟨 netbios-ns	14.76 KB	0%
🟦 https	11.96 KB	0%
🟦 bootps	3.93 KB	0%
🟦 ntp	76.00 Bytes	0%

Figure 23. Collector Screen Shot

Port Mirroring

Overview

Port mirroring is the capability on a network switch to send a copy of network packets seen on a switch port to a network monitoring device connected to another switch port. Port mirroring is also referred to as SPAN on Cisco switches. In vSphere 5.0, a distributed switch provides a port mirroring capability similar to that available on a physical network switch.

In this exercise, you will configure the port mirroring session such that it will provide complete visibility into the traffic flowing to and from a virtual machine. To demonstrate how network administrators can troubleshoot the virtual infrastructure traffic, evaluators are encouraged to create different internal traffic patterns and to use different destinations (virtual machine or uplink) to send mirror traffic. In this example environment, the virtual machine is configured as the mirror destination. On this destination the virtual machine Wireshark tool is installed to capture and analyze the mirror traffic.

Evaluation Environment for a Port Mirroring Session

The evaluation environment of Port Mirroring is similar to the one used during the evaluation of NetFlow feature. Figure 16 and Figure 17 show the different components in this environment:

1. Three virtual machines running Windows OS on Host1
2. Three virtual machines running Windows OS on Host2
3. Each virtual machine with the following software tool installed:
    a. JPerf tool

4. On one of the virtual machines on Host1, Wireshark is installed. Wireshark is a network protocol analyzer tool that will allow you to monitor the mirror traffic.

Configuring a Port Mirroring Session

A port mirroring session can be configured at the VDS level and needs the following parameters to perform the setup:

1. Source to be monitored: virtual machine dvPort number
2. Which traffic: ingress only, egress only, or both ingress and egress
3. Destination where the packet will be mirrored to: virtual machine, vmkNIC, or uplink dvPort number

Once you decide which virtual machine traffic you want to monitor, you can get the corresponding dvPort number using the following steps:

1. Switch to the **Home > Inventory > Networking** view.
2. Select dvSwitch and choose the **Ports** tab on the right panel. Scroll down to see the virtual machines and the associated port ID. Figure 24 shows virtual machines and port ID mapping.

In the example environment, the ingress traffic of virtual machine VM_01 is monitored and virtual machine VM_05 is the destination port to which packets will be mirrored. Both these virtual machine are on Host1.

The terms **ingress traffic** and **egress traffic** are with respect to the VDS. So when you want to monitor the traffic that is going out of the virtual machine towards the VDS, it is called ingress traffic. The traffic seeks **ingress** to the VDS and hence the source is called ingress. If you want to monitor traffic that is received by a virtual machine, then configure the port mirroring session with the source as **egress.**

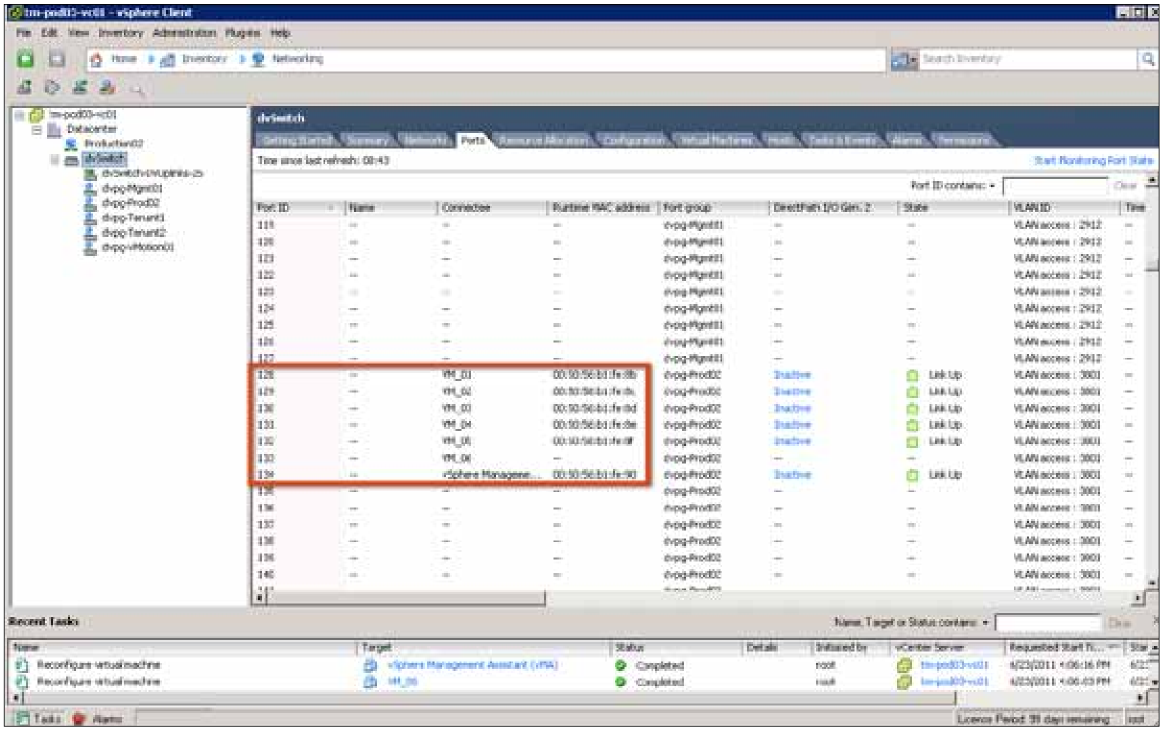

Time since last refresh: 00:43

Port ID	Name	Connectee	Runtime MAC address	Port group	DirectPath I/O Gen. 2	State	VLAN ID	Time
119	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
120	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
121	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
122	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
123	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
124	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
125	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
126	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
127	—	—	—	dvpg-Mgmt01	—	—	VLAN access : 2912	—
128	—	VM_01	00:50:56:b1:fe:8b	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
129	—	VM_02	00:50:56:b1:fe:8c	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
130	—	VM_03	00:50:56:b1:fe:8d	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
131	—	VM_04	00:50:56:b1:fe:8e	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
132	—	VM_05	00:50:56:b1:fe:8f	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
133	—	VM_06	—	dvpg-Prod02	—	—	VLAN access : 3001	—
134	—	vSphere Manageme...	00:50:56:b1:fe:90	dvpg-Prod02	Inactive	Link Up	VLAN access : 3001	—
135	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—
136	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—
137	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—
138	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—
139	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—
140	—	—	—	dvpg-Prod02	—	—	VLAN access : 3001	—

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Sta
Reconfigure virtual machine	vSphere Management Assistant (vMA)	Completed		root	tm-pod03-vc01	6/23/2011 4:06:16 PM	6/2
Reconfigure virtual machine	VM_05	Completed		root	tm-pod03-vc01	6/23/2011 4:06:03 PM	6/2

Figure 24. Port ID and Virtual Machine Mapping

3. Once you have identified the port IDs, it is time to configure the port mirroring session by selecting **dvSwitch** in **Networking** view. Right-click on dvSwitch and select **Edit Settings.**
4. Select the **Port Mirroring** tab and click **Add** as shown in Figure 25.

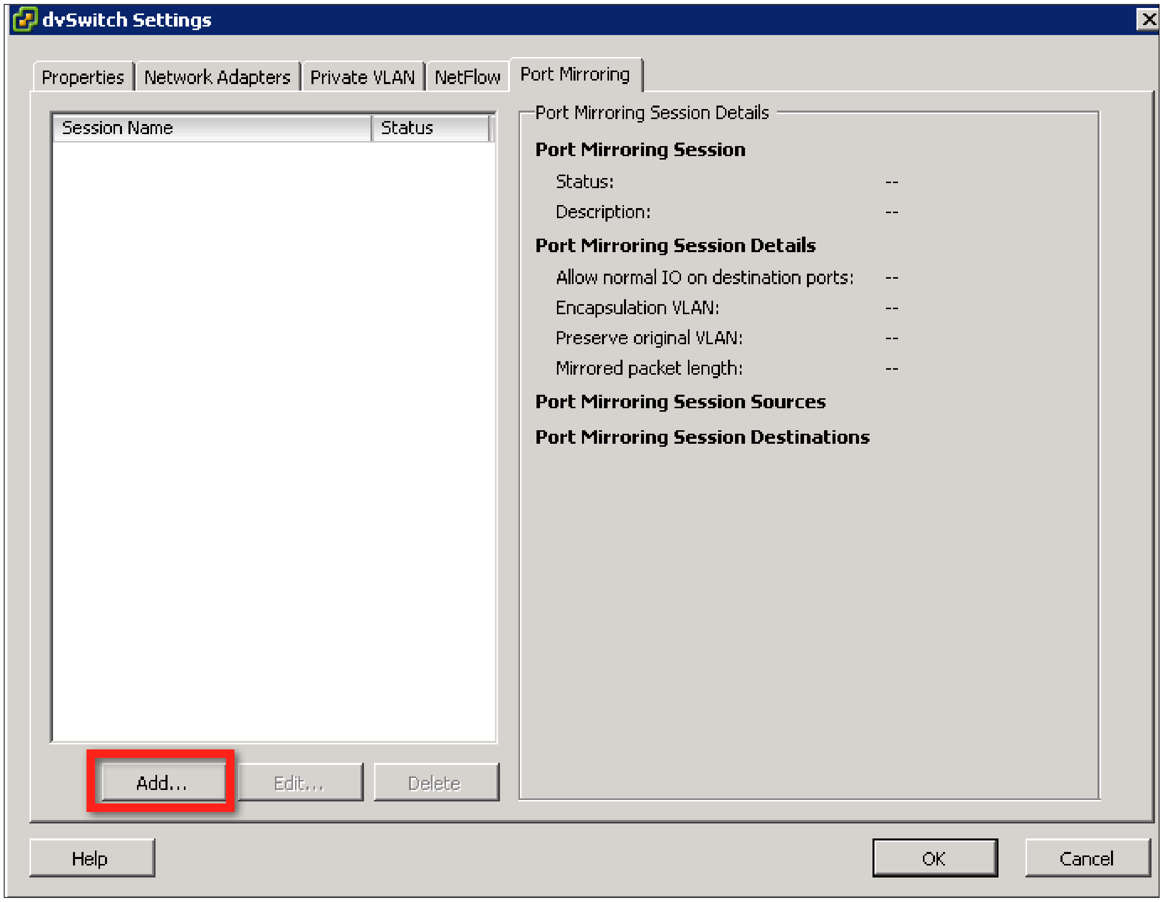

Port Mirroring Session Details

Port Mirroring Session

Status: --

Description: --

Port Mirroring Session Details

Allow normal IO on destination ports: --

Encapsulation VLAN: --

Preserve original VLAN: --

Mirrored packet length: --

Port Mirroring Session Sources

Port Mirroring Session Destinations

Figure 25. Adding a Port Mirroring Session

5. Choose a name for the session and click **Next.**

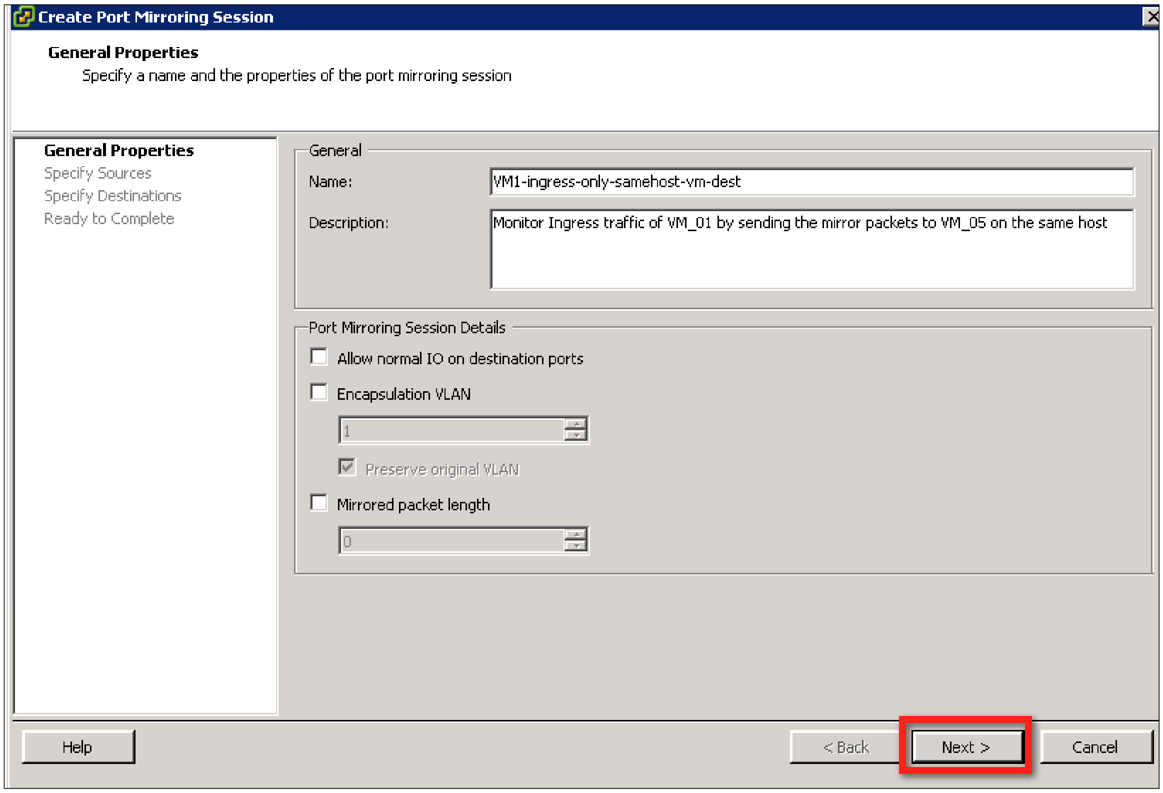

Create Port Mirroring Session

General Properties

Specify a name and the properties of the port mirroring session

General

Name: VM1-ingress-only-samehost-vm-dest

Description: Monitor Ingress traffic of VM_01 by sending the mirror packets to VM_05 on the same host

Port Mirroring Session Details

□ Allow normal IO on destination ports

□ Encapsulation VLAN

1

☑ Preserve original VLAN

□ Mirrored packet length

0

Figure 26. Port Mirroring Session Configuration (Continued)

6. Because you are monitoring ingress traffic on VM_01, select Ingress from the **Traffic direction** drop-down menu.

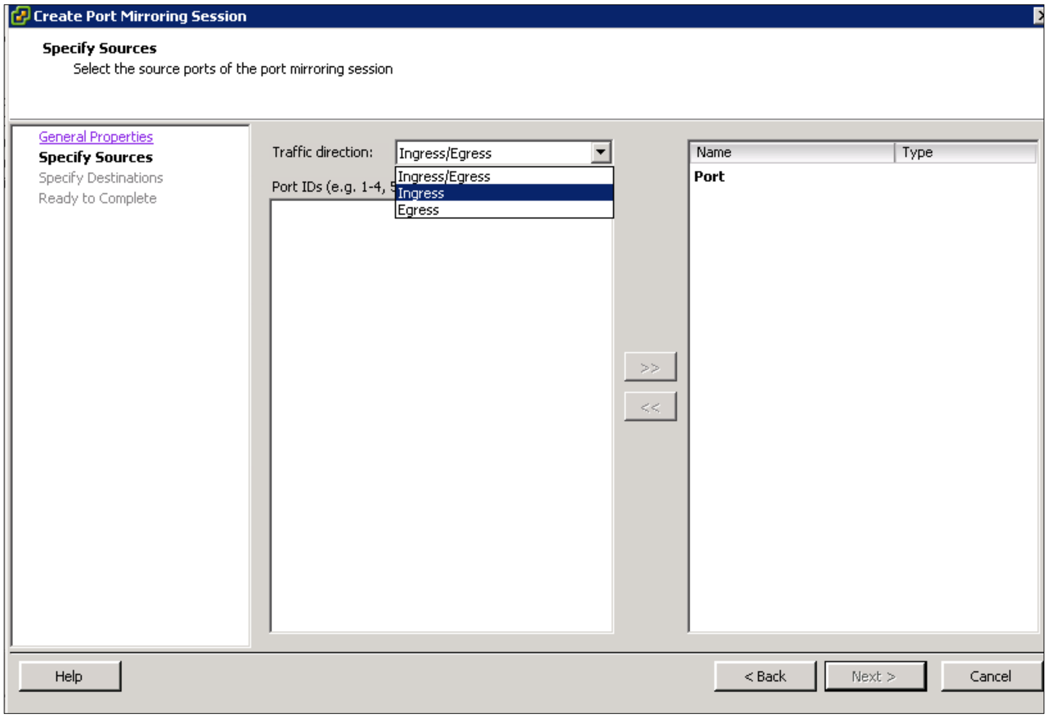

Create Port Mirroring Session

Specify Sources

Select the source ports of the port mirroring session

Traffic direction: Ingress/Egress

Figure 27. Port Mirroring Session Configuration (Continued)

