



Solution
Implementer's
Series™

Operations Guide

Using OneConnect 10Gb Ethernet Adapters with VMware vSphere

- Optimize and manage VMware deployments with OneConnect 10Gb Ethernet adapters and OneCommand Manager

 **EMULEX®**



Table of Contents

Table of Contents	2
Introduction	3
Components.....	3
OneConnect 10GbE Adapters.....	3
OneCommand Manager.....	4
Installation	4
OneConnect Ethernet Driver.....	4
OneCommand Manager Core Kit	4
Updating OneConnect Firmware.....	5
Verifying Firmware Installation	8
Creating Adapter Uplinks and Entities on a VMware Virtual Switch.....	8
VMware vSphere4.1 vSwitch	8
VMware vNetwork Distributed Switch	9
Configuring Emulex OneConnect 10GbE adapter as a vSwitch Uplink.....	9
Configuring Teaming, Failover and VLANs.....	12
Switch Commands to Set Up VLANs and Trunk Ports	18
BIOS Settings/Considerations for High Performance	20
OneConnect Performance: MTU, Heap Size, NetQueues.....	20
Maximum Transfer Unit.....	20
Network Heap Size in ESX Server	22
VMware NetQueues	23
NetQueue Performance Benefits	24
ESX Networking Tips.....	25
Configuring E1000, VMXNET2 and VMXNET 3 Adapters for a Windows Guest.....	25
Conclusion.....	28
References	28



Introduction

Server virtualization is a key technology to lower costs and optimize data center resources. New high-performance, multi-core servers are enabling many more virtual machines (VMs) per physical server and supporting virtualization of more demanding business-critical applications. The result is a requirement for higher I/O capacity that is driving the adoption of 10Gb/s Ethernet (10GbE) networks.

Emulex OneConnect™ adapters are widely available from major server and storage providers as an add-in Network Interface Card (NIC), mezzanine card for blade servers and a built-in LAN on Motherboard (LOM). OneConnect OCe10102-N adapters use the second generation BE2 controller and third-generation OCe11102-N adapters use the BE3 controller. Both are particularly well-suited for deployments with VMware vSphere, the clear market leader in server virtualization for enterprise-class data centers.

To help IT administrators fully leverage server virtualization, the Emulex Technical Marketing group conducted an analysis of how to optimize deployments of OneConnect 10GbE adapters with VMware vSphere 4.1. The result is this guide that provides information on the following topics:

- Installation of Emulex OneConnect 10GbE adapters with VMware vSphere 4.1
- Configuring a VMware virtual switch (vSwitch) uplink
- Configuring teaming, failover and VLANs
- BIOS and OneConnect adapter settings for high performance
- Tuning with MTU, heap size and VMware NetQueue
- Configuring E1000, VMXNET2 and VMXNET 3 Adapter for Microsoft® Windows® guests
- NetQueue performance comparisons
- OneConnect performance with Windows VM guests

Components

This deployment guide is based on the following components:

- Multi-core server with MSI-X and PCI-Express 2.0
- VMware vSphere 4.1
- Emulex OneConnect OCe10102-N or OCe11102-N 10GbE adapter
- Emulex OneCommand™ Manager

OneConnect 10GbE Adapters

The Emulex OneConnect product family provides optimized connectivity to 10GbE networks and includes options to support protocol offload for iSCSI and Fibre Channel over Ethernet (FCoE).



OneCommand Manager

OneCommand Manager provides centralized management of Emulex Fibre Channel Host Bus Adapters (HBAs) and OneConnect 10GbE network adapters, including support for iSCSI and FCoE protocols. OneCommand Manager provides powerful adapter provisioning and diagnostic capabilities to increase administration efficiency in a VMware environment.

Installation

OneConnect Ethernet Driver

Access the Emulex website (link below) to check for updates to the OneConnect driver that are included with the ESX 4.1 distribution: <http://www.emulex.com/downloads/emulex/cnas-and-hbas/drivers/vmware.html>

Installation instructions will be available from the same link.

Management Applications for VMware vSphere 4.1

UCNA and HBAs Application Kit

The Application core kit must be downloaded from the Emulex web site (link below) and installed to run specific commands to update firmware and boot code.

<http://www.emulex.com/downloads/emulex/vmware/vsphere-41/management.html>

To install the Application kit

- 1) Log in as root
- 2) Copy the elxocmcore kit contents to a directory on the ESX server (for example, /tmp)
- 3) `cd /tmp`
- 4) `rpm -U elxocmcore-esx41-<kit version>.x86_64.rpm`
- 5) `rpm -qa | grep elx` (to confirm the elxocmcore kit loaded successfully)
- 6) `cd /usr/sbin/hbanyware` (to run OneCommand Manager CLI commands)

To remotely manage adapters on an ESX server, the TCP/IP port used by OneCommand Manager must be opened to allow TCP traffic to pass through the ESX server firewall. By default port number 23333 is used. The port number must be the same for all hosts that are being managed remotely. If necessary, the port number can be changed from the OneCommand Manager Discovery tab.

Using the default port number, enter the following to open a port for inbound and outbound traffic:

```
esxcfg-firewall --openPort 23333,tcp,in,OneCommand
esxcfg-firewall --openPort 23333,tcp,out,OneCommand
```



To verify the correct port is open, enter:

```
esxcfg-firewall -q
```

Refer to the following links below for more information on how to configure ESX 4.1 firewall settings and OneCommand Manager for out-of band management:

http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_server_config.pdf

http://www-dl.emulex.com/support/utilities/onecommand/50/corekit_user_manual.pdf

Updating OneConnect Firmware

Important: Update the firmware and boot code after installing a new driver.

New firmware and boot code can be upgraded using OneCommand Manager CLI or from a remote host also running OneCommand Manager.

Download firmware and boot code from the following Emulex webpage:

<http://www.emulex.com/downloads.html>

Updating OneConnect Firmware from Local Server

Look up the MAC address for the adapter by entering:

```
./hbacmd listhba pt=NIC
```

Note: Because OneConnect adapters are capable of supporting FCoE, commands reference HBAs.



```
[root@denver hpanyware]# ./hbaconf listhba pt=NIC

Manageable HBA List

Permanent MAC : 00-00-c9-91-4f-3a
Current MAC   : 00-00-c9-91-4f-3a
Flags        : 80000700
Host Name    : denver
Mfg         : Emulex Corporation
Serial No.   : BT93744532
Port Number  : 0
Mode        : Initiator
PCI Function : 0
Port Type   : NIC
Model       : OneConnect OCe10100

Permanent MAC : 00-00-c9-91-4f-3c
Current MAC   : 00-00-c9-91-4f-3c
Flags        : 80000700
Host Name    : denver
Mfg         : Emulex Corporation
Serial No.   : BT93744532
Port Number  : 1
Mode        : Initiator
PCI Function : 1
Port Type   : NIC
Model       : OneConnect OCe10100

[root@denver hpanyware]#
```

Figure 1. Adapter list

Download the new firmware image to the adapter by entering:

```
./hbaconf download <MAC address> /<directory>/<firmware file name>
```

After the download completes, reboot the ESX server (see Figure 2).

Updating OneConnect Firmware from a Remote Computer

Prerequisite: Make sure OneCommand Manager is installed on the remote computer and the Full Management option is selected.

From the initial OneCommand Manager screen, execute the following steps:

- 1) Click on the Discovery Tab
- 2) Select Add Remote TCP/IP Host
- 3) Enter host name or IP address as shown in Figure 2

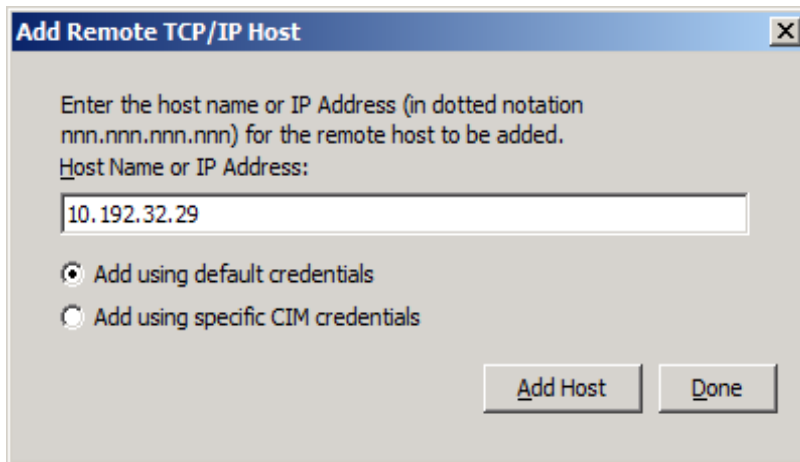


Figure 2. Enter host name or IP address to add remote TCP/IP host

After the discovery process is complete, OneCommand Manager displays a list of servers and adapters:

