



Using ESVA iSCSI-host Storage Systems in VMware vSphere 4

Application Note

Abstract

This application note explains the configure details of using Infortrend ESVA iSCSI-host storage systems in VMware vSphere 4 to deliver a virtualized data center featuring high efficiency, flexibility and availability.

VMware Virtualization

The concept of virtualization originated in 1960s but was not applied to the x86 architecture until 1990s. Since 1980s, x86 servers have been widely adopted in IT environment because they are much cheaper than mainframe computers. This distributed system of computing reduces TCO but gives birth to other challenges, such as low infrastructure utilization, increasing physical infrastructure costs, increasing IT management costs, insufficient failover and disaster protection, and etc. Virtualization is found an effective way to deal with these challenges.

In VMware's virtualization technology, ESX Server is the foundation of virtualized environments.

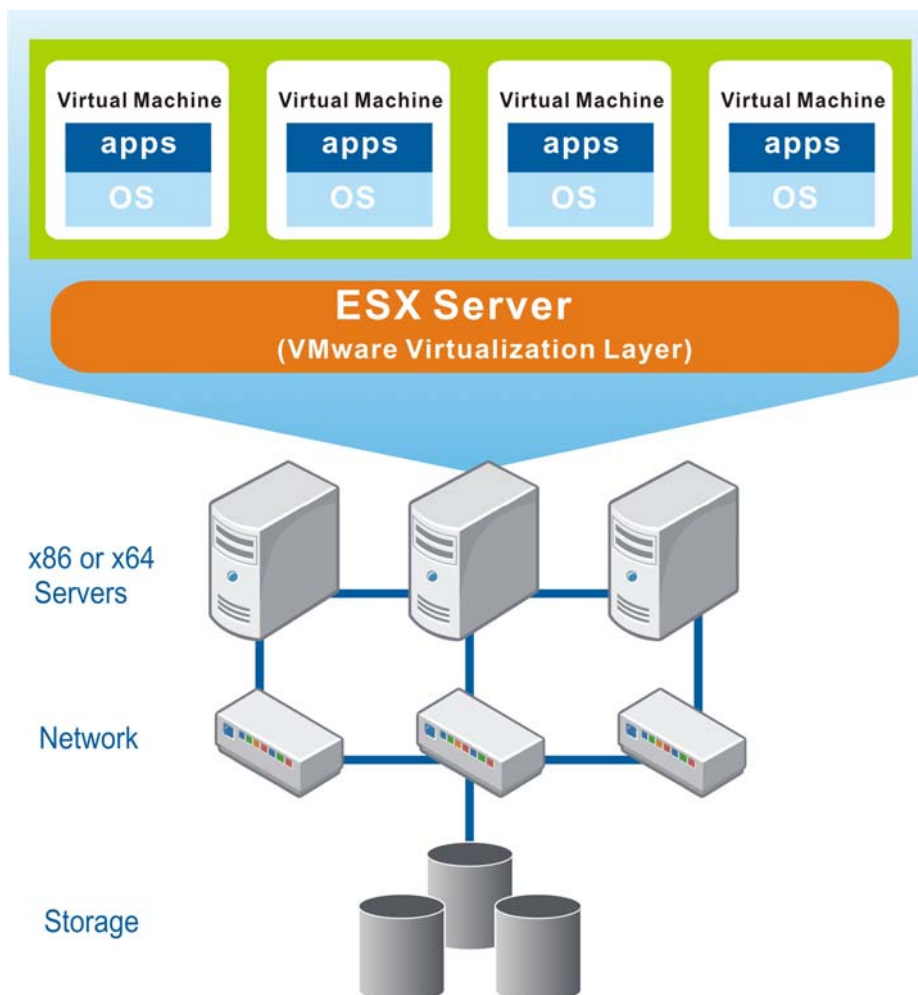


Figure 1. VMware ESX Server

Installed on an x86 or x64-based server, VMware ESX Server provides a virtualization layer on the host OS to consolidate all hardware resources, including processors, memories, storage and networking, and apply them to virtual

machines (i.e. virtual servers) running on the physical server. Each virtual machine can have its own OS and applications. By sharing hardware resources across multiple virtual machines, users can improve resource utilization and in turn greatly reduce the costs spent on building datacenter's physical infrastructure.