7. Specify the source by providing the port ID of VM_01 in the Port IDs field and then move it to the right field, under Port. Figure 29 shows the screen shot after the source port is entered. Click **Next.**

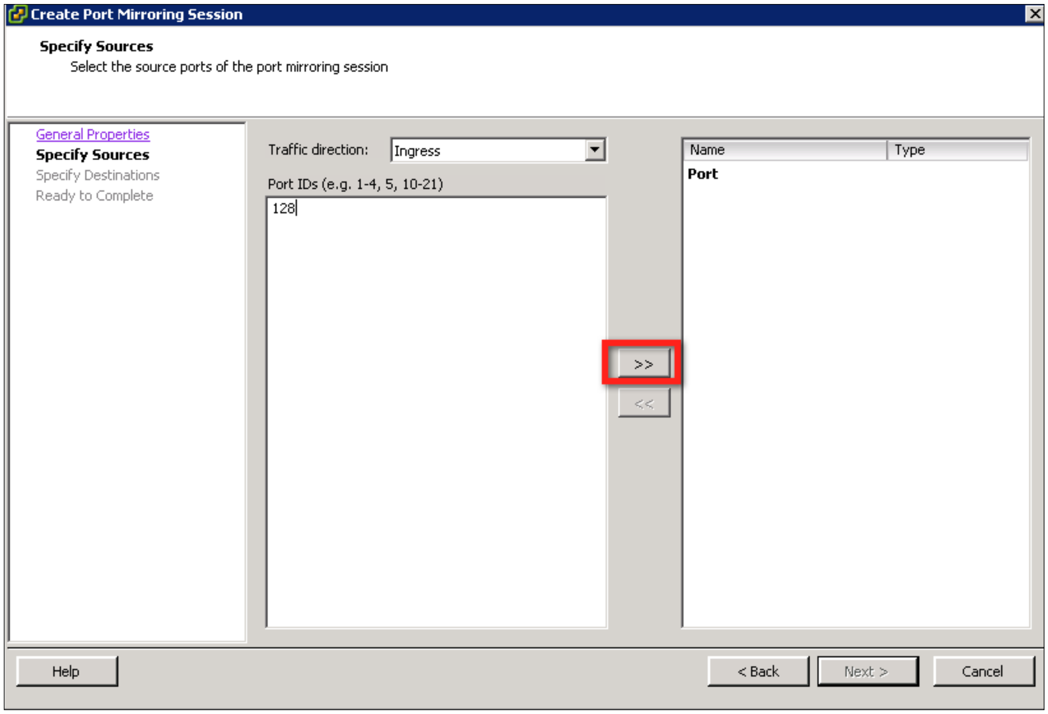

Create Port Mirroring Session

Specify Sources

Select the source ports of the port mirroring session

General Properties
Specify Sources
Specify Destinations
Ready to Complete

Traffic direction: Ingress

Port IDs (e.g. 1-4, 5, 10-21)

128

Name	Type
Port	

Figure 28. Port Mirroring Session Configuration (Continued)

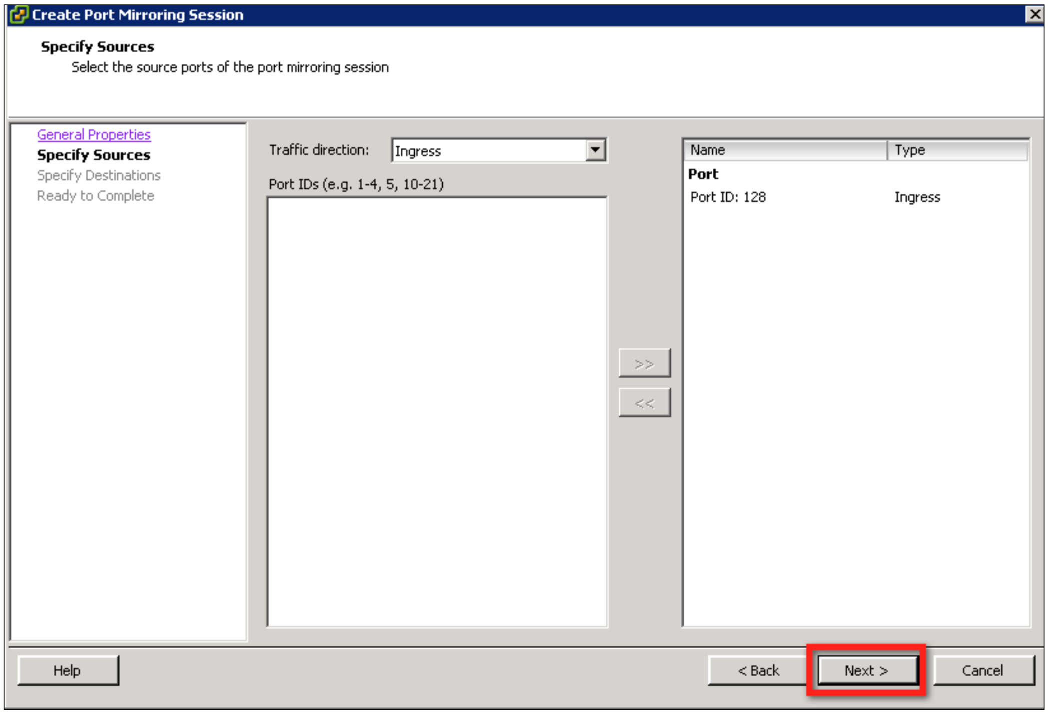

Create Port Mirroring Session

Specify Sources

Select the source ports of the port mirroring session

General Properties
Specify Sources
Specify Destinations
Ready to Complete

Traffic direction: Ingress

Port IDs (e.g. 1-4, 5, 10-21)

Name	Type
Port	
Port ID: 128	Ingress

Figure 29. Port Mirroring Session Configuration (Continued)

8. Specify destination by selecting either Port or Uplink from the Destination type drop-down menu. In this example, you are sending the mirror traffic to virtual machine VM_05, which is running on the same Host1 where virtual machine VM_01 is running. Select **Port** from the drop-down menu. You also have an option to mirror the traffic to an uplink port by selecting Uplink under Destination type.

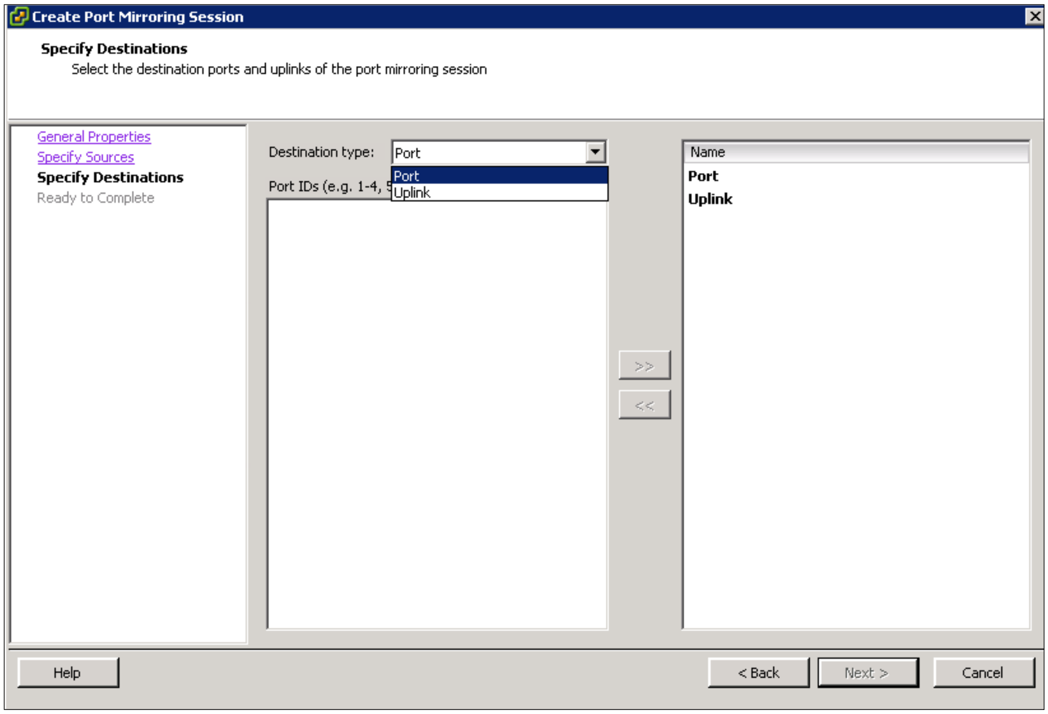

Create Port Mirroring Session

Specify Destinations

Select the destination ports and uplinks of the port mirroring session

Destination type: Port

Port IDs (e.g. 1-4, 5

Name
Port
Uplink

Figure 30. Port Mirroring Session Configuration (Continued)

9. Enter the port ID number of VM_05 in the Port IDs field and move it to the right under the Port field. Figure 32 shows the screen shot after the VM_05 port is selected as the destination.

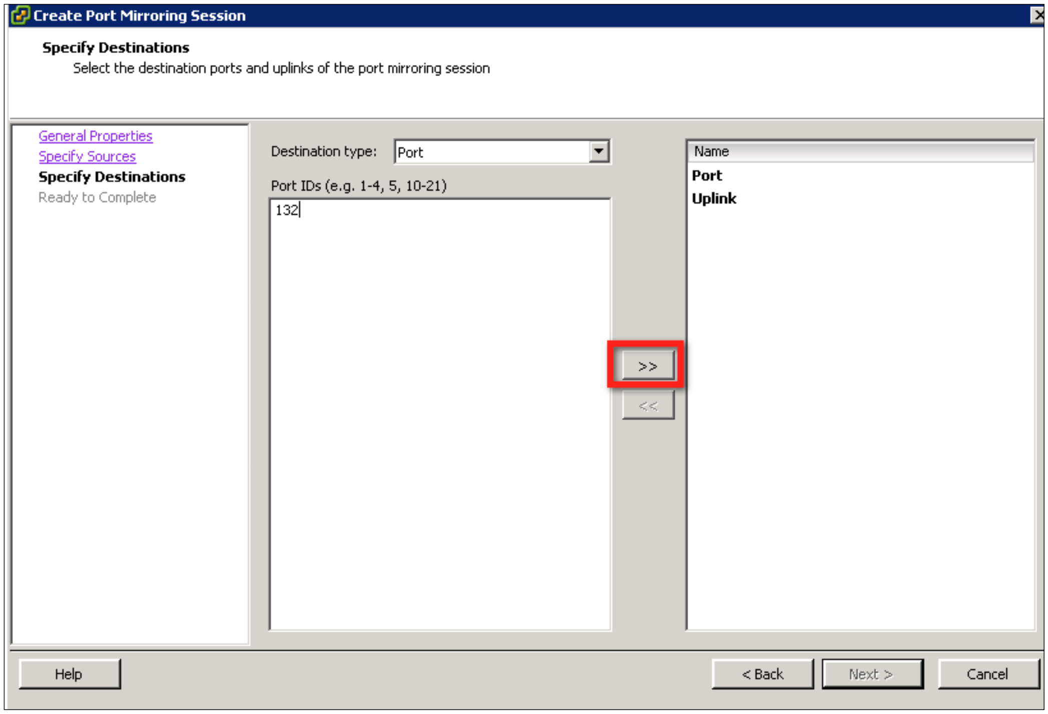

Specify Destinations

Select the destination ports and uplinks of the port mirroring session

Destination type: Port

Port IDs (e.g. 1-4, 5, 10-21)

132

Name

Port

Uplink

Figure 31. Port Mirroring Session Configuration (Continued)

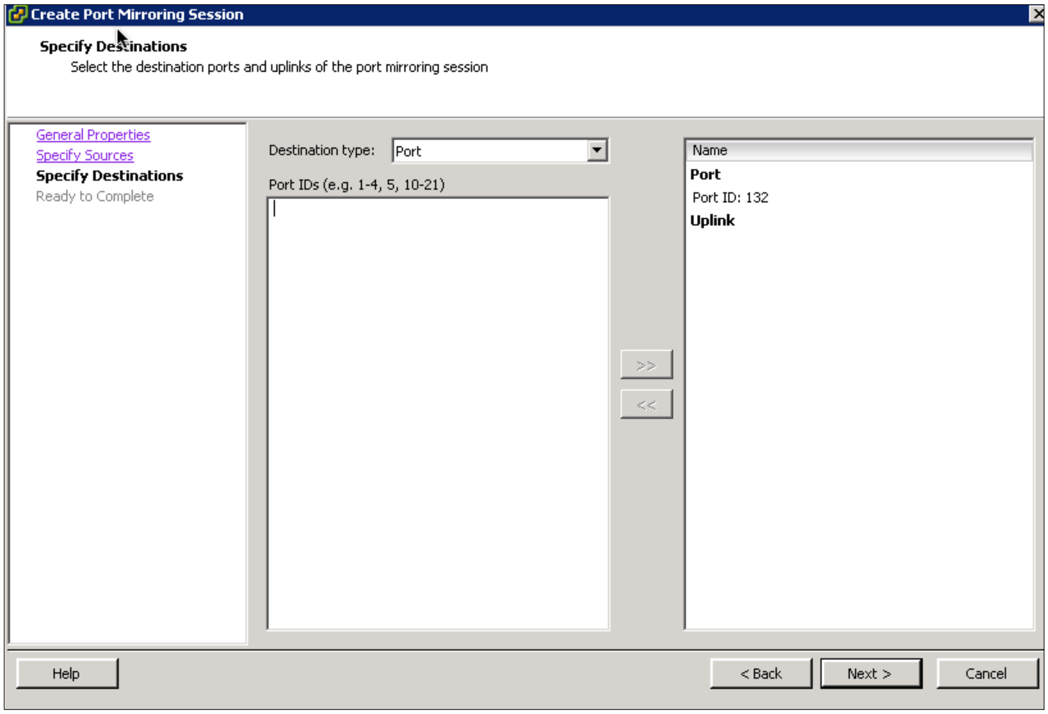

Specify Destinations

Select the destination ports and uplinks of the port mirroring session

Destination type: Port

Port IDs (e.g. 1-4, 5, 10-21)

Name

Port

Port ID: 132

Uplink

Figure 32. Port Mirroring Session Configuration (Continued)

10. This completes the creation of the port mirroring session. As shown in Figure 33, the status of the session is **Disabled**.

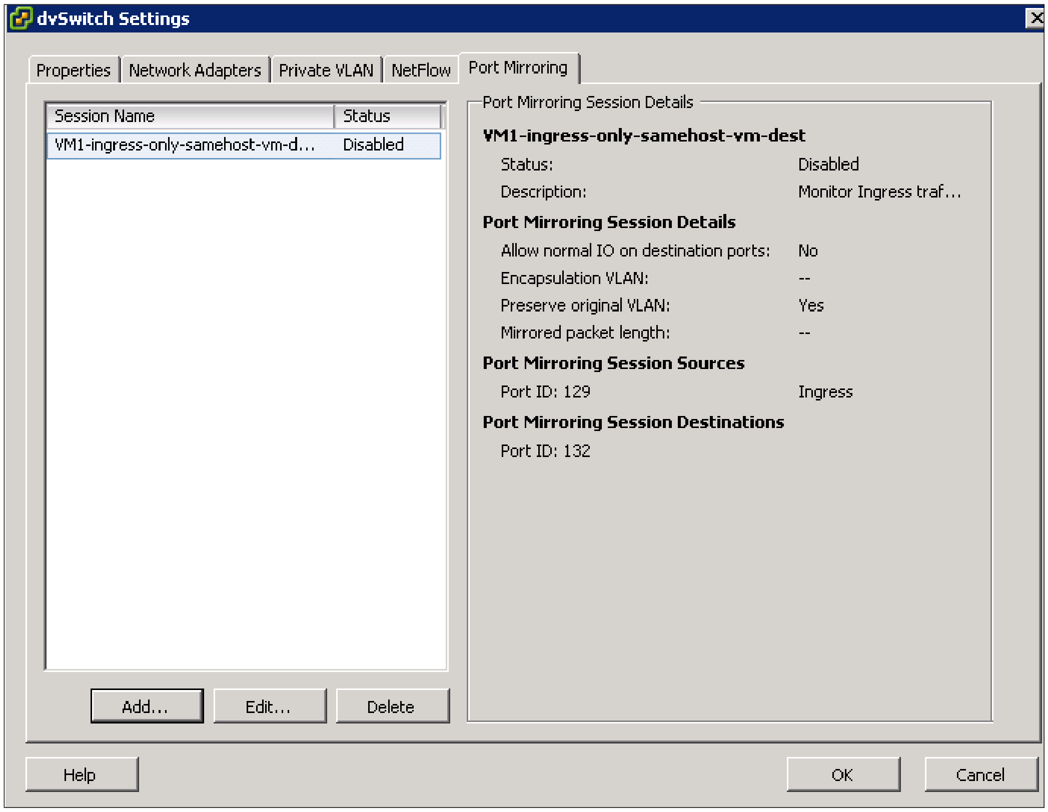

Session Name	Status
VM1-ingress-only-samehost-vm-d...	Disabled

Port Mirroring Session Details

VM1-ingress-only-samehost-vm-dest

Status: Disabled

Description: Monitor Ingress traf...