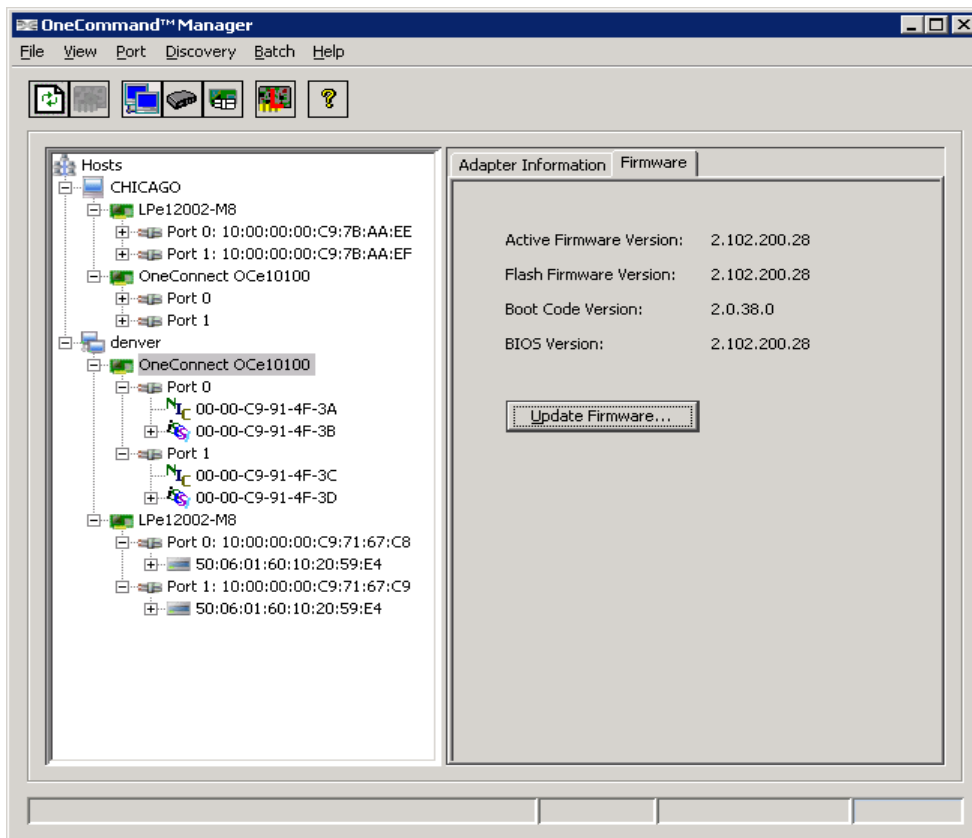


Figure 3. List of servers and adapters

To update firmware, execute the following steps:



- 1) Select adapter
- 2) Click on "Update Firmware"
- 3) Select firmware file

When the firmware download is complete, reboot the ESX server.

Verifying Firmware Installation

To confirm firmware installation, enter the following:

```
# /usr/sbin/hbanywhere/hbacmd hbaattribs <MAC Address>
```

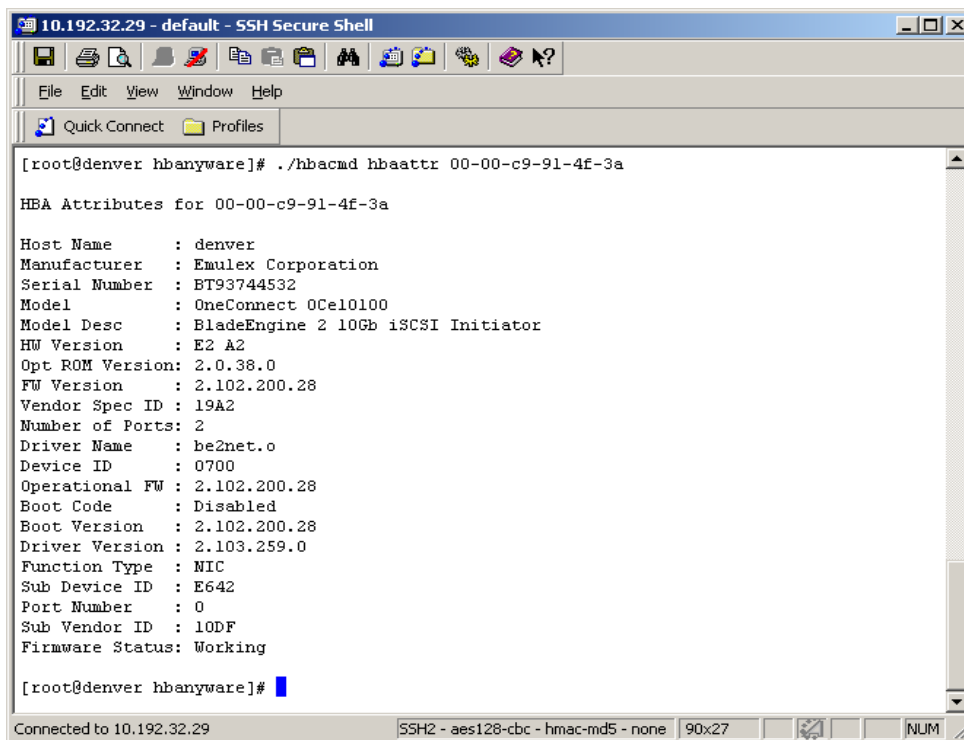


Figure 4. Firmware version and status for MAC address 00-00-c9-91-4f-3a

Creating Adapter Uplinks and Entities on a VMware Virtual Switch

The next step is to create physical uplinks on a VMware Virtual Switch (vSwitch) as described below.

VMware vSphere4.1 vSwitch

The core of a virtual network is the vSwitch which is essentially a virtualized Layer 2 switch that connects VMs to a virtual network and allows access to a physical network. Physical Network Interface Cards (pNICs) are used as vSwitch uplinks to a physical network. vSwitches can perform



complex operations including support for virtual LANs (VLANs), Security Policy Exceptions, Traffic Shaping and NIC Teaming functions. A virtual network interface card (vNIC) is connected to the VM side of a vSwitch and connects a VM to the virtual network. Each vNIC has its own MAC address and can be assigned one or more IP addresses, similar to a physical NIC in a non-virtualized computer.

VMware vNetwork Distributed Switch

The Cisco 1000V Series switch is a VMware vNetwork Distributed Switch (VDS) and a full Layer 2 virtual switch that runs inside the ESX hypervisor and uses the Cisco NX-OS operating system. It has security and mechanisms to support 802.3ad, 802.1q and other protocols used in today's switching networks. The following link provides additional details on VMware VDS and Cisco Nexus 1000V Series switches:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/data_sheet_c78-492971.html

Configuring Emulex OneConnect 10GbE adapter as a vSwitch Uplink

Select the vSphere Client Configuration tab to show available Network Adapters (Figure 5).

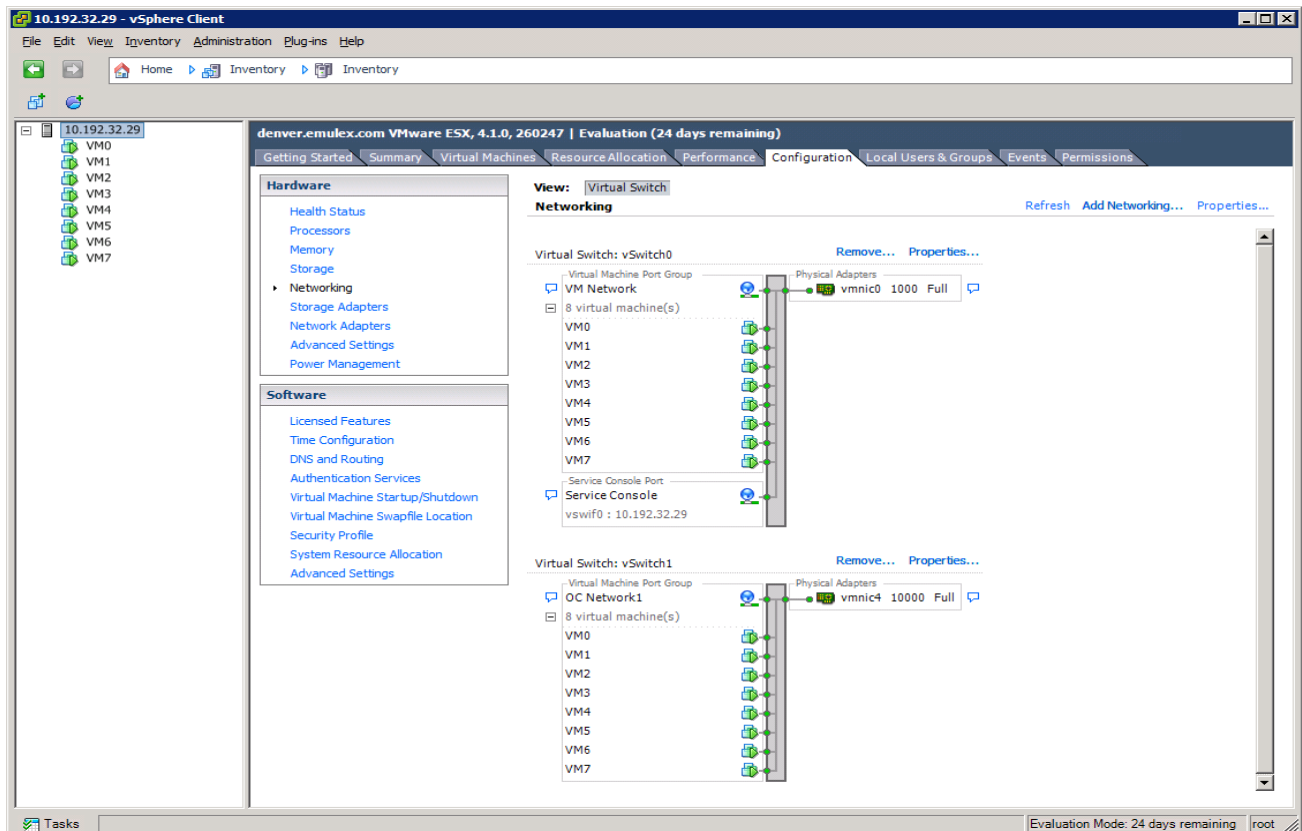


Figure 5. vSphere Client view of vSwitch1



In this example, vmnic0 was previously configured to vSwitch1. vmnic1 is a OneConnect adapter and is presently not configured to any vSwitches. Eight VMs (VM0 – VM7) are in OC Network1 and configured to handle VM network traffic. To add vmnic1 to vSwitch1, complete the following from the vSphere Client:

- 1) Click on vSwitch1 Properties
- 2) Select Network Adapters (Figure 6)




Figure 6. vSwitch1 properties

- 3) Click on Add. An Unclaimed Adapters list is shown (Figure 7):

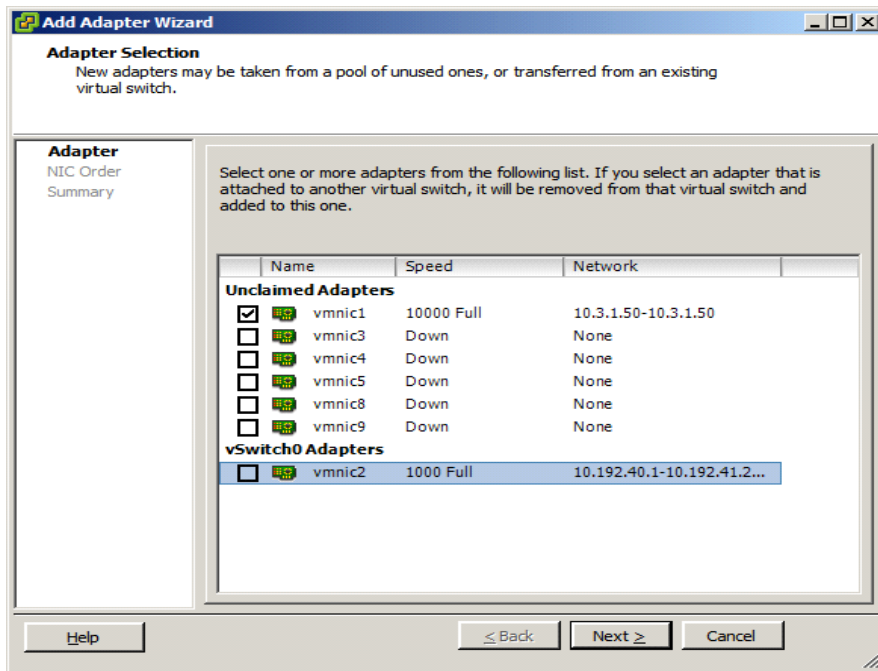


Figure 7. Unclaimed Adapter List

- 4) Select vmnic1 from the Unclaimed Adapter List. vmnic1 now appears in the list of Active Adapters and is configured to vSwitch1 as shown in Figure 8 and Figure 9.

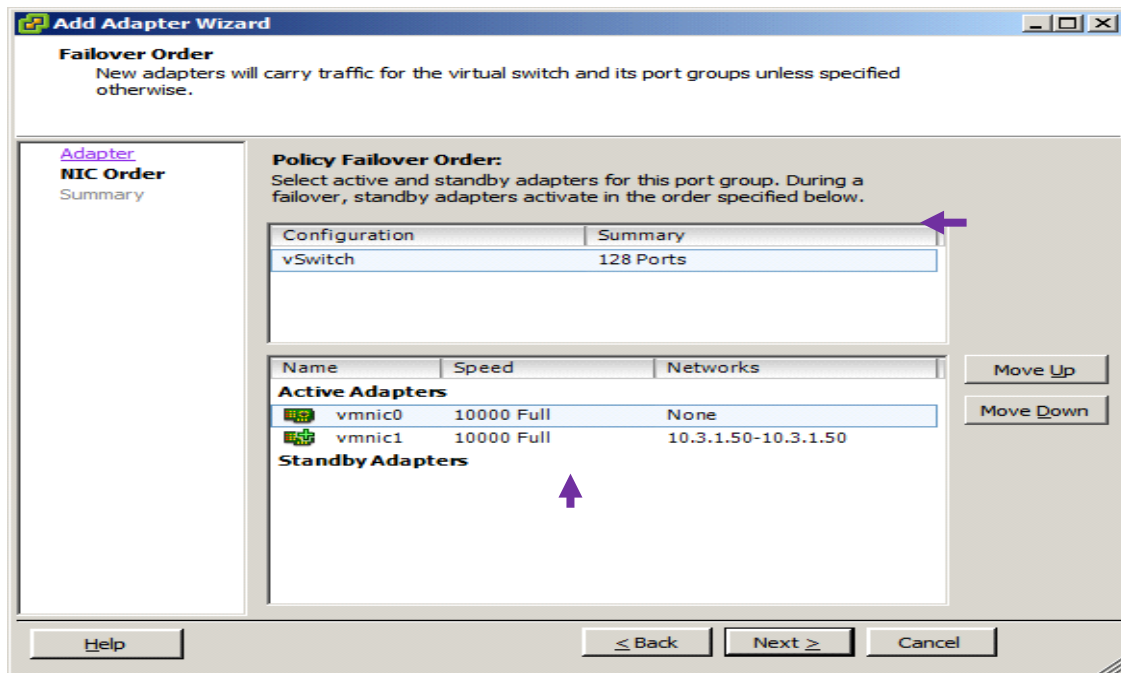


Figure 8. vmnic1 as active adapter configured to vSwitch1

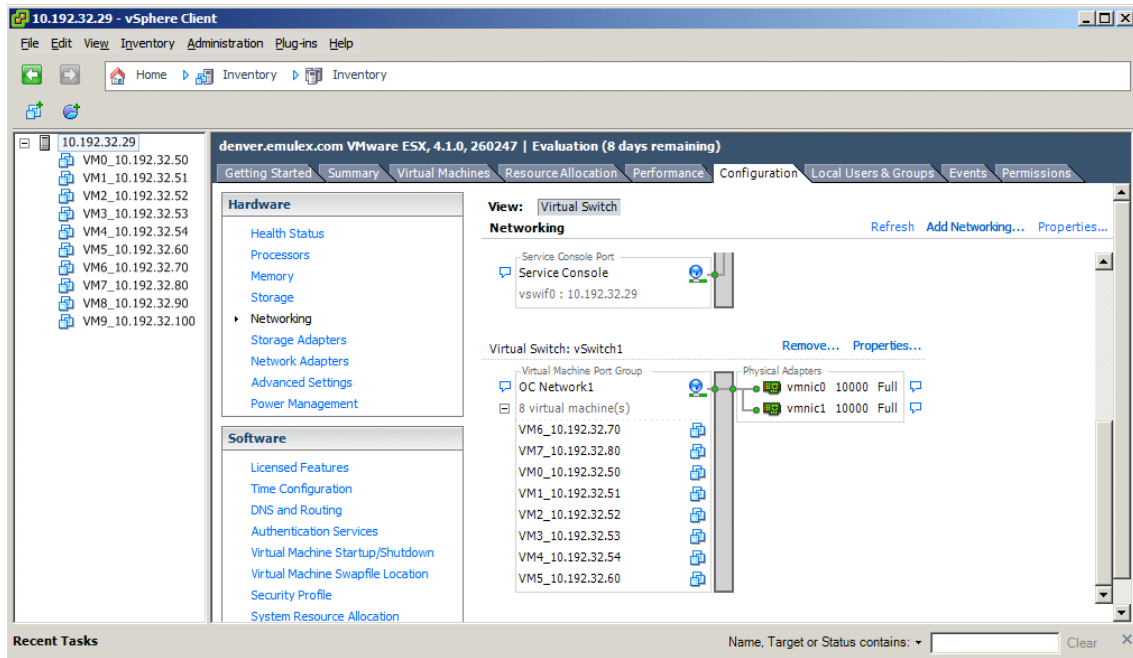


Figure 9. vSphere Client view of vmnic1

Configuring Teaming, Failover and VLANs

NIC teaming supports redundant NIC adapters that are used as vSwitch uplinks (VMNICs) to a physical switch. Configuring redundant uplinks provides high availability and can increase performance by leveraging multiple physical NIC ports and different network paths. NIC teaming offers several configuration options which can be implemented per-vSwitch or per Port Group:

- Active/Standby.
- Active/Active with load balancing based on VM Port-ID.
- Active/Active with load balancing based on the VM MAC-address hash.
- Active/Active with load balancing based on the hash of the source and destination IP address.