Besides the optimized resource utilization, VMware virtualization also reduces deployment efforts and simplifies management tasks. Free from the physical considerations and requirements, deploying virtual machines can be done in minutes or hours. After the deployment, managers can easily monitor the entire virtual datacenter through a unified management pane. When any of the physical device fails, the high availability features of VMware can ensure continuous system operation.

Infortrend ESVA (Enterprise Scalable Virtualized Architecture) Storage Systems in VMware Virtualized

Environments

The Infortrend ESVA (Enterprise Scalable Virtualized Architecture) Series is a leading-edge storage solution designed for mid-range enterprise Fibre Channel or iSCSI SAN. At affordable price, it meets mission-critical storage demands for performance, scalability and reliability with advanced hardware design and comprehensive data services. On the innovative Enterprise Scalable Virtualized Architecture, various features, including storage virtualization, thin provisioning, distributed load balancing, automatic data migration, prioritized volume accessibility, and array-based snapshot and replication, are consolidated to realize optimal business benefits. Making ESVA virtualized storage complement VMware's virtual infrastructure, users can enjoy enhanced benefits of optimized returns of investment, simplified storage infrastructure and maximized application productivity.

Optimized Returns of Investment

With storage virtualization technology, the capacity and computing power of multiple ESVA systems can be consolidated into a storage pool. For the most efficient utilization of pooled storage capacity, ESVA arrays support thin-provisioning on its virtual architecture. Just-in-time capacity is dynamically allocated to applications when data is written. ESVA lowers operating costs by

minimizing the space, power and cooling expenses wasted on the large and under-utilized data volumes commonly seen in traditional storage environments. In addition to increasing capacity, ESVA also ensures efficient utilization of bandwidth. Knowing that not all applications are created equal, Infortrend designed ESVA with an intelligent access prioritizing mechanism. This mechanism ensures that all applications connecting to the same storage pool can achieve their ideal service levels

Simplified Storage Infrastructure

ESVA simplifies storage management by enabling a single point of administration. Scaling the ESVA storage pool is also an easy task. Expansion enclosures can be connected to the ESVA system for increased capacity. If you want to increase capacity and performance at the same time, you can scale out the virtualized foundation by adding additional ESVA systems. All scaling and configuration tasks can be done online. When a new ESVA system is added, the distributed load balancing technology will dynamically balance workloads among storage systems for increased processing power. Power is increased with capacity expansion, allowing it to handle even the most demanding high-performance applications. If you remove a system from the pool, the load-balancing technology will also automatically migrate data to maintain optimum performance without disrupting service.

Maximized Application Productivity

In the competitive business world, downtime can lead to profit loss, damage to a corporation's reputation and threaten business continuity. The revolutionary ESVA architecture eliminates downtime for storage scaling. ESVA also includes storage-based replication capabilities. Space-efficient snapshots can serve as granular recovery points based on which files can be restored and data can be rolled back. As to the full data copies created within or across storage pool(s), they can be readily leveraged by host applications to resume production when the original data is corrupted. By strategically deploying snapshot images and full data copies, users are able to obtain the highest data availability with a minimum of service downtime in case the storage is damaged by logical errors, physical errors or disasters.

To know more about Infortrend ESVA Series, please visit <http://esva.infortrend.com/>.

Using ESVA iSCSI-host Storage Systems in vSphere 4

With the Infortrend ESVA iSCSI Series, users can leverage the current investment on IP network to realize easy SAN consolidation. When data exponentially grow along with business development, the powerful, reliable and flexible iSCSI SAN can help users meet diverse and changing application needs with great cost-efficiency. The ESVA iSCSI Series provides two models: ESVA-E20 and ESVA-E60. The E20 accommodates large-capacity SATA drives to store and protect data with the optimal dollar-per-gigabyte advantage, while the E60 houses high-performance to handle demanding transactional workloads for top-tier applications. Both the models have passed the compatibility test ensuring their seamless integration into an VMware vSphere 4 environment: <http://www.vmware.com/resources/compatibility/search.php?action=search&deviceCategory=san&productId=1&advancedORbasic=advanced&maxDisplayRows=50&key=&release%5B%5D=13&datePosted=-1&partnerId%5B%5D=121&arrayTypeld%5B%5D=1&rorre=0>.