Port Mirroring Session Details

Allow normal IO on destination ports: No

Encapsulation VLAN: --

Preserve original VLAN: Yes

Mirrored packet length: --

Port Mirroring Session Sources

Port ID: 129 Ingress

Port Mirroring Session Destinations

Port ID: 132

Figure 33. Port Mirroring Session Configuration Is Complete

11. To enable the port mirroring session, click **Edit,** as shown in Figure 33. This will pop up the panel shown in Figure 34. Select **Enabled** as the status.

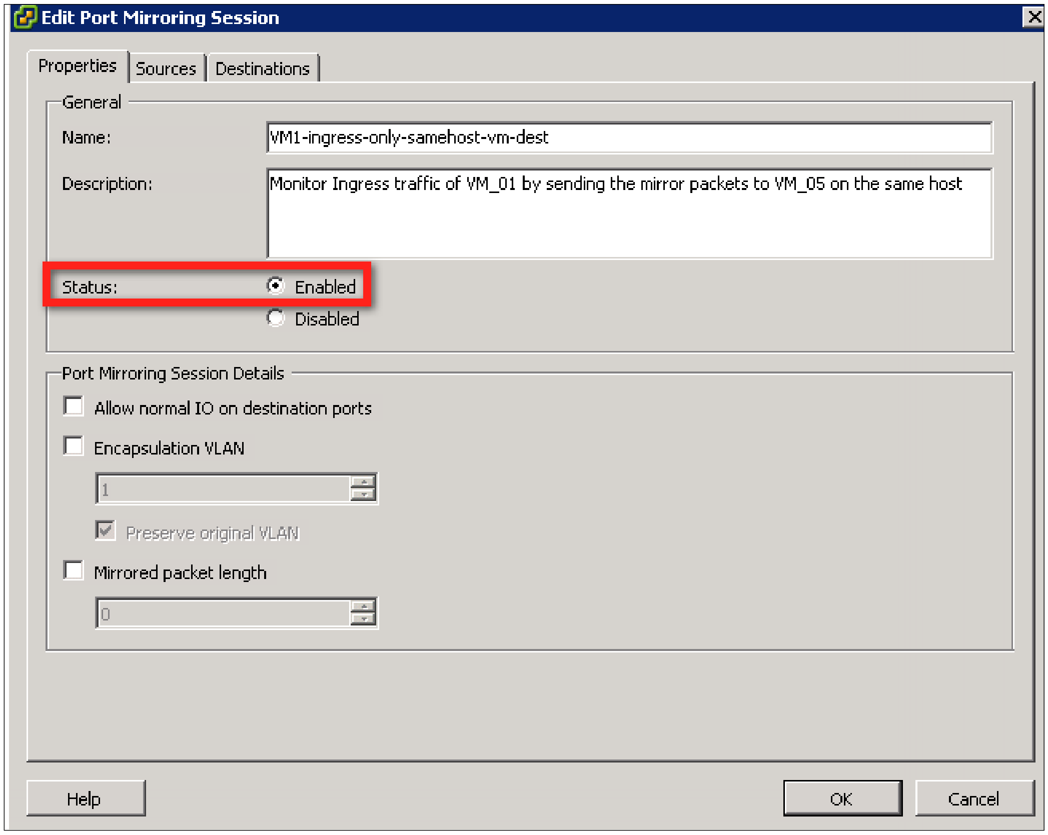

General

Name: VM1-ingress-only-samehost-vm-dest

Description: Monitor Ingress traffic of VM_01 by sending the mirror packets to VM_05 on the same host

Status: ● Enabled ○ Disabled

Port Mirroring Session Details

□ Allow normal IO on destination ports

□ Encapsulation VLAN

1

☑ Preserve original VLAN

□ Mirrored packet length

0

Figure 34. Enable Port Mirroring Session

Generating Traffic

After the port mirroring session configuration, you can test the way in which the VDS mirrors the packets to the destination port by generating some traffic using a standard traffic generator. Also, to demonstrate how the port mirroring capability provides visibility to all packets that a virtual machine sends or receives, you can create traffic between two virtual machines as in the following example:

1. VM_01 (10.91.35.60) to VM_03 (10.91.35.73) TCP session running on Host1

You can configure the JPerf server on VM_03 and the JPerf client on VM_01. For more details on how to configure the JPerf client and server, you can refer to the steps described in the NetFlow evaluation section, along with Figure 20 and Figure 21.

Checking Mirrored Traffic Using Wireshark

To check the mirrored traffic on the destination VM_05, you must install the Wireshark tool. This tool helps in analyzing network traffic. You can download this tool from www.wireshark.org. After installing the Wireshark tool, you can configure it to monitor the TCP traffic that you have mirrored from VM_01.

1. Click **How to Capture** to configure the filter for TCP traffic.

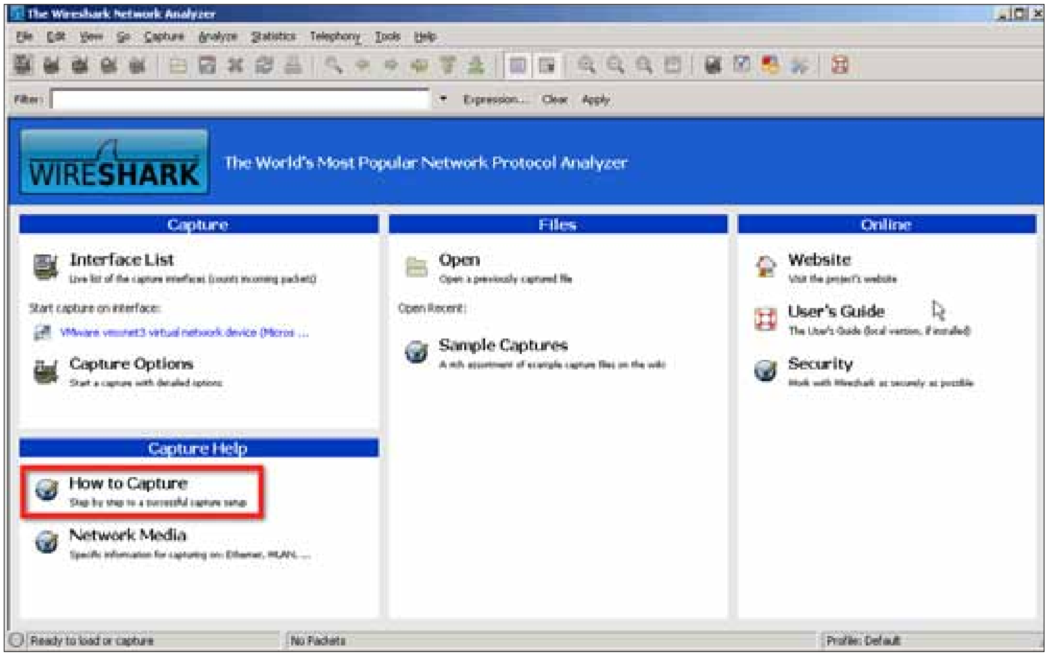
The Wireshark Network Analyzer
File Edit View Go Capture Analyze Statistics Telephony Tools Help
Filter: Expression... Clear Apply
WIRESHARK
The World's Most Popular Network Protocol Analyzer
Capture
Interface List
Live list of the capture interfaces (counts incoming packets)
Start capture on interface:
VMware vmxnet3 virtual network device (Micros ...
Capture Options
Start a capture with detailed options
Capture Help
How to Capture
Step by step to a successful capture setup
Network Media
Specific information for capturing on: Ethernet, WLAN, ...
Files
Open
Open a previously captured file
Open Recent:
Sample Captures
A rich assortment of example capture files on the wiki
Online
Website
Visit the project's website
User's Guide
The User's Guide (local version, if installed)
Security
Work with Wireshark as securely as possible
Ready to load or capture
No Packets
Profile: Default

Figure 35. Wireshark

2. After Wireshark's Capture Options panel pops up, you can click the **Capture Filter** tab to choose the traffic that you want to monitor.

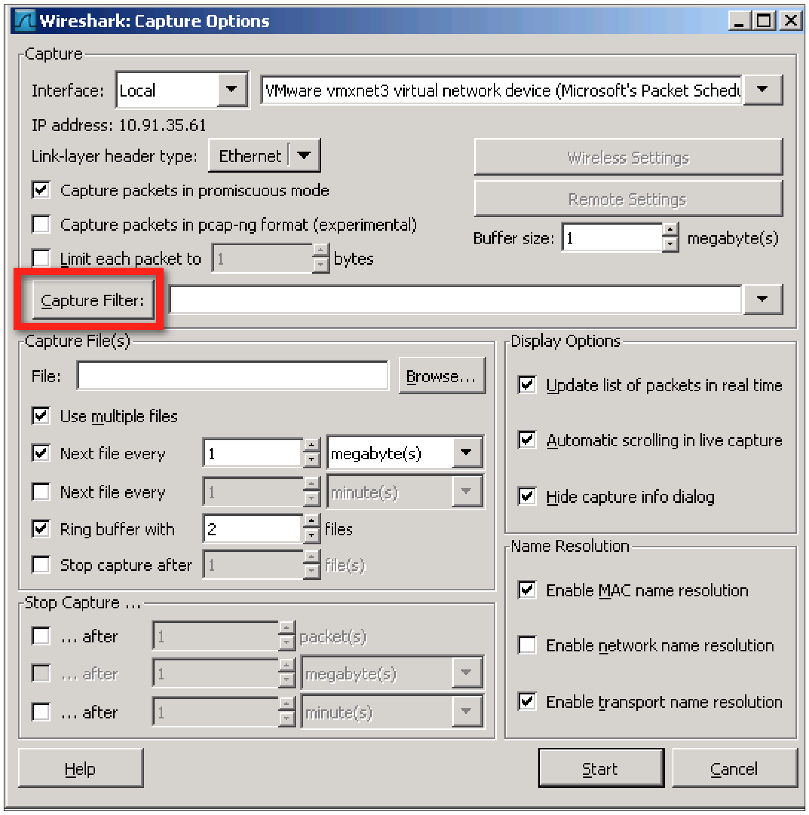

Capture

Interface: Local VMware vmxnet3 virtual network device (Microsoft's Packet Schedu

IP address: 10.91.35.61

Link-layer header type: Ethernet

☑ Capture packets in promiscuous mode

☐ Capture packets in pcap-ng format (experimental)

☐ Limit each packet to 1 bytes

Buffer size: 1 megabyte(s)

Capture Filter:

Capture File(s)

File:

☑ Use multiple files

☑ Next file every 1 megabyte(s)

☐ Next file every 1 minute(s)

☑ Ring buffer with 2 files

☐ Stop capture after 1 file(s)

Stop Capture ...

☐ ... after 1 packet(s)

☐ ... after 1 megabyte(s)

☐ ... after 1 minute(s)

Display Options

☑ Update list of packets in real time

☑ Automatic scrolling in live capture

☑ Hide capture info dialog

Name Resolution

☑ Enable MAC name resolution

☐ Enable network name resolution

☑ Enable transport name resolution

Figure 36. Wireshark Filter Configuration

3. Figure 37 shows the Capture Filter panel with different protocol options. You can select the **TCP only** option because the packets that will be mirrored from VM_01 will be TCP packets.

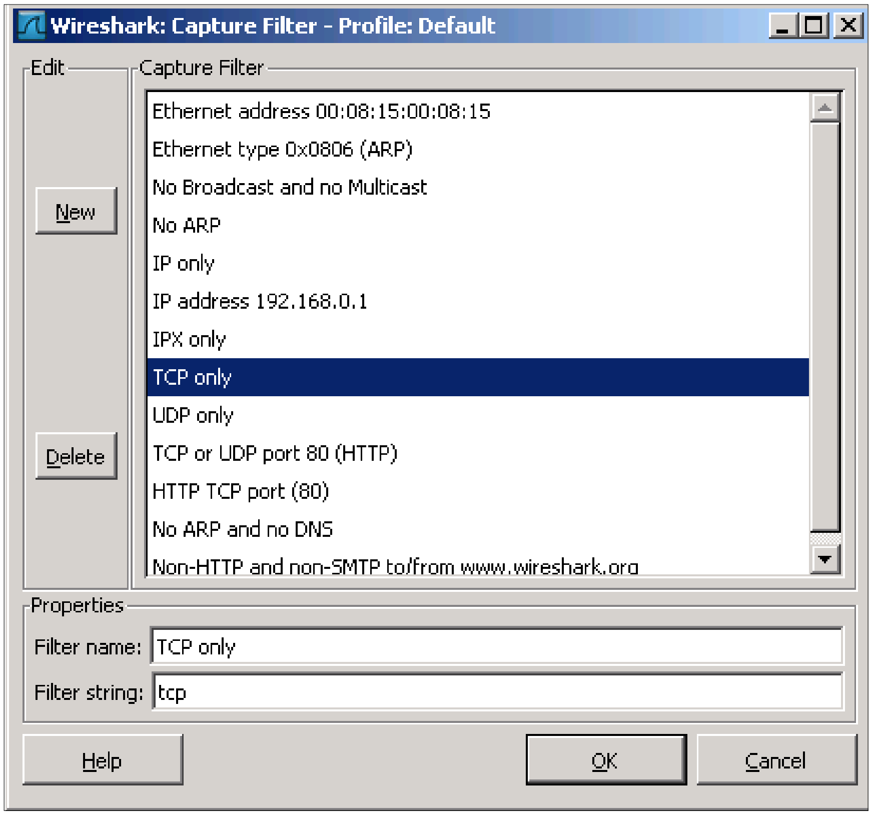

Ethernet address 00:08:15:00:08:15
Ethernet type 0x0806 (ARP)
No Broadcast and no Multicast
No ARP
IP only
IP address 192.168.0.1
IPX only
TCP only
UDP only
TCP or UDP port 80 (HTTP)
HTTP TCP port (80)
No ARP and no DNS
Non-HTTP and non-SMTP to/from www.wireshark.org

Filter name: TCP only
Filter string: tcp

Figure 37. Wireshark Filter Configuration

4. Click **Start** to start the capture process.

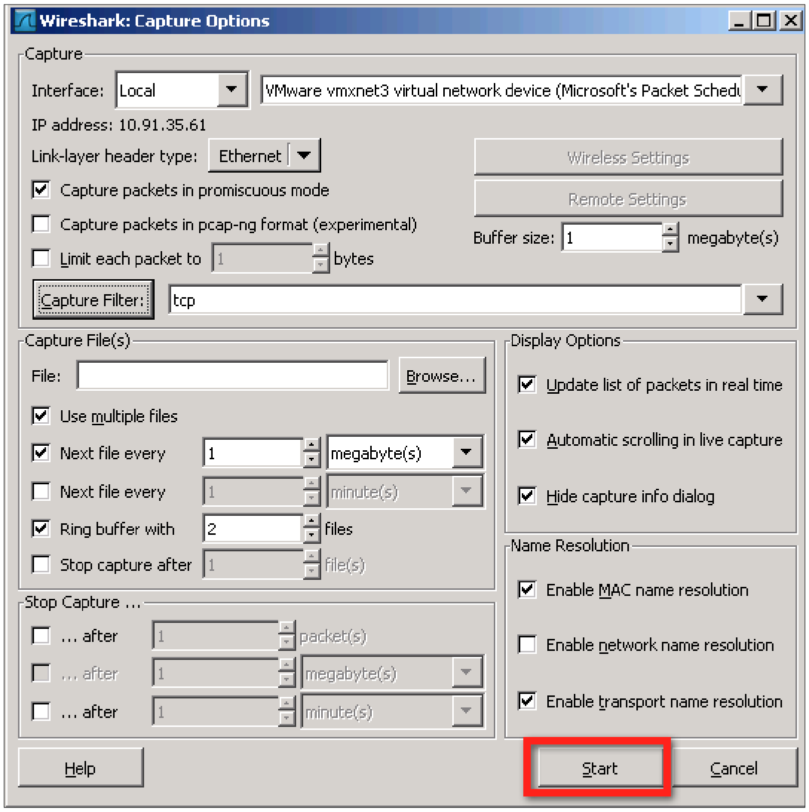

Wireshark: Capture Options

Capture

Interface: Local | VMware vmxnet3 virtual network device (Microsoft's Packet Schedu

IP address: 10.91.35.61

Link-layer header type: Ethernet

☑ Capture packets in promiscuous mode

☐ Capture packets in pcap-ng format (experimental)

☐ Limit each packet to 1 bytes

Capture Filter: tcp

Buffer size: 1 megabyte(s)

Capture File(s)

File:

☑ Use multiple files

☑ Next file every 1 megabyte(s)

☐ Next file every 1 minute(s)

☑ Ring buffer with 2 files

☐ Stop capture after 1 file(s)

Stop Capture ...