To configure Emulex adapters for teaming and failover, complete the following steps using VMware vCenter 4.1:

- 1) Click on vSwitch1 Properties
- 2) Click on Ports

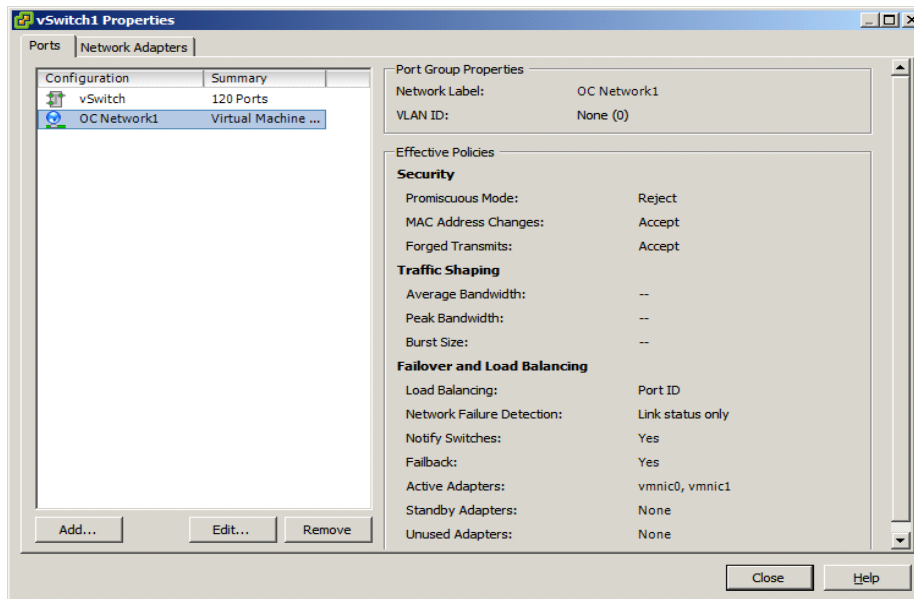


Figure 10. Configure OneConnect adapter for NIC teaming and failover

- 3) Select OC Network1
- 4) Click on “Edit” (Figure 10)

Click on the “NIC Teaming Tab” to access Policy Exceptions. Failover can be changed by clicking on “Override vSwitch Failover Order” to specify Active, Standby or Unused, selecting an adapter and using the Move Up/ Down buttons.

Load Balancing

ESX offers three different methods for load balancing:

- Route based on the originating port ID (this is the default). Uplink based on the virtual port where the traffic entered the vSwitch (also known as vSwitch port-based load balancing).
- Route based on source MAC hash. Uplink based on a hash of the source Ethernet address.
- Route based on IP hash. Uplink based on a hash of the source and destination IP address.

The link below provides additional information on configuring NIC Teaming in ESX 4.1:\

http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_server_config.pdf

The following example uses routing based on source MAC hash value, with VM outbound traffic mapped to a pNIC based on the vNIC MAC address. Other settings include Network Failover Detection, Switch Notification and Failback.

To change adapter Failover Order, check Override vSwitch Failover Order. Configure the adapter as an Active, Standby or Unused adapter by clicking on the Move up or Move down button (Figure 11).

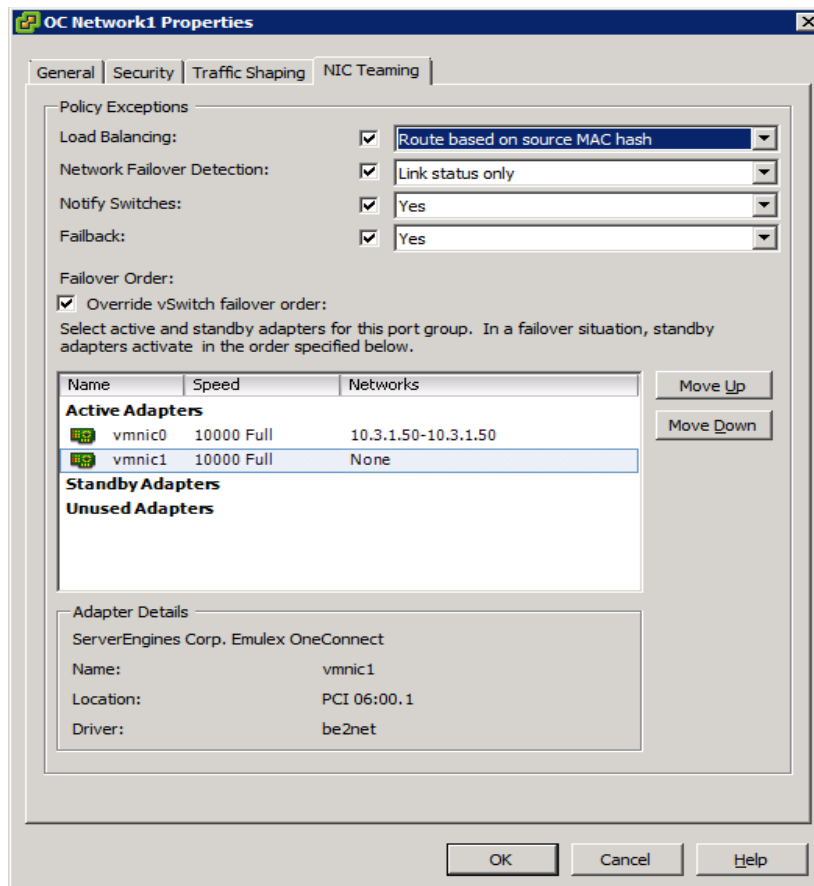


Figure 11. Setting Override vSwitch Failover Order

Virtual Switch Tagging (VST) allows the vSwitch to perform 802.1q tagging. This allows binding a single network interface to multiple VLANs and broadcast domains using a 32-bit field that is injected into the frame of each packet. The tag is inspected by a switch and then stripped out of the frame before it reaches its destination. The switch routes the packet to the appropriate broadcast domain. This ensures that traffic remains isolated and secured from the rest of the network.

VLAN ID can be assigned with the following steps. From the VMware vSphere client, click on the Configuration tab and select Networking (Figure 12).

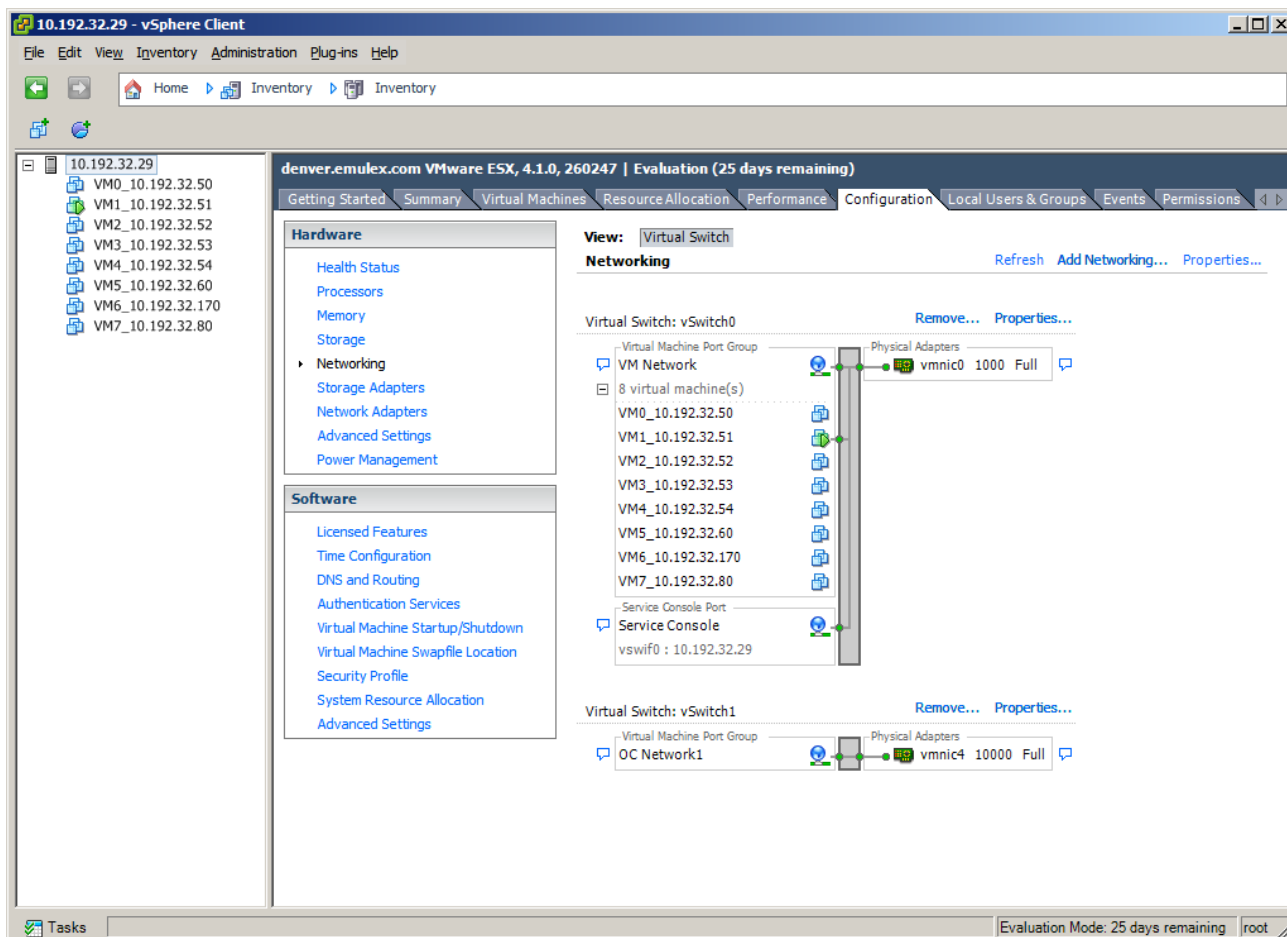


Figure 12. Select Networking from the Configurations Tab



Click on the Ports tab and select Virtual Machine Port Group (OCNetwork1 in Figure 13):