☐ ... after 1 packet(s)

☐ ... after 1 megabyte(s)

☐ ... after 1 minute(s)

Display Options

☑ Update list of packets in real time

☑ Automatic scrolling in live capture

☑ Hide capture info dialog

Name Resolution

☑ Enable MAC name resolution

☐ Enable network name resolution

☑ Enable transport name resolution

Figure 38. Start Capture

1. The screen shot in Figure 39 provides the packets captured by the Wireshark tool. These are the packets that are sent out by the virtual machine VM_01 as part of the TCP session. VDS then mirror these packets to the virtual machine VM_05. TCP session traffic consists of packets that flow to and from the two end points. In this example, VM_01 (10.91.35.60) and VM_03 (10.91.35.73) are the two end points of the TCP session. As you can see from the screen shot, the Wireshark Analyzer captures only traffic going out of VM_01. This is because the port mirroring session was configured to mirror **only ingress traffic** to the destination.

You can also define port mirroring sessions to mirror either both traffic or egress only traffic to the destination. Check the mirrored traffic on the destination using the Wireshark tool.

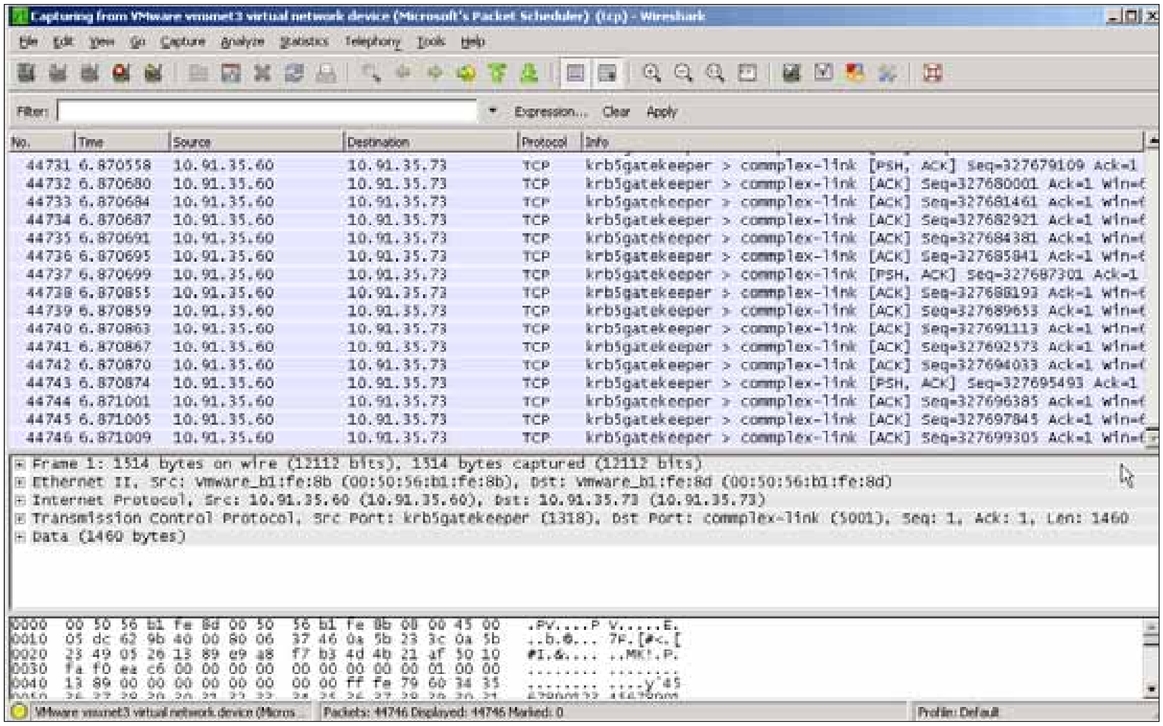
Capturing from VMware vmnet3 virtual network device (Microsoft's Packet Scheduler) (tcp) - Wireshark
File Edit View Go Capture Analyze Statistics Telephony Tools Help
Filter:
Expression... Clear Apply
No. Time Source Destination Protocol Info
VMware vmnet3 virtual network device (Micros...
Packets: 44746 Displayed: 44746 Marked: 0
Profile: Default

```
44731 6.870558  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [PSH, ACK] Seq=327679109 Ack=1
44732 6.870680  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327680001 Ack=1 Win=(
44733 6.870684  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327681461 Ack=1 Win=(
44734 6.870687  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327682921 Ack=1 Win=(
44735 6.870691  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327684381 Ack=1 Win=(
44736 6.870695  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327685841 Ack=1 Win=(
44737 6.870699  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [PSH, ACK] Seq=327687301 Ack=1
44738 6.870855  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327688193 Ack=1 Win=(
44739 6.870859  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327689653 Ack=1 Win=(
44740 6.870863  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327691113 Ack=1 Win=(
44741 6.870867  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327692573 Ack=1 Win=(
44742 6.870870  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327694033 Ack=1 Win=(
44743 6.870874  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [PSH, ACK] Seq=327695493 Ack=1
44744 6.871001  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327696385 Ack=1 Win=(
44745 6.871005  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327697845 Ack=1 Win=(
44746 6.871009  10.91.35.60  10.91.35.73  TCP  krb5gatekeeper > commplex-link [ACK] Seq=327699305 Ack=1 Win=(
```

```
⊞ Frame 1: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits)
⊞ Ethernet II, Src: Vmware_b1:fe:8b (00:50:56:b1:fe:8b), Dst: Vmware_b1:fe:8d (00:50:56:b1:fe:8d)
⊞ Internet Protocol, Src: 10.91.35.60 (10.91.35.60), Dst: 10.91.35.73 (10.91.35.73)
⊞ Transmission Control Protocol, Src Port: krb5gatekeeper (1318), Dst Port: commplex-link (5001), Seq: 1, Ack: 1, Len: 1460
⊞ Data (1460 bytes)
```

```
0000   00 50 56 b1 fe 8d 00 50  56 b1 fe 8b 08 00 45 00   .PV....P V.....E.
0010   05 dc 62 9b 40 00 80 06  37 46 0a 5b 23 3c 0a 5b   ..b.@... 7F.[#<.[
0020   23 49 05 26 13 89 e9 a8  f7 b3 4d 4b 21 af 50 10   #I.&.... ..MK!.P.
0030   fa f0 ea c6 00 00 00 00  00 00 00 00 00 01 00 00   ........ ........
0040   13 89 00 00 00 00 00 00  00 00 ff fe 79 60 34 35   ........ ....y`45
```

Figure 39. Captured Traffic

NIOC

Network I/O Control (NIOC) is the advanced feature of the VDS that provides traffic management capability. Network traffic management provides the required control and guarantee for different traffic types in the consolidated I/O environment. In the vSphere 5.0 platform, NIOC supports traffic management capabilities for the system, the virtual machine, and user-defined traffic types.

NIOC Rationale

Applications have different CPU, memory, and network I/O resource requirements. Business-critical applications have high resource requirements and higher Service Level Agreements (SLAs) as compared to noncritical applications. In the virtual infrastructure, where business-critical applications run along with noncritical applications, it becomes critical that resources are allocated according to the individual workload requirements.

The vSphere virtual platform provides you the capability to manage CPU, memory, and network resources. Network resources are managed through the NIOC feature on the VDS. When NIOC is enabled, the VDS traffic is divided into the following predefined network resource pools: VMware Fault Tolerance (FT) traffic, iSCSI traffic, vMotion traffic, management traffic, NFS traffic, and virtual machine traffic. The vSphere 5.0 release enhances the NIOC by enabling you to create user-defined network resource pools for any traffic type.

Evaluation Overview

In this example environment, first you will see the impact on the network I/O when only the virtual machine network resource pool is used and shared among different workloads. After that demonstration, you will configure user-defined network resource pools for individual workloads and see how it improves the network I/O performance.

Evaluation Environment for NIOC

The evaluation environment is same as that used during the NetFlow and port mirroring evaluation. It consists of the following components as shown in Figure 1:

1. Three virtual machines running Windows OS on Host1
2. Three virtual machines running Windows OS on Host2
3. Each virtual machine has the following software tool installed:
   a. JPerf tool
4. vSphere Management Assistant to provide remote command access

Challenges When Using the Virtual Machine Traffic Type for Multiple Workloads

In this example environment, you will use a predefined virtual machine traffic resource pool to allocate shares and limits to the virtual machine traffic from different workloads. In this approach all virtual machines (workloads) share the network resources allocated to the virtual machine traffic type. To begin, you have to configure the NIOC with the virtual machine traffic's shares and limits parameters.

First, enable the NIOC and configure the Virtual Machine Traffic type:

1. Enable the NIOC by selecting the **Properties** under the **Resource Allocation** tab. Figure 41 shows the panel where you should check the box.

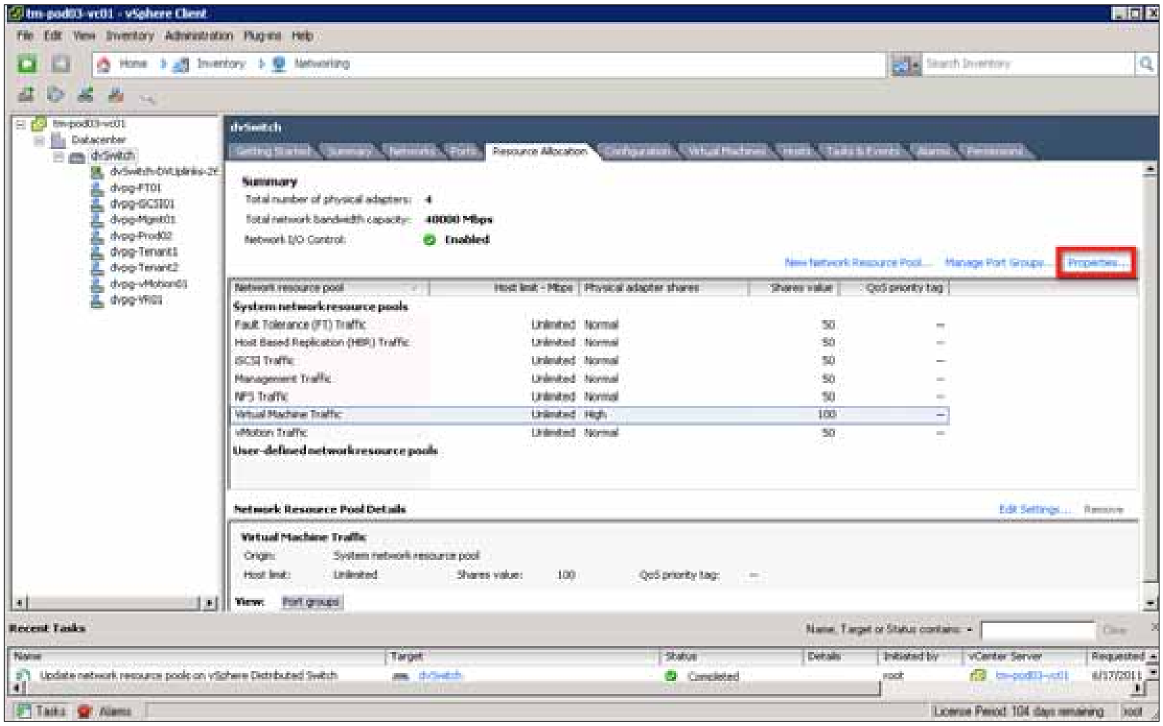

Summary

Total number of physical adapters: 4

Total network bandwidth capacity: 40000 Mbps

Network I/O Control: Enabled

Network resource pool	Host limit - Mbps	Physical adapter shares	Shares value	QoS priority tag
System network resource pools				
Fault Tolerance (FT) Traffic	Unlimited	Normal	50	--
Host Based Replication (HBR) Traffic	Unlimited	Normal	50	--
iSCSI Traffic	Unlimited	Normal	50	--
Management Traffic	Unlimited	Normal	50	--
NFS Traffic	Unlimited	Normal	50	--
Virtual Machine Traffic	Unlimited	High	100	--
vMotion Traffic	Unlimited	Normal	50	--
User-defined network resource pools				

Network Resource Pool Details

Virtual Machine Traffic

Origin: System network resource pool

Host limit: Unlimited Shares value: 100 QoS priority tag: --

Name	Target	Status	Details	Initiated by	vCenter Server	Requested
Update network resource pools on vSphere Distributed Switch	dvSwitch	Completed		root	tm-pod03-vc01	6/17/2011

Figure 40. How to Enable NIOC

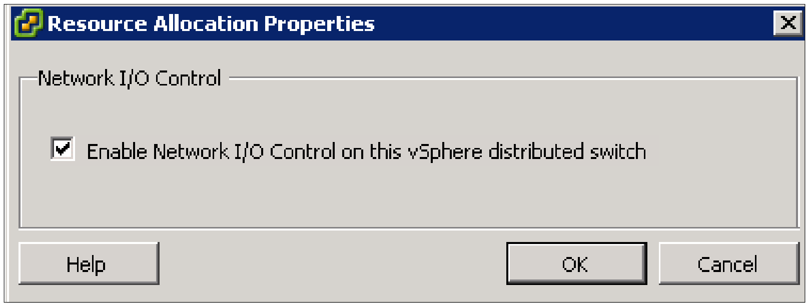

Resource Allocation Properties

Network I/O Control

☑ Enable Network I/O Control on this vSphere distributed switch

Figure 41. Enabling NIOC

2. Change the parameters of the virtual machine traffic type by clicking **Edit Settings** as shown in Figure 42.

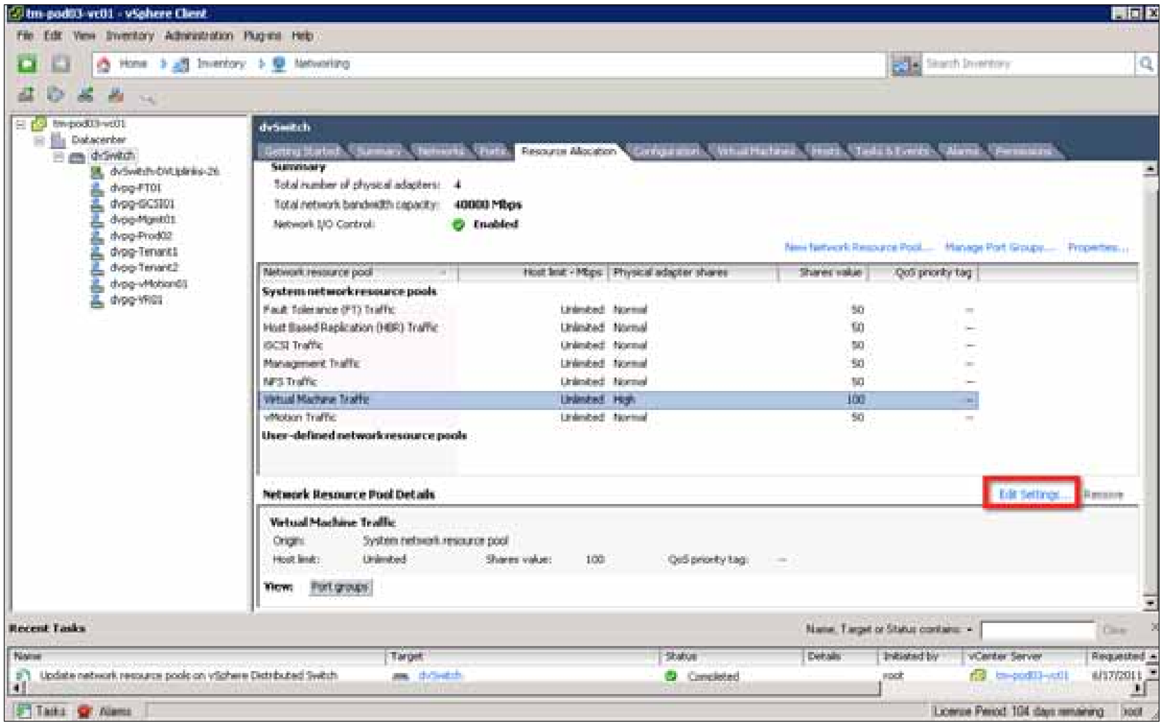

tm-pod03-vc01 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Networking
Search Inventory
tm-pod03-vc01
Datacenter
dvSwitch
dvSwitch-DVUplinks-26
dvpg-FT01
dvpg-iSCSI01
dvpg-Mgmt01
dvpg-Prod02
dvpg-Tenant1
dvpg-Tenant2
dvpg-vMotion01
dvpg-VR01
dvSwitch
Getting Started Summary Networks Ports Resource Allocation Configuration Virtual Machines Hosts Tasks & Events Alarms Permissions
New Network Resource Pool... Manage Port Groups... Properties...
Edit Settings... Remove
Tasks Alarms

Summary

Total number of physical adapters: 4

Total network bandwidth capacity: 40000 Mbps

Network I/O Control: Enabled

Network resource pool	Host limit - Mbps	Physical adapter shares	Shares value	QoS priority tag
System network resource pools				
Fault Tolerance (FT) Traffic	Unlimited	Normal	50	--
Host Based Replication (HBR) Traffic	Unlimited	Normal	50	--
iSCSI Traffic	Unlimited	Normal	50	--
Management Traffic	Unlimited	Normal	50	--
NFS Traffic	Unlimited	Normal	50	--
Virtual Machine Traffic	Unlimited	High	100	--
vMotion Traffic	Unlimited	Normal	50	--
User-defined network resource pools				

Network Resource Pool Details

Virtual Machine Traffic

Origin: System network resource pool

Host limit: Unlimited Shares value: 100 QoS priority tag: --

View: Port groups

Recent Tasks

Name, Target or Status contains: Clear

Name	Target	Status	Details	Initiated by	vCenter Server	Requested
Update network resource pools on vSphere Distributed Switch	dvSwitch	Completed		root	tm-pod03-vc01	6/17/2011

License Period: 104 days remaining root

Figure 42. Edit I/O Shares and Limits

3. You can limit the virtual machine traffic type by entering the bandwidth number in the host limit field. In this example, the virtual machine traffic is limited to 20Mbps. In this example environment, there are three VMs on Host 1 (VM_01, VM_02, VM_03) that will share this 20Mbps of network I/O capacity.

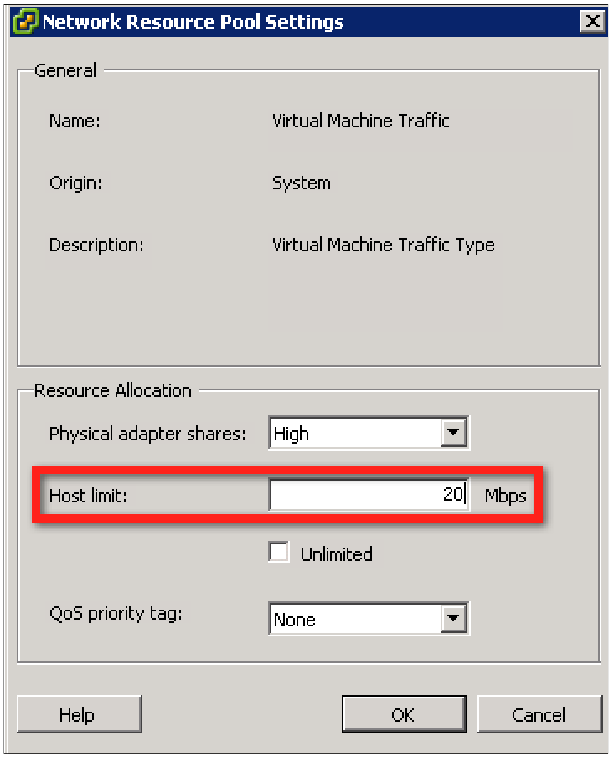

Network Resource Pool Settings

General

Name: Virtual Machine Traffic

Origin: System

Description: Virtual Machine Traffic Type

Resource Allocation

Physical adapter shares: High

Host limit: 20 Mbps

□ Unlimited

QoS priority tag: None

Figure 43. Limit Configuration

In the normal operation it is not recommended to limit the traffic. However, in this example, to minimize the amount of traffic you have to generate, the limit configuration is used.

After completing the NIOC configuration, you can generate traffic to simulate two workload scenarios. In the example environment, the traffic is generated using the JPerf tool between virtual machines running on different hosts as follows:

1. VM_01 on Host1 to VM_02 on Host2 TCP session
2. VM_03 on Host1 to VM_04 on Host2 TCP session

You can configure the JPerf server on VM_02 and VM_04 and the JPerf client on VM_01 and VM_03. For more details on how to configure the JPerf client and server, you can refer to steps described in the NetFlow evaluation section along with Figure 20 and Figure 21.

Once the setup of the two TCP sessions between the virtual machines is complete, you can monitor the performance of these flows from the JPerf client view. Figure 44 and Figure 45 shows the JPerf client screens of VM_03 and VM_01, respectively.

When 20Mbps of bandwidth is shared between two workloads, you can see how the bandwidth is unevenly distributed between the two TCP sessions. In this situation, if VM_03 utilizes more bandwidth, VM_01 suffers and vice versa. This ultimately impacts the performance of both TCP sessions. Consider this deployment with an important application workload that shares traffic with other workloads that are bursty in nature. Application traffic will suffer and consequently users will see the performance degradation in terms of response time and availability.

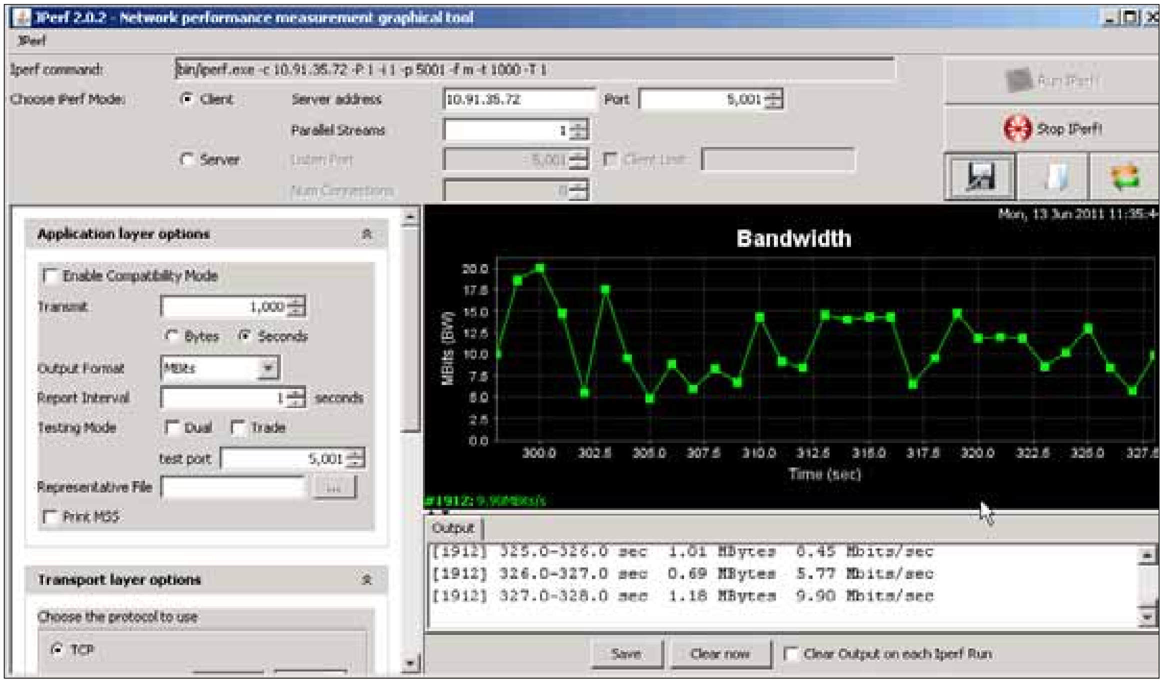

Iperf command: bin/iperf.exe -c 10.91.35.72 -P 1 -i 1 -p 5001 -f m -t 1000 -T 1

Choose iPerf Mode: ● Client Server address 10.91.35.72 Port 5,001

Parallel Streams 1

○ Server Listen Port 5,001 □ Client Limit

Num Connections 0

□ Enable Compatibility Mode

Transmit 1,000 ○ Bytes ● Seconds

Output Format MBits

Report Interval 1 seconds

Testing Mode □ Dual □ Trade

test port 5,001

Representative File

□ Print MSS

Choose the protocol to use

● TCP

```
[1912] 325.0-326.0 sec  1.01 MBytes  8.45 Mbits/sec
[1912] 326.0-327.0 sec  0.69 MBytes  5.77 Mbits/sec
[1912] 327.0-328.0 sec  1.18 MBytes  9.90 Mbits/sec
```

□ Clear Output on each Iperf Run

Figure 44. Client1-Side Bandwidth

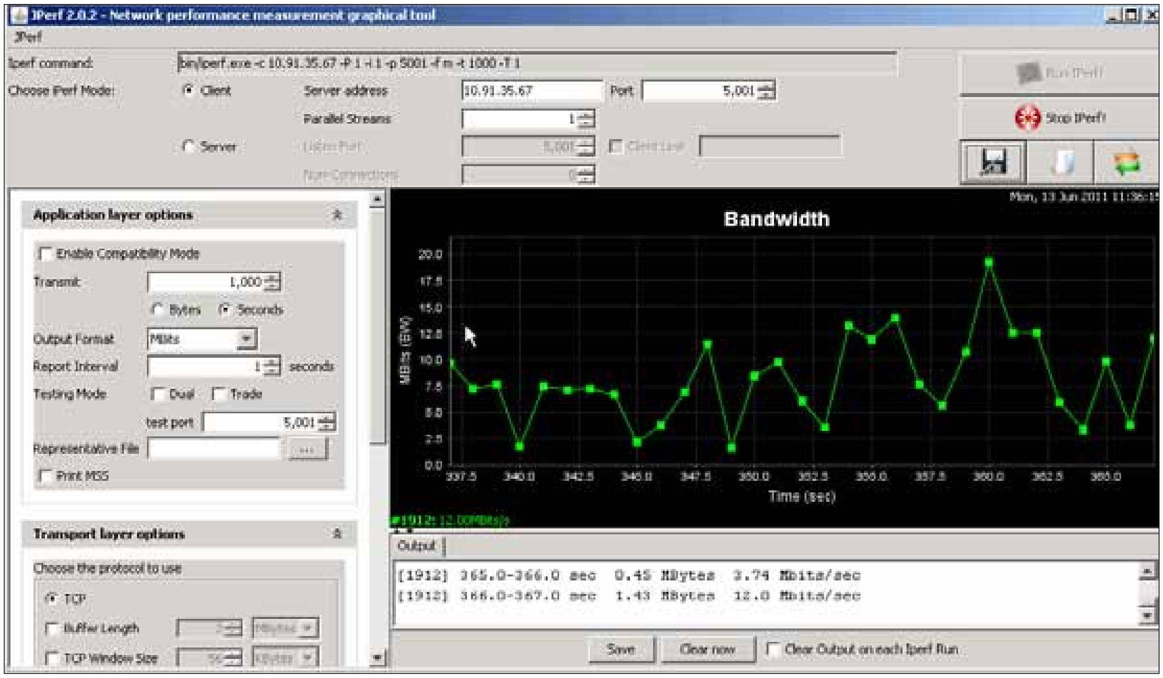

Iperf command: bin/iperf.exe -c 10.91.35.67 -P 1 -i 1 -p 5001 -f m -t 1000 -T 1

Choose iPerf Mode: ● Client Server address 10.91.35.67 Port 5,001

Parallel Streams 1

○ Server Listen Port 5,001 □ Client Limit

Num Connections 0

□ Enable Compatibility Mode

Transmit 1,000 ○ Bytes ● Seconds

Output Format MBits

Report Interval 1 seconds

Testing Mode □ Dual □ Trade

test port 5,001

Representative File

□ Print MSS

Choose the protocol to use

● TCP

□ Buffer Length 2 MBytes

□ TCP Window Size 56 KBytes

```
[1912] 365.0-366.0 sec  0.45 MBytes  3.74 Mbits/sec
[1912] 366.0-367.0 sec  1.43 MBytes  12.0 Mbits/sec
```

□ Clear Output on each Iperf Run

Figure 45. Client2-Side Bandwidth

This network I/O issue is addressed through the new user-defined resource pools feature in the vSphere 5.0 platform. The following section will demonstrate the advantage of this new capability.

Configuring User-Defined Resource Pools

In this step, you will make use of the user-defined resource pools and allocate those resource pools to different workloads. This approach of allocating resources to individual workloads eliminates the problems faced by the use of the virtual machine traffic type for different workloads.

To define user-defined resource pools, follow these steps:

1. Switch to the **Home > Inventory > Networking** view.
2. Select the dvSwitch and choose the **Resource Allocation** tab on the right panel.
3. Click **New Network Resource Pool.**

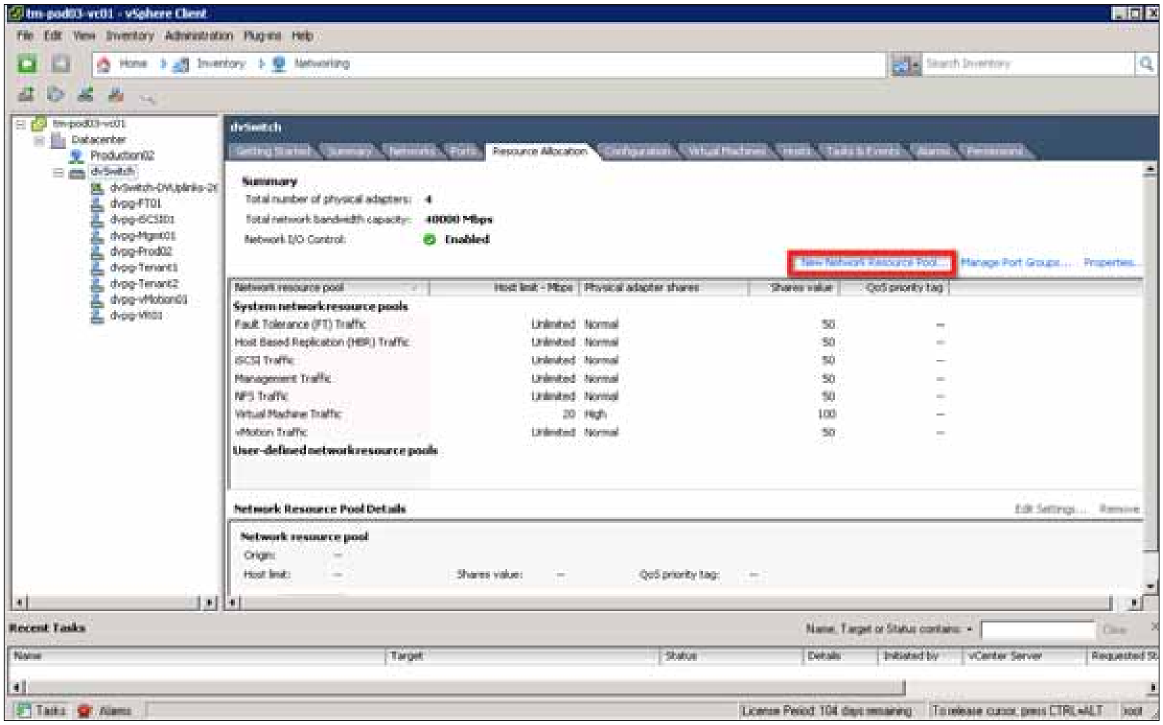

Summary

Total number of physical adapters: 4

Total network bandwidth capacity: 40000 Mbps

Network I/O Control: Enabled

Network resource pool	Host limit - Mbps	Physical adapter shares	Shares value	QoS priority tag
System network resource pools				
Fault Tolerance (FT) Traffic	Unlimited	Normal	50	--
Host Based Replication (HBR) Traffic	Unlimited	Normal	50	--
iSCSI Traffic	Unlimited	Normal	50	--
Management Traffic	Unlimited	Normal	50	--
NFS Traffic	Unlimited	Normal	50	--
Virtual Machine Traffic	20	High	100	--
vMotion Traffic	Unlimited	Normal	50	--
User-defined network resource pools				

Network Resource Pool Details

Network resource pool

Origin: --

Host limit: -- Shares value: -- QoS priority tag: --

Name	Target	Status	Details	Initiated by	vCenter Server	Requested St

Figure 46. User-Defined Resource Pool Configuration

4. In the panel shown in Figure 47, provide a name for the new network resource pool. In this example, this resource pool will be associated with the VM_01 workload. The host limit is set to 10Mbps. The option of QoS priority tag helps in tagging the packets with the 802.1p tag. You can use this option so that the network infrastructure treats the packets according to the priority and thus provides End-to-End QoS. In this example, the packets are not tagged.