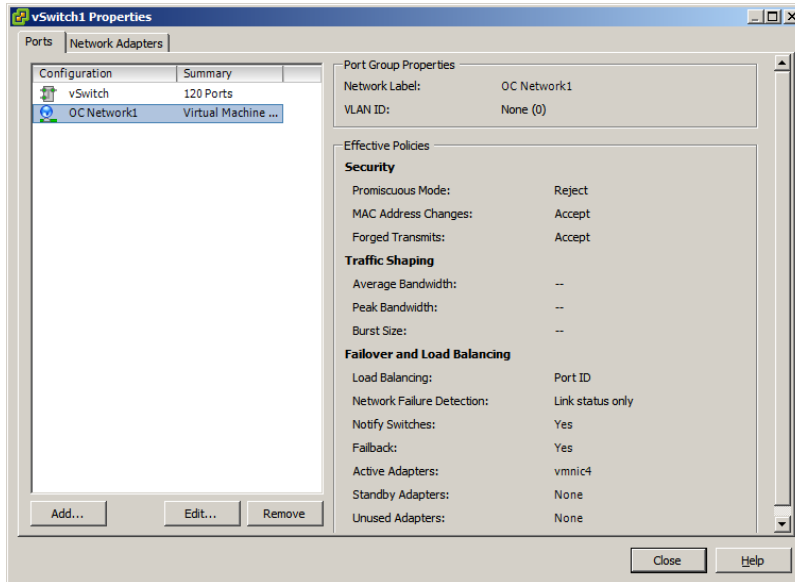


Figure 13. Select Virtual Machine Port Group

Click on Edit and enter the VLAN ID:

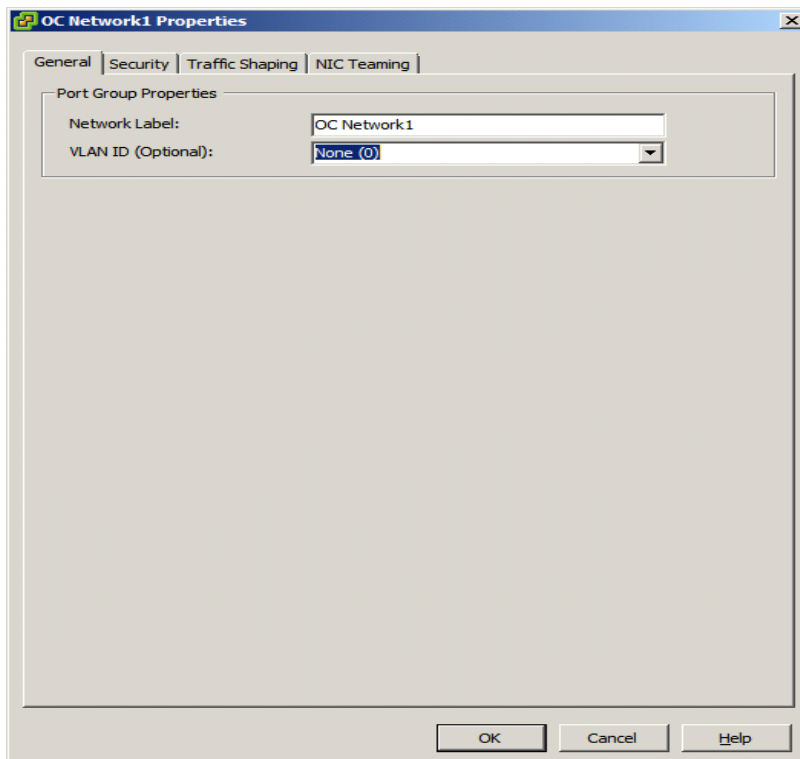


Figure 14. Assign VLAN ID



The vSwitch1 Properties, Port Group Properties displays VLAN ID 1001 (Figure 15).

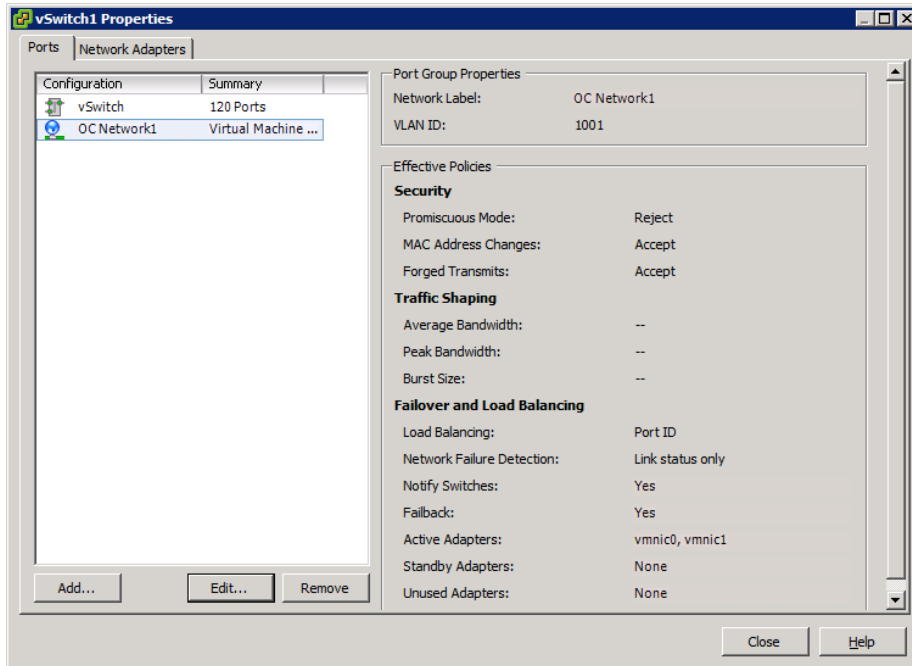


Figure 15. View VLAN ID 1001

With this setup, virtual machines VM1 – VM7 are now in VLAN 1001 (Figure 16).

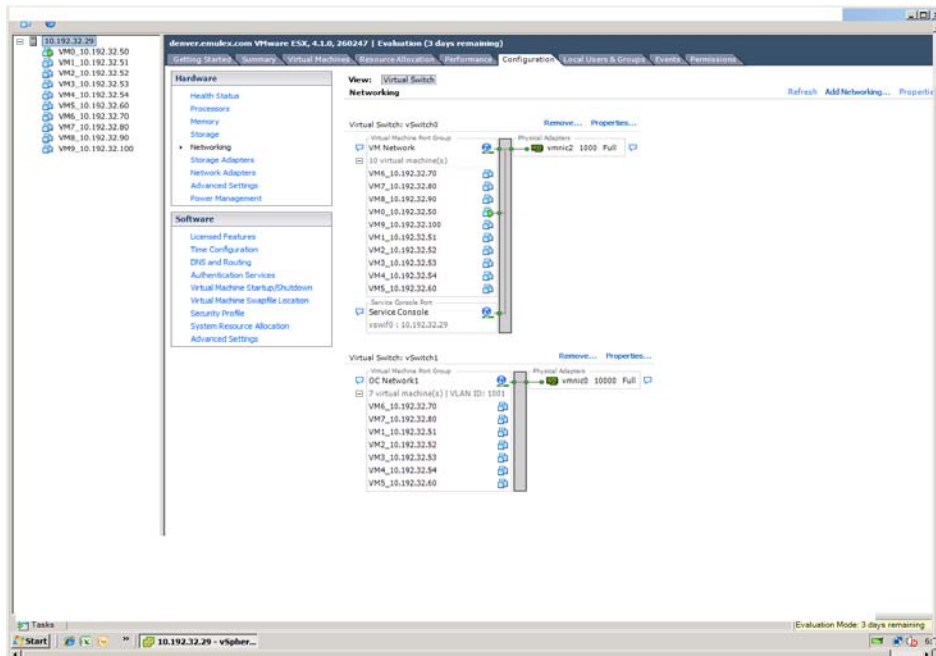


Figure 16. OC Network1 has VLAN ID 1001



VLANS across Virtual and Physical Network Environments

Figure 17 shows an ESX configuration with three vSwitches having unique VLAN IDs 1001, 1002 and 1003.