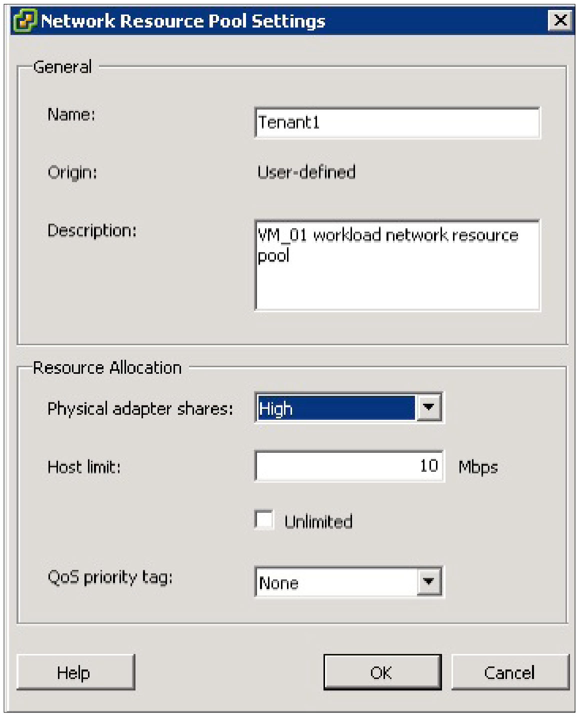

Network Resource Pool Settings

General

Name: Tenant1

Origin: User-defined

Description: VM_01 workload network resource pool

Resource Allocation

Physical adapter shares: High

Host limit: 10 Mbps

□ Unlimited

QoS priority tag: None

Figure 47. Tenant1 Configuration

5. Repeat Step 4 and define another resource pool for the VM_02 workload with the same shares and limits parameters.

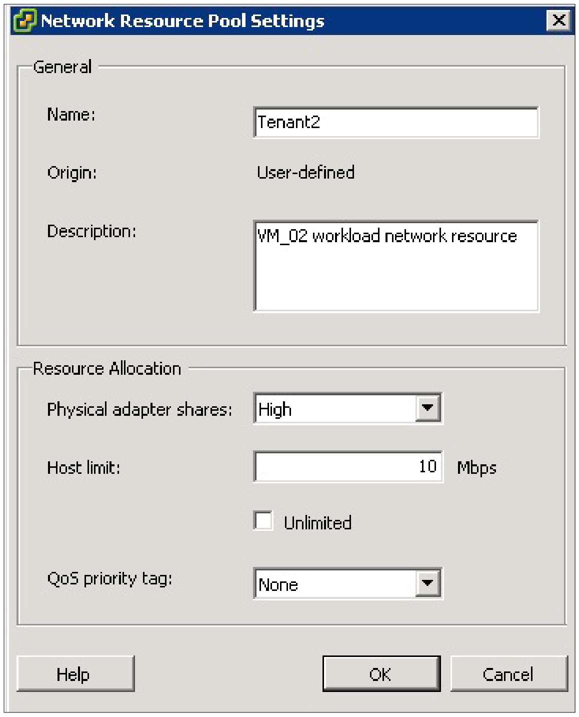

Network Resource Pool Settings

General

Name: Tenant2

Origin: User-defined

Description: VM_02 workload network resource

Resource Allocation

Physical adapter shares: High

Host limit: 10 Mbps

□ Unlimited

QoS priority tag: None

Figure 48. Tenant2 Configuration

6. Figure 49 shows the screen shot of the Resource Allocation tab view after custom resources are created.

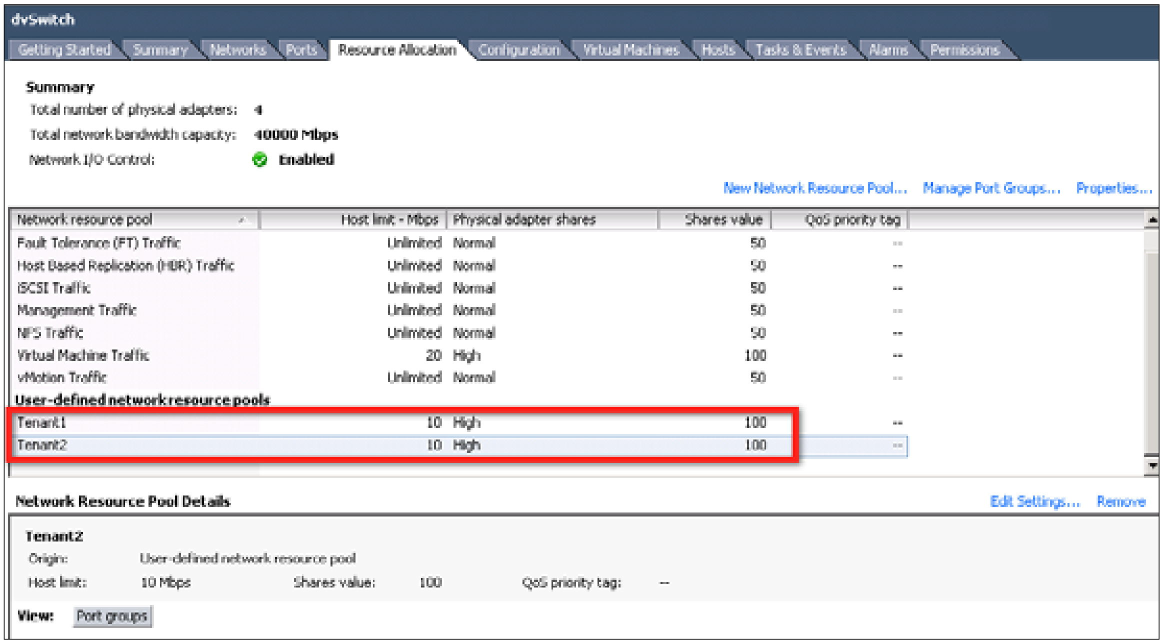

Summary

Total number of physical adapters: **4**

Total network bandwidth capacity: **40000 Mbps**

Network I/O Control: ✓ **Enabled**

Network resource pool	Host limit - Mbps	Physical adapter shares	Shares value	QoS priority tag
Fault Tolerance (FT) Traffic	Unlimited	Normal	50	--
Host Based Replication (HBR) Traffic	Unlimited	Normal	50	--
iSCSI Traffic	Unlimited	Normal	50	--
Management Traffic	Unlimited	Normal	50	--
NFS Traffic	Unlimited	Normal	50	--
Virtual Machine Traffic	20	High	100	--
vMotion Traffic	Unlimited	Normal	50	--
User-defined network resource pools				
Tenant1	10	High	100	--
Tenant2	10	High	100	--

Network Resource Pool Details

Tenant2

Origin: User-defined network resource pool

Host limit: 10 Mbps Shares value: 100 QoS priority tag: --

View: Port groups

Figure 49. Network I/O Resource Allocation View

Associating New Resource Pools with Tenant Port Groups

After creating the custom resource pools, you have to associate them with DV port groups. Once the resource pool is associated with a DV port group, the virtual machine connected to the DV port group gets the allocated network I/O resources.

In this example environment, you have two new resource pools, Tenant1 and Tenant2. Also, you have already defined two port groups named dvpg-Tenant1 and dvpg-Tenant2. Follow the steps to associate the Tenant1 resource pool with the dvpg-Tenant1 port group and the Tenant2 resource pool with the dvpg-Tenant2 port group.

1. Under the **Home > Inventory > Networking** view, select **dvpg-Tenant1.**
2. Right-click **dvpg-Tenant1** and select **Edit Settings.**

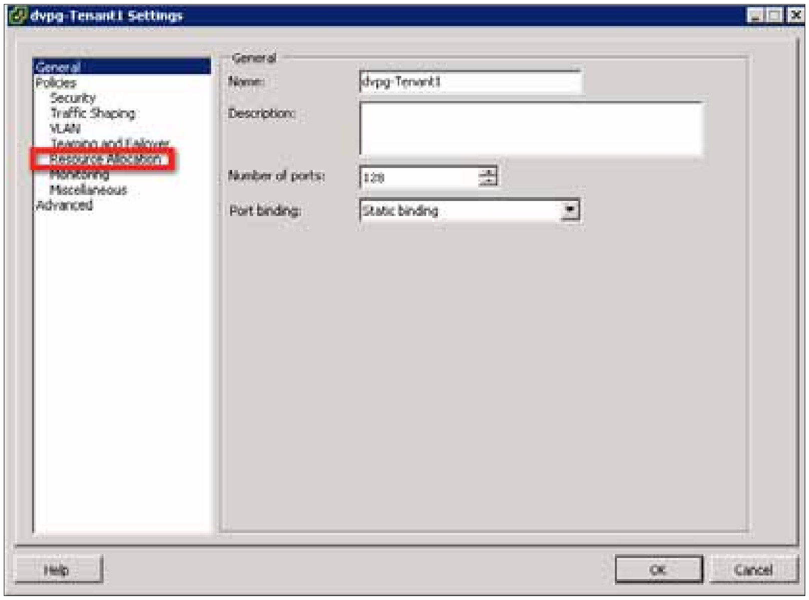

dvpg-Tenant1 Settings

General

Name: dvpg-Tenant1

Description:

Number of ports: 128

Port binding: Static binding

Figure 50. Tenant1 Port Group

3. Select **Resource Allocation** in the left panel and in the Policies pane on the right, click the **Network Resource Pool** pull-down menu. You will see the two new resource pools that were created in earlier steps.

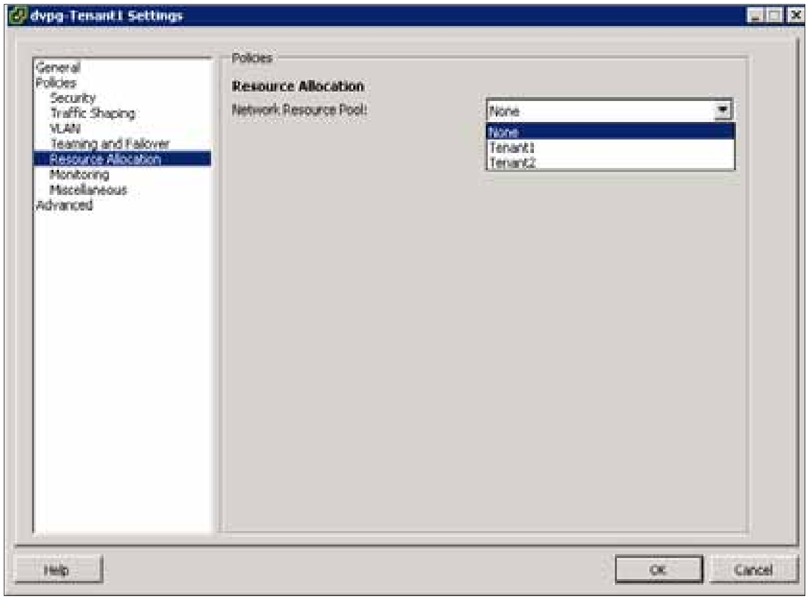

dvpg-Tenant1 Settings

Policies

Resource Allocation

Network Resource Pool: None

Figure 51. Associating Tenant1 Port Group with Tenant1 Resource Pool

4. Choose the **Tenant1** resource pool and click **OK.**

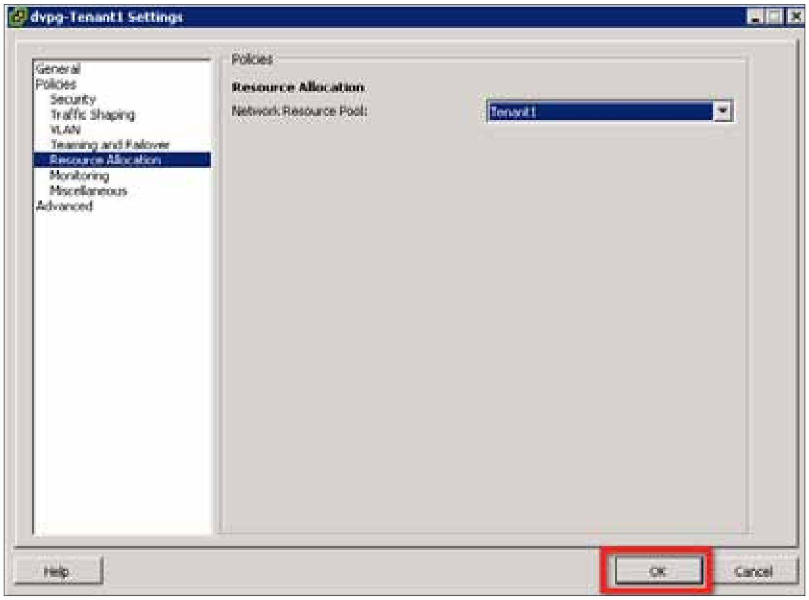
dvpg-Tenant1 Settings

Policies

Resource Allocation

Network Resource Pool: Tenant1

Figure 52. Association in Progress

5. Repeat steps 2 to 4 to associate the dvpg-Tenant2 port group with the Tenant2 resource pool.

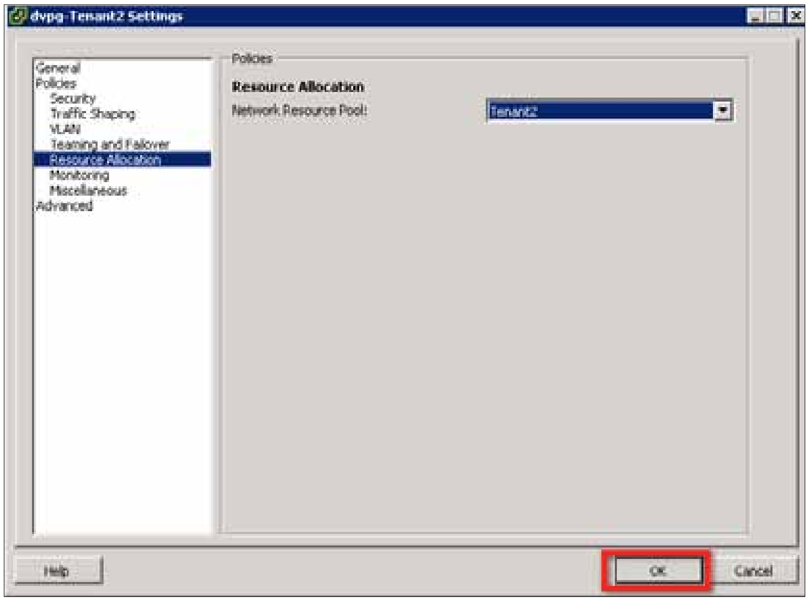
dvpg-Tenant2 Settings

Policies

Resource Allocation

Network Resource Pool: Tenant2

Figure 53. Association Complete

Moving Virtual Machines to the Tenant Port Groups

In this step, you will move virtual machines VM_01 and VM_03 running on Host 1 from the dvpg-Prod02 port group to new tenant port groups. When you move VM_01 to the dvpg-Tenant1 port group, you allocate VM_01 with network I/O resources defined by the Tenant1 resource pool. Similarly, moving VM_03 to the dvpg-Tenant2 port group allocates network I/O resources defined by the Tenant2 resource pool.

Follow these steps to perform the virtual machine transition to new port groups.

1. Switch to the **Home > Inventory > Hosts and Clusters** view.
2. Select **VM_01** and click **Edit Settings.**

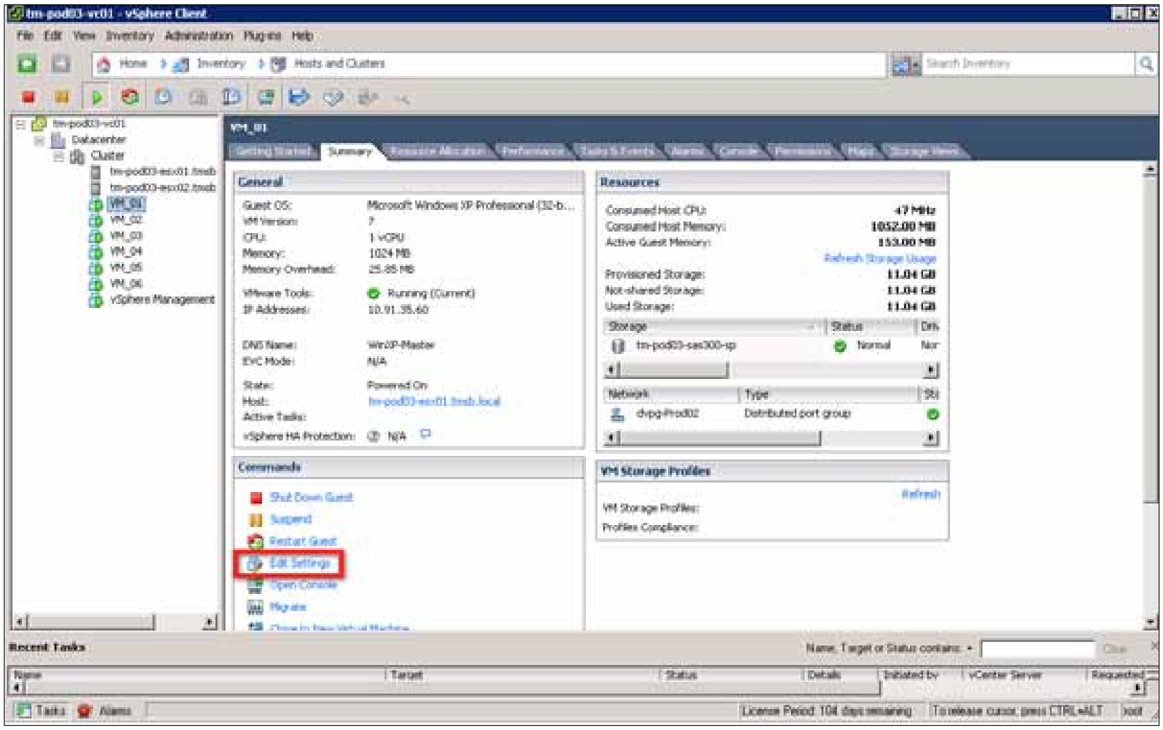

General

Guest OS: Microsoft Windows XP Professional (32-b...
VM Version: 7
CPU: 1 vCPU
Memory: 1024 MB
Memory Overhead: 25.85 MB
VMware Tools: Running (Current)
IP Addresses: 10.91.35.60
DNS Name: WinXP-Master
EVC Mode: N/A
State: Powered On
Host: tm-pod03-esx01.tmsb.local
Active Tasks:
vSphere HA Protection: N/A

Resources

Consumed Host CPU: 47 MHz
Consumed Host Memory: 1052.00 MB
Active Guest Memory: 153.00 MB
Refresh Storage Usage
Provisioned Storage: 11.04 GB
Not-shared Storage: 11.04 GB
Used Storage: 11.04 GB

Storage	Status	Dri
tm-pod03-sas300-sp	Normal	Nor

Network	Type	Sta
dvpg-Prod02	Distributed port group	✓

VM Storage Profiles

Refresh
VM Storage Profiles:
Profiles Compliance:

Name	Target	Status	Details	Initiated by	vCenter Server	Requested

Figure 54. Changing Virtual Machine Port Group

3. Choose **Network adapter 1,** and from the **Network label** drop-down menu, select **dvpg-Tenant1.**

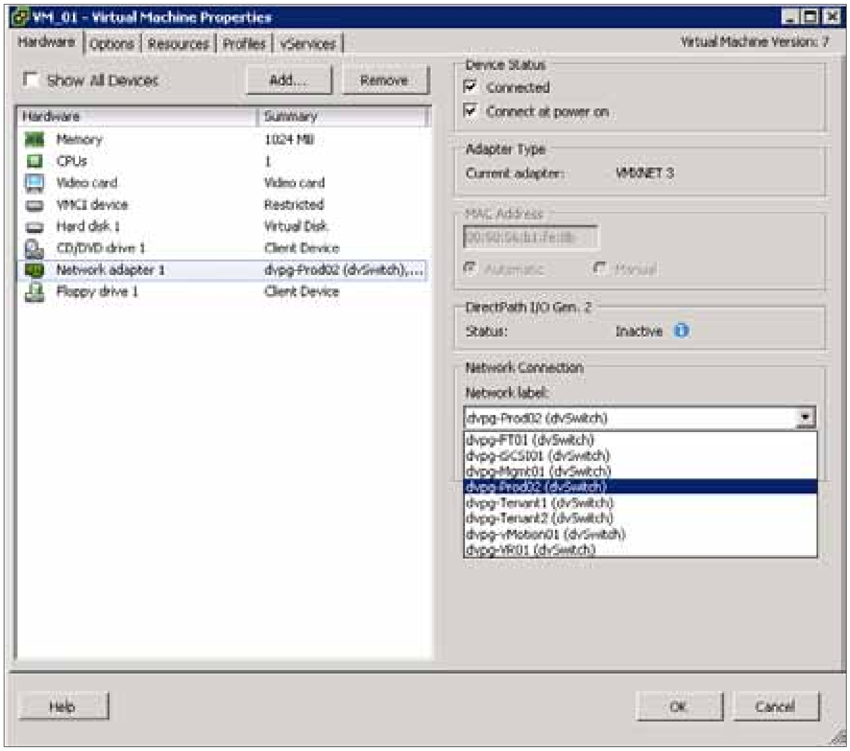

VM_01 - Virtual Machine Properties

Hardware | Options | Resources | Profiles | vServices — Virtual Machine Version: 7

□ Show All Devices — Add... — Remove

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
Hard disk 1	Virtual Disk
CD/DVD drive 1	Client Device
Network adapter 1	dvpg-Prod02 (dvSwitch),...
Floppy drive 1	Client Device

Device Status
☑ Connected
☑ Connect at power on

Adapter Type
Current adapter: VMXNET 3

MAC Address
00:50:56:b1:fe:8b
● Automatic ○ Manual

DirectPath I/O Gen. 2
Status: Inactive

Network Connection
Network label:
dvpg-Prod02 (dvSwitch)

- dvpg-FT01 (dvSwitch)
- dvpg-iSCSI01 (dvSwitch)
- dvpg-Mgmt01 (dvSwitch)
- dvpg-Prod02 (dvSwitch)
- dvpg-Tenant1 (dvSwitch)
- dvpg-Tenant2 (dvSwitch)
- dvpg-vMotion01 (dvSwitch)
- dvpg-VR01 (dvSwitch)

Help — OK — Cancel

Figure 55. Changing It to dvpg-Tenant1

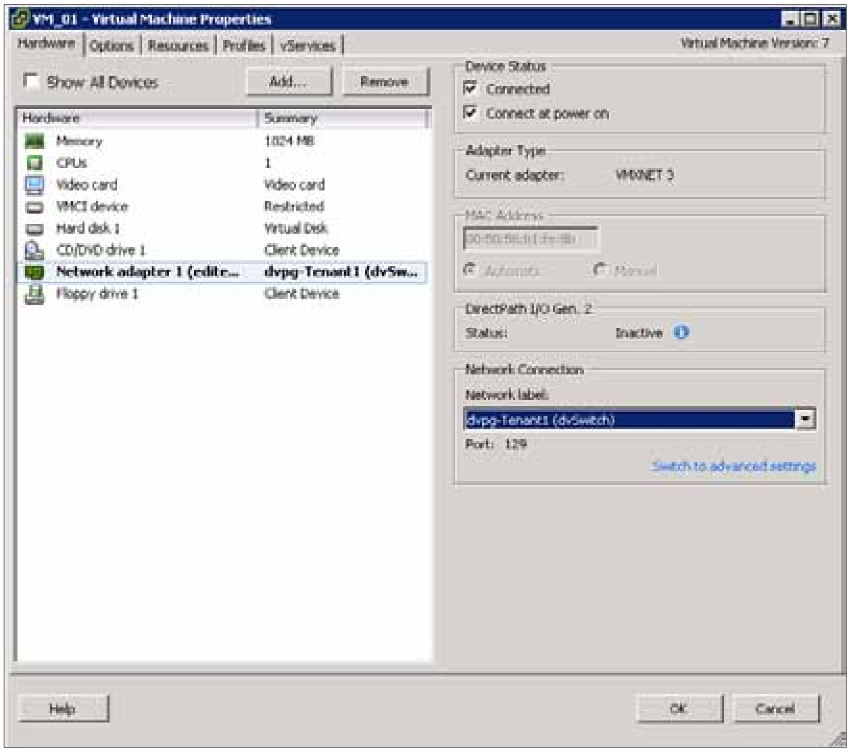

VM_01 - Virtual Machine Properties

Hardware | Options | Resources | Profiles | vServices — Virtual Machine Version: 7

□ Show All Devices — Add... — Remove

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
Hard disk 1	Virtual Disk
CD/DVD drive 1	Client Device
Network adapter 1 (edite...	**dvpg-Tenant1 (dvSw...**
Floppy drive 1	Client Device

Device Status
☑ Connected
☑ Connect at power on

Adapter Type
Current adapter: VMXNET 3

MAC Address
00:50:56:b1:fe:8b
● Automatic ○ Manual

DirectPath I/O Gen. 2
Status: Inactive

Network Connection
Network label:
dvpg-Tenant1 (dvSwitch)
Port: 129
Switch to advanced settings

Help — OK — Cancel

Figure 56. Migration Complete

4. Figure 57 shows the screen shot after VM_01 is moved to the dvpg-Tenant1 port group.

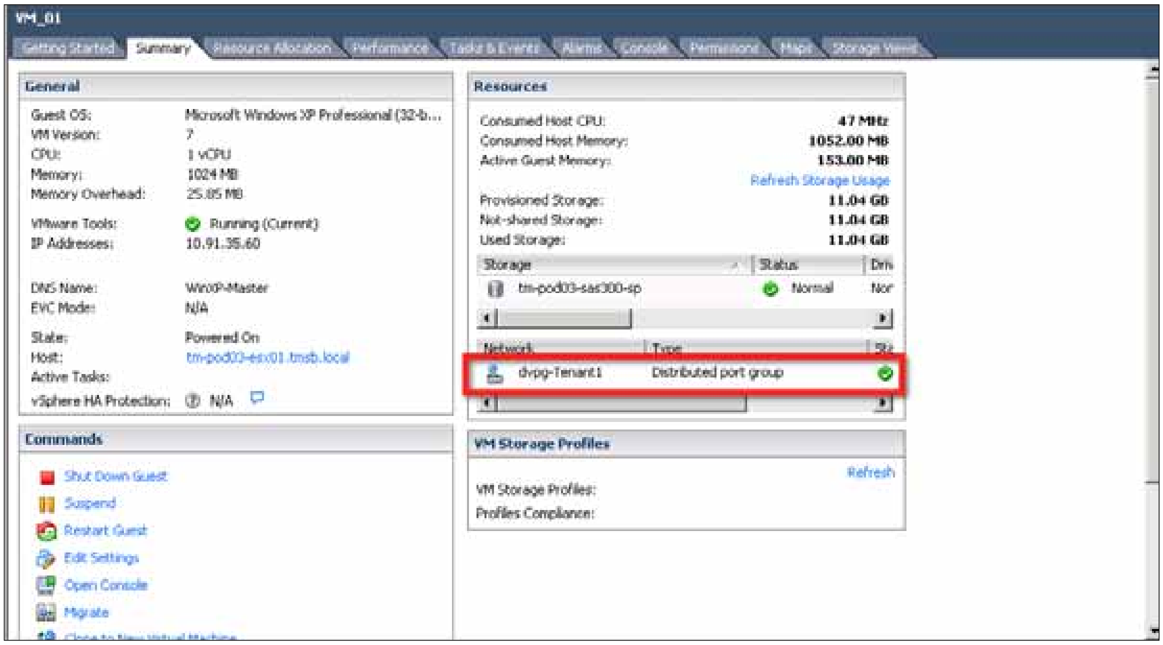

General

Guest OS: Microsoft Windows XP Professional (32-b...
VM Version: 7
CPU: 1 vCPU
Memory: 1024 MB
Memory Overhead: 25.85 MB

VMware Tools: Running (Current)
IP Addresses: 10.91.35.60

DNS Name: WinXP-Master
EVC Mode: N/A

State: Powered On
Host: tm-pod03-esx01.tmsb.local
Active Tasks:
vSphere HA Protection: N/A

Commands

- Shut Down Guest
- Suspend
- Restart Guest
- Edit Settings
- Open Console
- Migrate
- Clone to New Virtual Machine

Resources

Consumed Host CPU: 47 MHz
Consumed Host Memory: 1052.00 MB
Active Guest Memory: 153.00 MB
Refresh Storage Usage
Provisioned Storage: 11.04 GB
Not-shared Storage: 11.04 GB
Used Storage: 11.04 GB

Storage	Status	Dri
tm-pod03-sas300-sp	Normal	Nor

Network	Type	Sta
dvpg-Tenant1	Distributed port group	✓

VM Storage Profiles

Refresh
VM Storage Profiles:
Profiles Compliance:

Figure 57. VM_01 on dvpg-Tenant1 Port Group

5. Repeat steps 2 to 4 to move VM_03 to the dvpg-Tenant2 port group.

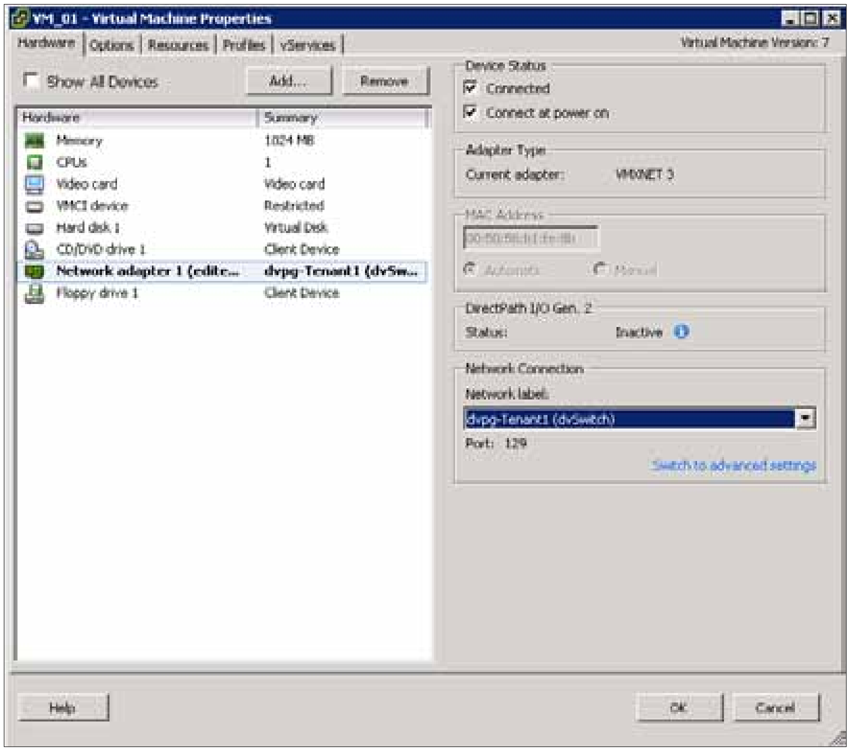

Virtual Machine Version: 7

□ Show All Devices

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
Hard disk 1	Virtual Disk
CD/DVD drive 1	Client Device
Network adapter 1 (edite...	**dvpg-Tenant1 (dvSw...**
Floppy drive 1	Client Device

Device Status
☑ Connected
☑ Connect at power on

Adapter Type
Current adapter: VMXNET 3

MAC Address
00:50:56:b1:fe:8b
● Automatic ○ Manual

DirectPath I/O Gen. 2
Status: Inactive

Network Connection
Network label:
dvpg-Tenant1 (dvSwitch)
Port: 129
Switch to advanced settings