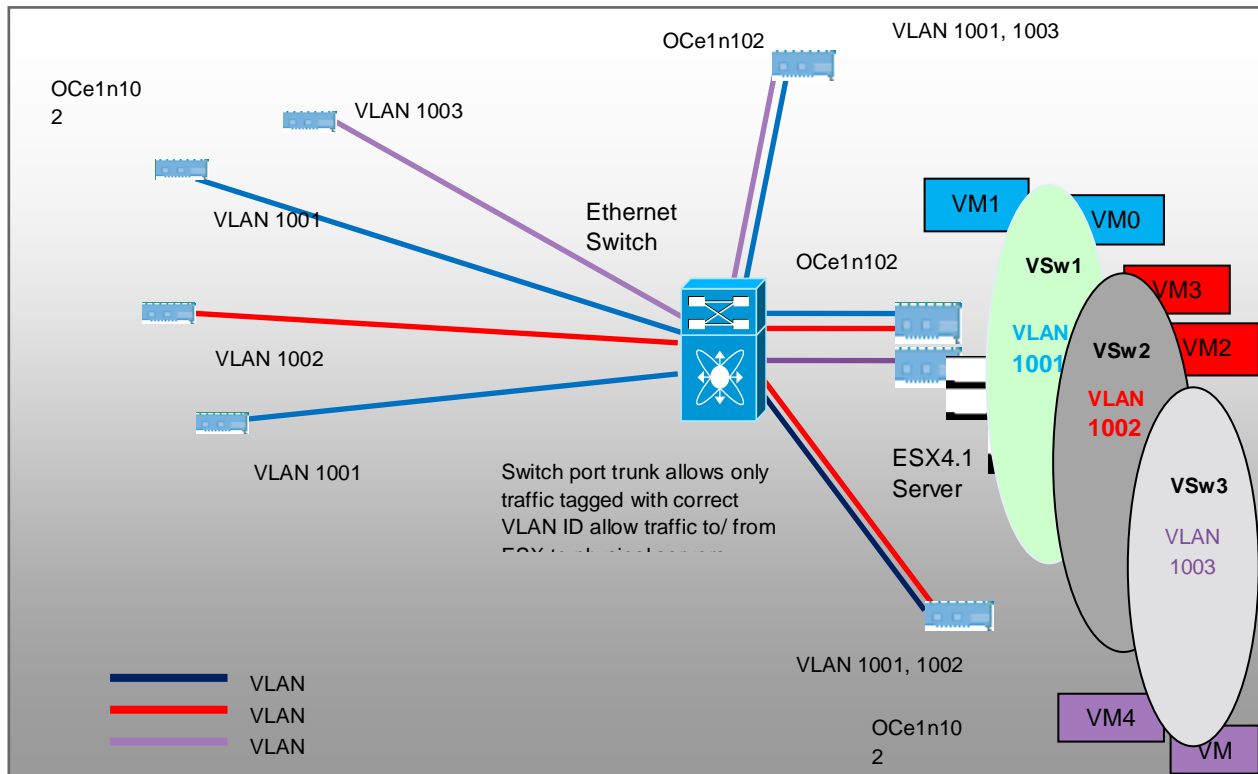


Figure 17. ESX configuration with three vSwitches

Only virtual hosts connected to a vSwitch in VLAN 1001 are allowed to communicate with physical servers in VLAN 1001. VMs connected to the same vSwitch can communicate with other VMs on the same vSwitch. For example, VM3 can communicate with VM2 through the vSwitch. VMs on a different vSwitch can only communicate through the physical switch and both vSwitches must be configured with the same VLAN ID.

Switch Commands to Set Up VLANs and Trunk Ports

The first step is creation of VLANs. The switch commands to create a range of VLANs are as follows:

```
# configure terminal  
# vlan {vlan-id | vlan-range}  
# no shut
```

The following example shows configuration of a switch to make VLANs 1001-1003 active. Confirm by entering the “show vlan” command.



```

File Edit Setup Control Window Help
sjtechmkt#
sjtechmkt# configure terminal
sjtechmkt(config)# vlan 1001,1002,1003
sjtechmkt(config-vlan)# show vlan

```

VLAN	Name	Status	Ports
1	default	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/9, Eth1/10 Eth1/13, Eth1/14, Eth1/15 Eth1/16, Eth1/17, Eth1/18 Eth1/19, Eth1/20, Eth1/21 Eth1/22, Eth1/23, Eth1/24 Eth1/25, Eth1/26, Eth1/27 Eth1/28, Eth1/29, Eth1/30 Eth1/31, Eth1/32, Eth1/33 Eth1/34, Eth1/35, Eth1/36 Eth1/37, Eth1/38, Eth1/39 Eth1/40
1001	VLAN1001	active	
1002	VLAN1002	active	
1003	VLAN1003	active	

Figure 18. Configure switch to make VLANs active

The next step is configuration of Ethernet ports for VLANs. The commands are as follows:

```

# configure terminal
# interface Ethernet 1/7
# switchport mode trunk

```

Figure 19 shows configuration Ethernet ports 1/7, 1/8 and 1/9 for VLANs 1001-1003.

```

File Edit Setup Control Window Help
sjtechmkt# configure terminal
sjtechmkt(config)# interface Ethernet1/7
sjtechmkt(config-if)# switch mode trunk
sjtechmkt(config-if)# switch trunk allowed vlan 1001
sjtechmkt(config-if)# no shut
sjtechmkt(config-if)# interface Ethernet1/8
sjtechmkt(config-if)# switch mode trunk
sjtechmkt(config-if)# switch trunk allowed vlan 1002
sjtechmkt(config-if)# no shut
sjtechmkt(config-if)# interface Ethernet1/9
sjtechmkt(config-if)# switch mode trunk
sjtechmkt(config-if)# switch trunk allowed vlan 1003
sjtechmkt(config-if)# no shut
sjtechmkt(config-if)# exit
sjtechmkt(config)#

```

Figure 19. Configuring Ethernet ports for VLANs



BIOS Settings/Considerations for High Performance

In most circumstances the default hardware basic input/output system (BIOS) settings are not set for optimal performance and will benefit from modifications. ESX supports Enhanced Intel SpeedStep and Enhanced AMD PowerNOW CPU power management technologies that save power when a host is not fully utilized, but these settings can reduce performance. Consider disabling these options to achieve the best overall system performance.

Because there are many server models and hardware types, it is difficult to determine the best options for every server. However, the following recommendations should be considered:

- Confirm that the most recent version of the system BIOS is being used.
- Confirm that the BIOS is set to enable all populated sockets and all cores in each socket.
- Enable “Turbo Mode” for processors that support it.
- Enable hyper-threading in the BIOS.
- Some NUMA-capable systems provide an option to disable NUMA by enabling node interleaving. In most cases best performance is achieved by disabling node interleaving.
- Confirm that hardware-assisted virtualization features (VT-x, AMD-V, EPT, RVI) are enabled in the BIOS.
- Disable C1E halt state in the BIOS. (See note below regarding performance considerations versus power considerations.)

NOTE: After these changes are made some systems may need a complete power down before the changes take effect. See <http://communities.vmware.com/docs/DOC-8978> (“Trouble-shooting Intel VT-x Issues”) for details.

- Disable any other power-saving mode in the BIOS.
- Disable any unneeded devices from the BIOS, such as serial and USB ports.

For more information regarding configuring hardware BIOS Settings for VMware platforms, see Performance Best Practices for VMware vSphere4.1: http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.1.pdf

OneConnect Performance: MTU, Heap Size, NetQueues

Performance of OneConnect 10GbE adapters can be optimized with changes to the Maximum Transfer Unit (MTU), heap size and NetQueues, as described below.

Maximum Transfer Unit

Increasing the MTU to support jumbo frames can significantly enhance throughput and reduce CPU utilization for large file transfers by efficiently sending larger payloads per packet. The following example shows how to increase the MTU size of the vSwitch and the virtual Ethernet adapter in a guest OS.



As shown in Figure 20, to check current MTU on all vSwitches, enter:

```
# esxcfg-vswitch -l
```

```
[root@denver ~]# esxcfg-vswitch -l
Switch Name      Num Ports  Used Ports  Configured Ports  MTU    Uplinks
vSwitch0        128        6           128               1500   vmnic2

  PortGroup Name  VLAN ID  Used Ports  Uplinks
  VM Network      0        3           vmnic2
  Service Console 0        1           vmnic2

Switch Name      Num Ports  Used Ports  Configured Ports  MTU    Uplinks
vSwitch1        128        5           128               1500   vmnic0

  PortGroup Name  VLAN ID  Used Ports  Uplinks
  OC Network1    1001    3           vmnic0
```

Figure 20. Check current MTU on all vSwitches

As shown in Figure 21, to change vSwitch1 MTU to support jumbo frames up to 9000 blocks, enter:

```
# esxcfg-vswitch<N> -m 9000
```

```
[root@denver ~]# esxcfg-vswitch vSwitch1 -m 9000
[root@denver ~]# esxcfg-vswitch -l
Switch Name      Num Ports  Used Ports  Configured Ports  MTU    Uplinks
vSwitch0        128        6           128               1500   vmnic2

  PortGroup Name  VLAN ID  Used Ports  Uplinks
  VM Network      0        3           vmnic2
  Service Console 0        1           vmnic2

Switch Name      Num Ports  Used Ports  Configured Ports  MTU    Uplinks
vSwitch1        128        5           128               9000   vmnic0

  PortGroup Name  VLAN ID  Used Ports  Uplinks
  OC Network1    1001    3           vmnic0
```

Figure 21. Change MTU on a virtual switch

The MTU value can be changed in a Windows Guest OS to use jumbo packets by configuring the adapter properties as shown in Figure 22.

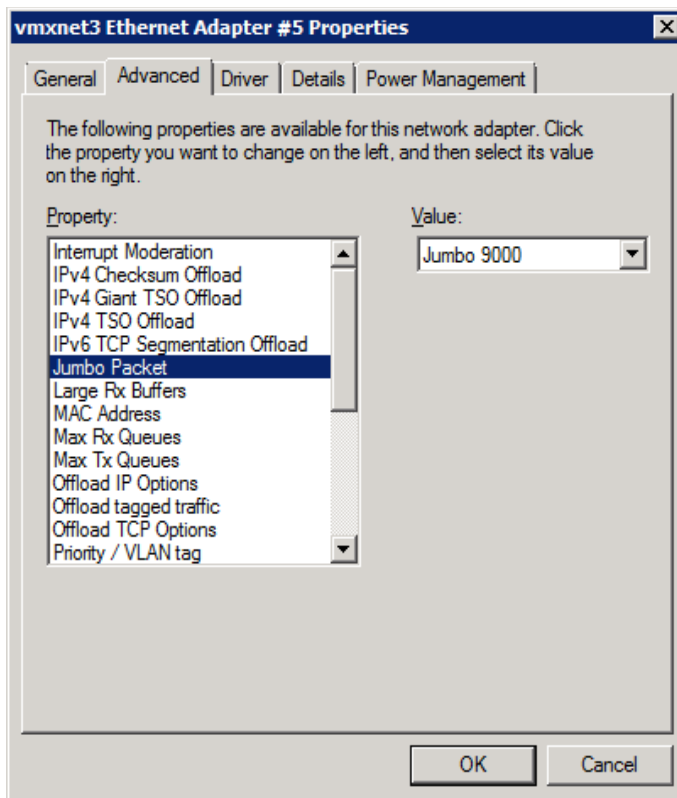


Figure 22. Change MTU in a Windows Guest OS

MTU size can also be changed in a Linux Guest OS. To set the MTU to support jumbo frames up to 9000 blocks, enter:

```
# ifconfig eth<N> mtu 9000
```

Network Heap Size in ESX Server

By default, the ESX server network stack allocates 64MB of buffers to handle network data. If the network load requires more than 64MB of memory, the OneConnect driver cannot allocate more memory automatically to handle the additional load. When this happens, the driver logs messages in the file `/proc/vmware/log` on the ESX server indicating that the `alloc_skb()` call failed. This situation can negatively impact network performance.

To read the current size of the network heap on the ESX server, enter:

```
# esxcfg-advcfg -j netPktHeapMaxSize
netPktHeapMaxSize = 0
```



The value returned is 0. The following example shows changing the heap size to 128MB and then verifying the change was made (new value takes effect after a reboot):

```
# esxcfg-advcfg -k 128 netPktHeapMaxSize  
# esxcfg-advcfg -j netPktHeapMaxSize  
netPktHeapMaxSize = 128
```

VMware NetQueues

NetQueue is a VMware performance technology that significantly improves performance in 10GbE virtualized environments. NetQueue requires MSI-X support on the server platform, however support may be limited to specific system models. NetQueue is enabled by default in ESX 4.0 and ESX4.1. As a best practice, you should verify this on the ESX server or with vSphere Client.

Pinning a network queue interrupt to a CPU can reduce interrupt latency, increase VM network performance and reduce CPU utilization.

To verify NetQueue is enabled in an ESX vSphere 4.1 Server, enter the following and then reboot the server:

```
# esxcfg-advcfg -j netNetqueueEnabled  
  
netNetqueueEnabled = TRUE
```

To verify NetQueue is enabled in the vSphere Client, do the following (as shown in Figures 23 and 24):

- 1) Click on Configuration
- 2) Select Advanced Setting
- 3) Select VMkernel
- 4) Confirm box is checked: VMKernel.boot.netNETqueueEnabled

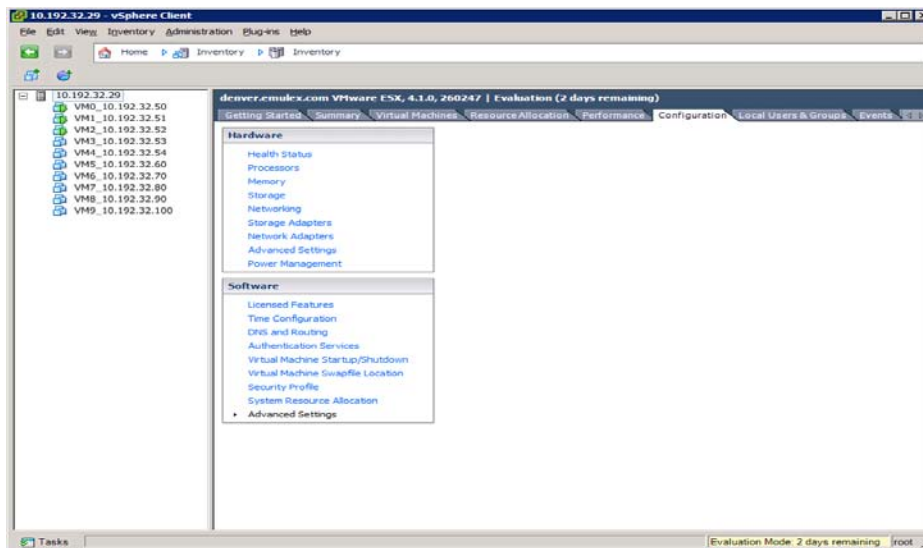


Figure 23. Verify NetQueue is enabled in the vSphere Client (steps 1-3)

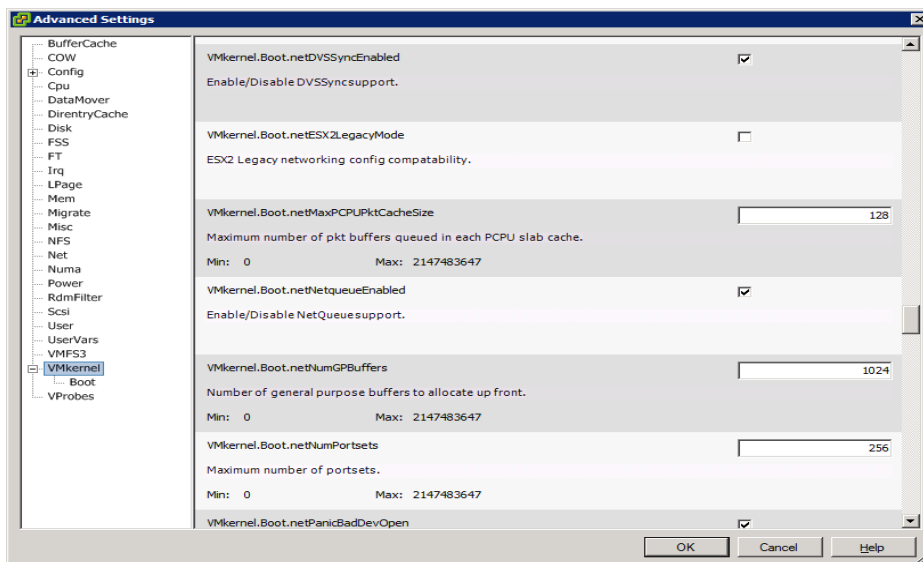


Figure 24. Verify NetQueue is enabled in the vSphere Client (steps 4 and 5)

NetQueue Performance Benefits

When NetQueue is enabled in ESX, the Emulex OneConnect adapter provides up to eight individual receive queues on each port and offloads layer 2 packet sorting and classification functions from the hypervisor. When Ethernet packets are routed to a OneConnect network controller, they are placed into specific queues, allowing processing to be scaled to multiple CPUs and significantly improving receive-side networking performance.



Network traffic coming in on a specific queue will trigger a MSI-X interrupt for one or more CPUs to complete the application task. By default, VMs will share all available cores in the system and interrupts generated by network queues are handled automatically in ESX.

ESX will do an excellent job balancing CPU loads and supporting multiple devices on most multi-processor systems. If additional performance tuning is required, adjusting affinity settings and pinning IRQ vectors can improve performance but may require some experimentation and will affect virtualization features such as VMotion and Dynamic Resource Scheduling.

ESX Networking Tips

- CPU utilization plays a significant role in network throughput. To process higher levels of throughput, more CPU resources are needed. The effect of CPU resource availability on the network throughput of virtualized applications is even more significant. Because insufficient CPU resources will reduce maximum throughput, it is important to monitor CPU utilization of high-throughput workloads.
- Use separate vSwitches, each connected to its own physical network adapter, to avoid contention between the ESX service console, the VMkernel and VMs (especially VMs running heavy networking workloads).
- To establish a network connection between two VMs that reside on the same ESX system, connect both VMs to the same vSwitch. If the VMs are connected to different vSwitches, traffic will go through the external physical network and incur unnecessary CPU and network overhead.
- The VMkernel network device driver defaults to auto-negotiation of speed and duplex settings. These settings will work correctly with network switches that are also set to auto-negotiate. If the switch is configured for specific speed and duplex settings, force the network driver to use the same speed and duplex settings (for Gigabit links, network settings should be set to auto-negotiate and not forced).
- The VMXNET2, VMXNET3 and E1000 devices support jumbo frames for better performance. **To enable jumbo frames**, set the **MTU** size to **9000** both in the **guest** and in the **vSwitch configuration**. The physical NICs at both ends and all the intermediate hops/routers/switches must also support jumbo frames.