Figure 58. Migrating VM_03

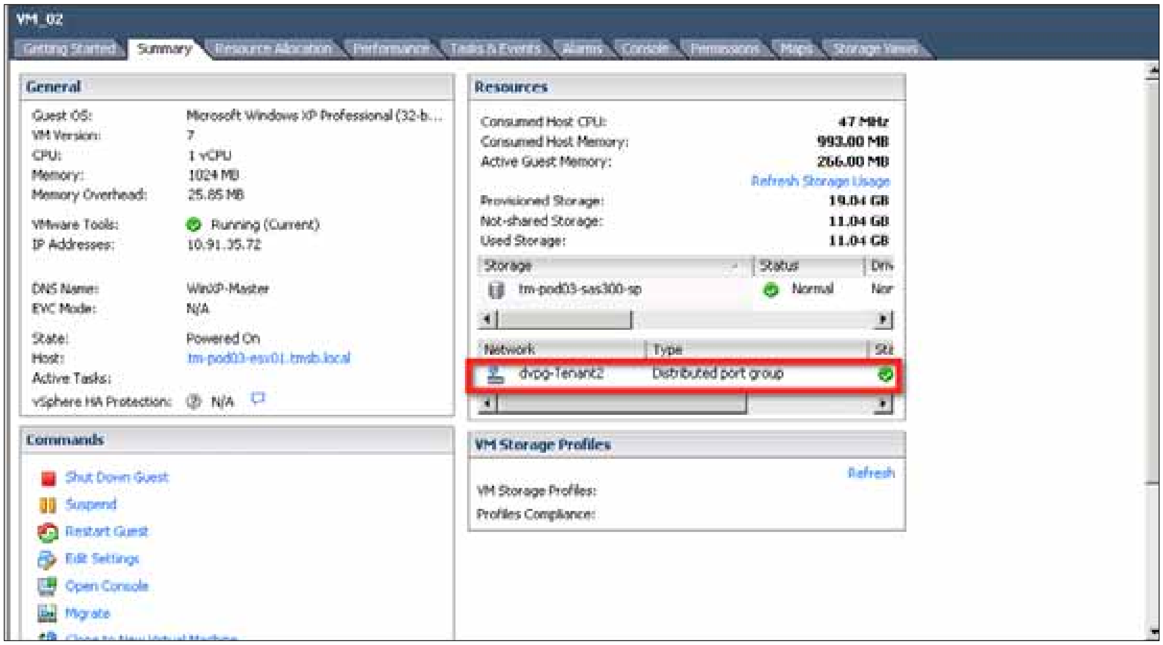

General

Guest OS: Microsoft Windows XP Professional (32-b...
VM Version: 7
CPU: 1 vCPU
Memory: 1024 MB
Memory Overhead: 25.85 MB

VMware Tools: Running (Current)
IP Addresses: 10.91.35.72

DNS Name: WinXP-Master
EVC Mode: N/A

State: Powered On
Host: tm-pod03-esx01.tmdb.local
Active Tasks:
vSphere HA Protection: N/A

Resources

Consumed Host CPU: 47 MHz
Consumed Host Memory: 993.00 MB
Active Guest Memory: 266.00 MB
Refresh Storage Usage
Provisioned Storage: 19.04 GB
Not-shared Storage: 11.04 GB
Used Storage: 11.04 GB

Storage	Status	Driv
tm-pod03-sas300-sp	Normal	Nor

Network	Type	Sta
dvpg-Tenant2	Distributed port group	✓

VM Storage Profiles

VM Storage Profiles:
Profiles Compliance:

Figure 59. Migration Complete

Testing the I/O Performance of the Virtual Machines

After completing the virtual machines' transition to tenant port groups, you can generate traffic to simulate two workload scenarios. In the example environment, the traffic is generated using the JPerf tool between virtual machines running on different hosts as follows:

1. VM_01 on Host1 to VM_02 on Host2 TCP session
2. VM_03 on Host1 to VM_04 on Host2 TCP session

You can configure the JPerf server on VM_02, VM_04 and the JPerf client on VM_01 and VM_03. For more details on how to configure the JPerf client and server, you can refer to steps described in the NetFlow evaluation section along with Figure 17 and Figure 18.

Once the setup of the two TCP sessions between the virtual machines is complete, you can monitor the performance of these flows from the JPerf client view. Figure 60 and Figure 61 show the JPerf client screens of VM_01 and VM_03, respectively.

The bandwidth charts in both the figures indicate that the two TCP sessions get uniform network I/O resources. This is because you have isolated the two workloads by assigning them to their individual resource pool. There is no sharing of bandwidth as it happens with one virtual machine traffic type. This demonstrates the advantage of using custom resource pools.

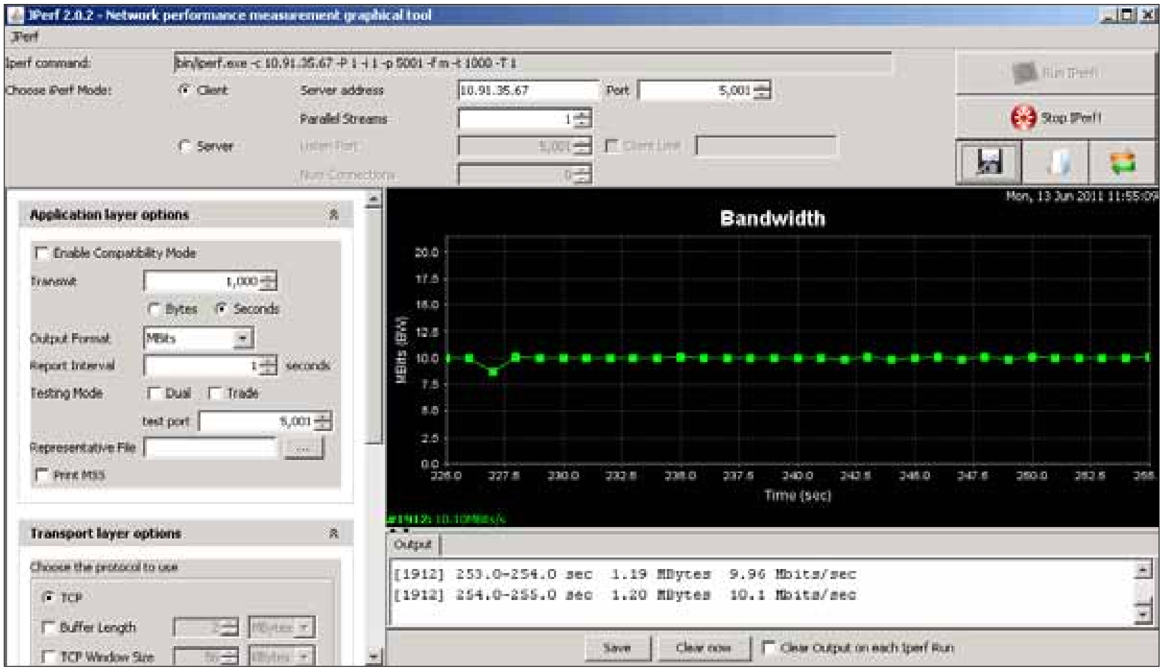

Iperf command: bin/iperf.exe -c 10.91.35.67 -P 1 -i 1 -p 5001 -f m -t 1000 -T 1

Choose iPerf Mode: ● Client   Server address 10.91.35.67   Port 5,001

Parallel Streams 1

○ Server   Listen Port 5,001   □ Client Limit

Num Connections 0

Application layer options

□ Enable Compatibility Mode

Transmit 1,000   ○ Bytes ● Seconds

Output Format MBits

Report Interval 1 seconds

Testing Mode □ Dual □ Trade

test port 5,001

Representative File

□ Print MSS

Transport layer options

Choose the protocol to use

● TCP

□ Buffer Length 2 MBytes

□ TCP Window Size 56 KBytes

Output

```
[1912] 253.0-254.0 sec  1.19 MBytes  9.96 Mbits/sec
[1912] 254.0-255.0 sec  1.20 MBytes  10.1 Mbits/sec
```

□ Clear Output on each Iperf Run

Figure 60. JPerf Client Screen of VM_01

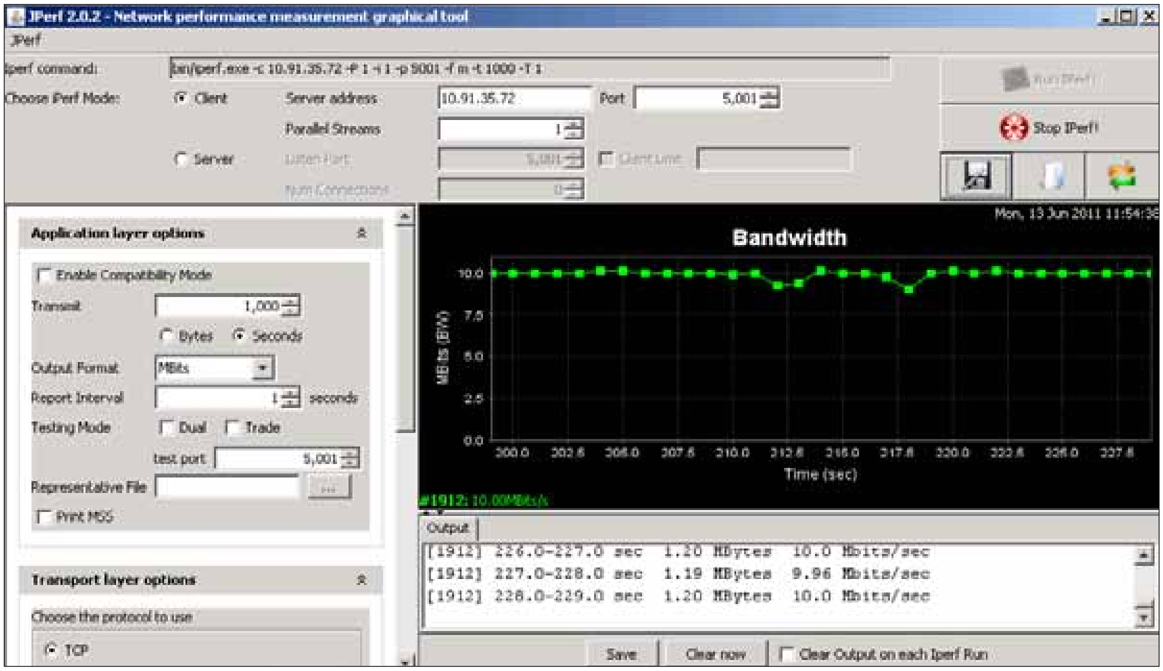

Iperf command: bin/iperf.exe -c 10.91.35.72 -P 1 -i 1 -p 5001 -f m -t 1000 -T 1

Choose iPerf Mode: ● Client   Server address 10.91.35.72   Port 5,001

Parallel Streams 1

○ Server   Listen Port 5,001   □ Client Limit

Num Connections 0

Application layer options

□ Enable Compatibility Mode

Transmit 1,000   ○ Bytes ● Seconds

Output Format MBits

Report Interval 1 seconds

Testing Mode □ Dual □ Trade

test port 5,001

Representative File

□ Print MSS

Transport layer options

Choose the protocol to use

● TCP

Output

```
[1912] 226.0-227.0 sec  1.20 MBytes  10.0 Mbits/sec
[1912] 227.0-228.0 sec  1.19 MBytes  9.96 Mbits/sec
[1912] 228.0-229.0 sec  1.20 MBytes  10.0 Mbits/sec
```

□ Clear Output on each Iperf Run

Figure 61. JPerf Client Screen of VM_03

You can also use vSphere Management Assitant to check the network I/O utilization. Run the resxtop command on Host1 to get the network utilization of the running virtual machines.

Run "resxtop –server=10.91.32.23 –vihost tm-pod03-esx01.tmsb.local" and then press **n** for network stats. Figure 62 shows the screen shot of the network utilization. You can see the bandwidth used by the VM_01 and VM_03 virtual machines.

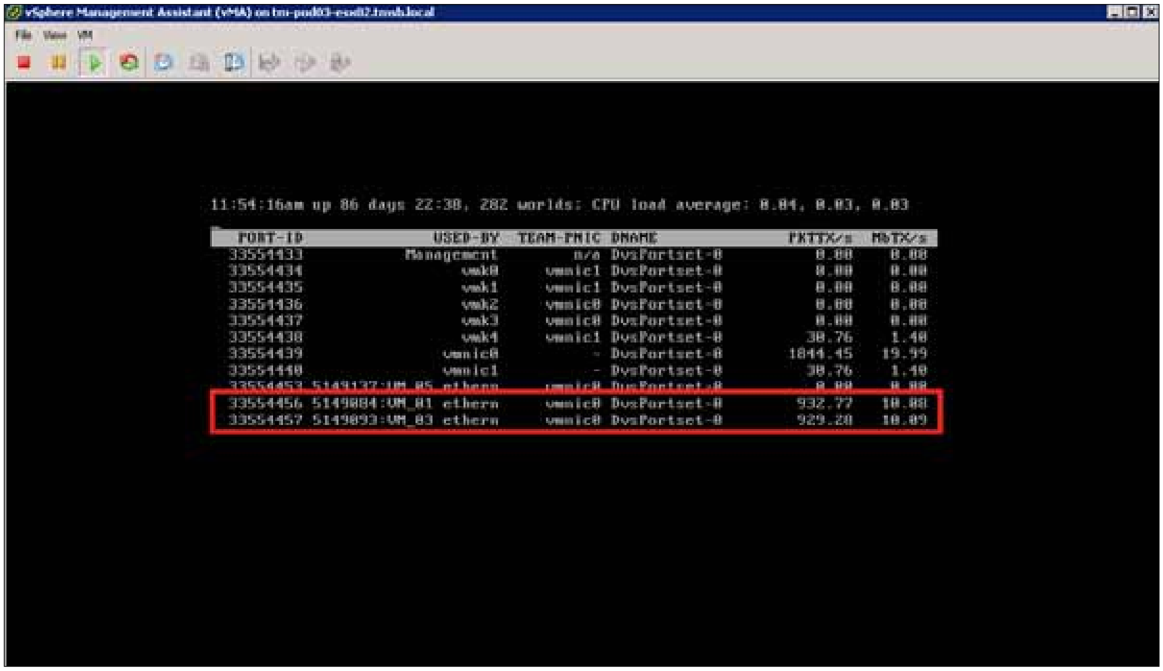
vSphere Management Assistant (vMA) on tm-pod03-esx02.tmsb.local
File
View
VM

```
11:54:16am up 86 days 22:30, 282 worlds; CPU load average: 0.04, 0.03, 0.03

   PORT-ID              USED-BY  TEAM-PNIC DNAME              PKTTX/s  MbTX/s
  33554433           Management        n/a DvsPortset-0          0.00    0.00
  33554434                 vmk0     vmnic1 DvsPortset-0          0.00    0.00
  33554435                 vmk1     vmnic1 DvsPortset-0          0.00    0.00
  33554436                 vmk2     vmnic0 DvsPortset-0          0.00    0.00
  33554437                 vmk3     vmnic0 DvsPortset-0          0.00    0.00
  33554438                 vmk4     vmnic1 DvsPortset-0         38.76    1.40
  33554439               vmnic0          - DvsPortset-0       1844.45   19.99
  33554440               vmnic1          - DvsPortset-0         38.76    1.40
  33554453 5149137:VM_05 ethern     vmnic0 DvsPortset-0          0.00    0.00
  33554456 5149084:VM_01 ethern     vmnic0 DvsPortset-0        932.77   10.08
  33554457 5149093:VM_03 ethern     vmnic0 DvsPortset-0        929.20   10.09
```

Figure 62. Screen Shot of the Network Utilization

Conclusion

The VMware vSphere 5.0 platform provides the visibility in virtual machine traffic through the NetFlow and port mirroring features and enhances network I/O control through user-defined resource pools. These new networking features help network administrators when troubleshooting network issues, and they provide advanced traffic management capability. This evaluation guide covered the step-by-step configuration of these new features and also provided simple exercises on how to test these features. After going through these evaluation exercises in this guide, you should be able to see how these new features can benefit your virtual infrastructure and cloud deployments.

Help and Support During the Evaluation

This guide is intended to provide an overview of the steps required to ensure a successful evaluation of VMware vSphere. It is not meant to substitute for product documentation. Please refer to the online product documentation for vSphere for more detailed information (see below for links). You may also consult the online knowledge base if you have any additional questions. Should you require further assistance, please contact a VMware sales representative or channel partner.

VMware vSphere and vCenter Resources:

- Product documentation:
  http://www.vmware.com/support/pubs/
- Online support:
  http://www.vmware.com/support/
- Support offerings:
  http://www.vmware.com/support/services
- Education services:
  http://mylearn1.vmware.com/mgrreg/index.cfm
- Support knowledge base:
  http://kb.vmware.com
- PowerCLI toolkit community:
  http://communities.vmware.com/community/developer/windows_toolkit
  (or type Get-VIToolkitCommunity within PowerCLI)
- PowerCLI blogs:
  http://blogs.vmware.com/vipowershell

VMware Contact Information

For additional information or to purchase VMware vSphere, VMware's global network of solutions providers is ready to assist you. If you would like to contact VMware directly, you can reach a sales representative at 1-877-4VMWARE (650-475-5000 outside North America) or email sales@vmware.com. When emailing, please include the state, country, and company name from which you are inquiring. You can also visit http://www.vmware.com/vmwarestore/.

Providing Feedback

We appreciate your feedback on the material included in this guide. In particular, we would be grateful for any guidance on the following topics:

- How useful was the information in this guide?
- What other specific topics would you like to see covered?
- Overall, how would you rate this guide?

Please send your feedback to the following address: tmdocfeedback@vmware.com, with "VMware vSphere 5.0 Evaluation Guide" in the subject line. Thank you for your help in making this guide a valuable resource.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com