For more information, see VMware Performance Best Practices:

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.1.pdf

Configuring E1000, VMXNET2 and VMXNET 3 Adapters for a Windows Guest

For best performance, VMware recommends using the VMXNET3 network adapter. This requires virtual hardware version 7 and installation of VMware tools in the guest OS. To add a VMXNET3 adapter, follow the instructions below (shown in Figures 25, 26 and 27):

- 1) Right click on the VM



- 2) Select Edit Settings
- 3) Click Add
- 4) Select Ethernet Adapter
- 5) Click Next
- 6) Select Adapter Type
- 7) Click Next

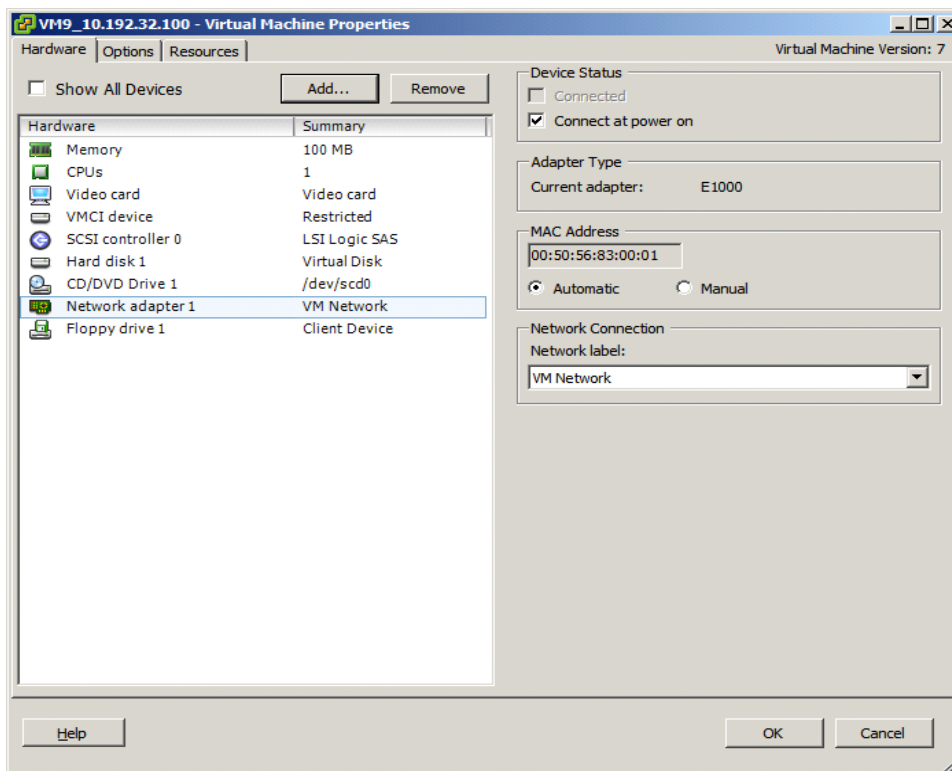


Figure 25. Add VMXNET3 adapter (steps 1-3)

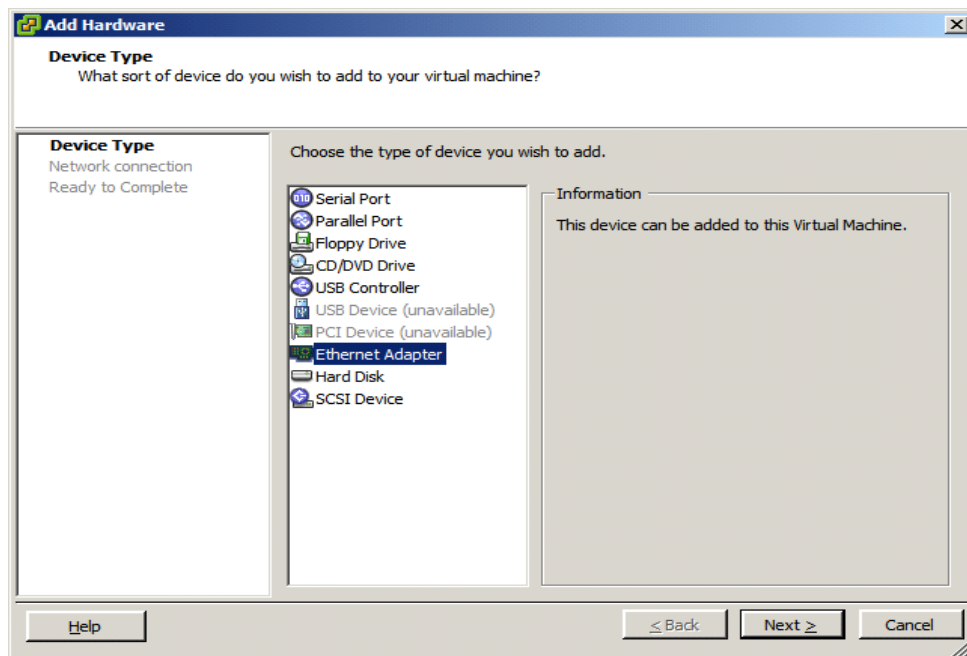


Figure 26. Add VMXNET3 adapter (steps 4 and 5)

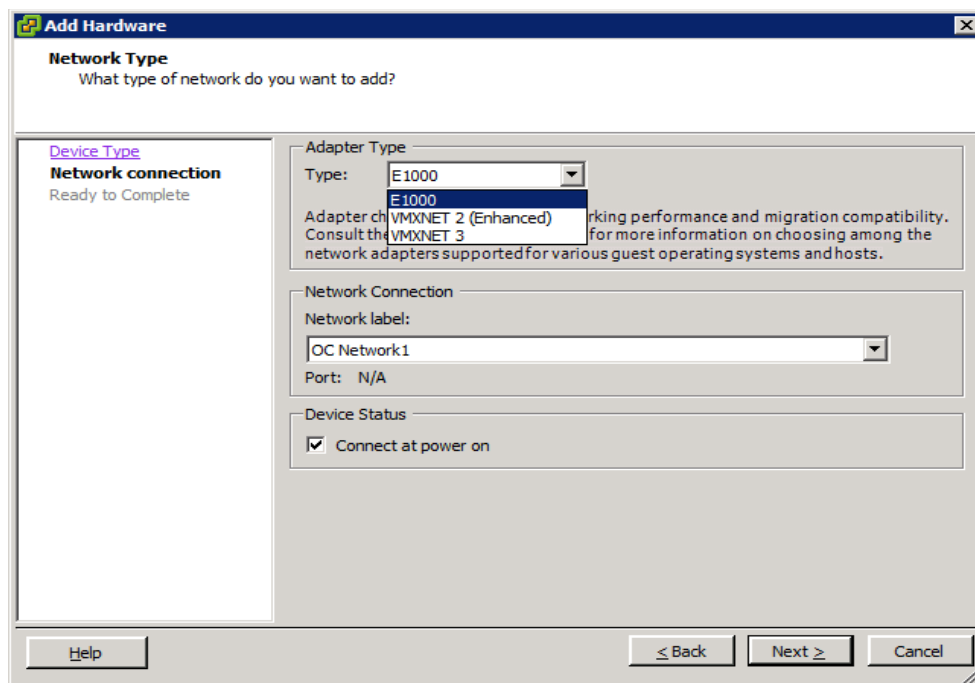


Figure 27. Add VMXNET3 adapter (steps 6 and 7)



Conclusion

IT organizations can use VMware vSphere 4.1 to create a virtualized infrastructure of server resources that are dynamically provisioned and optimized. By consolidating servers, dramatic savings can be realized for capital equipment and operating expenses for power, cooling, maintenance and system management.

Emulex OneConnect 10GbE adapters deliver high-performance, high-bandwidth network connectivity for virtualized servers. This deployment guide provided information that can help ensure a VMware/Emulex solution is installed with parameters and other features that are set to deliver best performance. Specific recommendations including enabling NetQueue and use of jumbo MTU for best throughput and performance.

References

Emulex VMware 4.1 User Manual

<http://www-dl.emulex.com/support/vmware/8217822/vmware41.pdf>

Emulex OneCommand Manager Command Line Interface

http://www-dl.emulex.com/support/utilities/onecommand/50/corekit_user_manual.pdf

VMware vSphere TM: The CPU Scheduler in VMware ESX 4.1

http://www.vmware.com/files/pdf/techpaper/VMW_vSphere41_cpu_schedule_ESX.pdf

VMware Compatibility I/O Guide (September 1, 2010)

<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=software>

VMware Virtual Network Components

<http://www.vmware.com/technical-resources/virtual-networking/networking-basics.html>

What's New in VMware vSphere 4: Virtual Networking


<http://www.vmware.com/files/pdf/vsphere-whatsnew-networking-wp.pdf>

Performance Best Practices for VMware vSphere 4.1

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.1.pdf

ESX Configuration Guide ESX4.1 vCenter Server 4.1

http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_server_config.pdf



www.emulex.com

World Headquarters 3333 Susan Street, Costa Mesa, California 92626 +1 714 662 5600
Bangalore, India +91 80 40156789 | Beijing, China +86 10 68499547
Dublin, Ireland +35 3 (0)1 652 1700 | Munich, Germany +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | Tokyo, Japan +81 3 5322 1348
Wokingham, United Kingdom +44 (0) 118 977 2929