Planning Considerations

Data Formats

To make the data volumes on ESVA arrays accessible to ESX servers, they have to be configured as either VMFS (Virtual Machine File System) volume or RDM (Raw Device Mapping) volume¹. VMFS is VMware's proprietary clustered file system. It is the most common access method. If users would like to allow multiple virtual machines to run on and multiple physical servers to access a single volume, they should configure the volume with the VMFS format. Another alternative to make virtual machines access data volumes on the storage is RDM. Virtual machines access VMFS volumes and RDM volumes in different ways. As shown in **Figure 3**, virtual machines can directly access a virtual disk in the VMFS format but their access to the RDM volume is enabled through a mapping file in the VMFS volume. This mapping file contains metadata that redirects disk access to the physical devices.

¹ Maximum size of an RDM volume in 2TB.




Figure 2. Different Ways of Accessing a VMFS Volume and an RDM Volume

Treating the RDM volume as a local disk, virtual machines could format it in a proper way. RDM is especially useful in the following two applications:

1. To perform SAN-based snapshot/volume copy or other layered applications on virtual machines.
2. To leverage Microsoft Clustering Services (MSCS) to implement virtual-to-virtual clusters or physical-to-virtual clusters. Clustered data and quorum disks have to be configured as RDM volumes.

Deployment of VMFS Volumes

The following guidelines direct users to properly deploy VMFS volumes for their applications.

1. Virtual machine boot disks and application data should be stored in separate VMFS volumes. Most I/Os issued to boot disks involves paging activities and are sensitive to response time. By separating boot disks from application data, the risk of prolonged response time due to application related I/O activities can be mitigated.
2. Database platforms for enterprise data management, such as Microsoft SQL Server or Oracle, often use active logs and/or recovery data structures to track data changes. In cases of unplanned application or operating system disruptions, these active logs or recovery data structures are critical in ensuring system recovery and data consistency. Therefore, all virtual machines supporting such database platforms should be provided with an independent VMFS volume for storing active log files and recovery data structures. Furthermore, if the files or structures are mirrored, the source and the target should be stored in separate VMFS volumes.

3. Application data, including database files, should be stored in a separate VMware file system. Furthermore, this file system should not contain any structures that are critical for application and/or database recovery.
4. It is recommended that the VMFS volumes are no more than about 80% full. This ensures that administrators would not suddenly run out of space to accommodate user data and VMware snapshots for virtual machines.

RAID Level

ESVA storage arrays allow users to protect their data volumes with various RAID levels, including RAID 1, RAID 3, RAID 10, RAID 5 and RAID 6. Data volumes in the same storage array can be protected with different RAID levels. The following are general guidelines for you to configure RAID levels for your data volumes in an VMware virtualized environment.

1. Virtual machine boot volumes are generally subject to low I/O rates. The boot volumes can be configured with RAID 5 protection.
2. For most applications, RAID 5 is a proper level to protect virtual disks with. However, if the application involves extensive logging, such as financial applications, RAID 10 may be a better option.
3. Infrastructure servers, such as Domain Name System (DNS), perform most of their activities utilizing CPU and RAM, and therefore are often subject to low I/O rates. If users use virtual machines as infrastructure servers, it is proper to provide them with RAID 5-protected volumes as storage space.
4. Log devices for databases should be RAID 10-protected volumes. Furthermore, if databases or application logs are mirrored, the source and the target should be located on separate sets of disks (in VMFS format, if applicable).
5. The virtual machines that generate high workloads of small-blocked, random Read I/O, such as Microsoft Exchange, should be allocated RAID 10-protected volumes for better performance.
6. Large file servers with vast majority of the storage consumed by static files can be provided with RAID 5-protected volumes since the I/O rates are expected to be low.

Example Configuration Steps

The below example explains how to make ESVA iSCSI-host storage available to ESX servers using VMware iSCSI software initiator.

Step 1: Add iSCSI VMkernel Ports to the vSwitch

In the vCenter GUI, select *Configuration* tab from the top menu and then click *Networking* under the *Hardware* panel. Click *Properties* for the vSwitch you would like to use for iSCSI traffic.

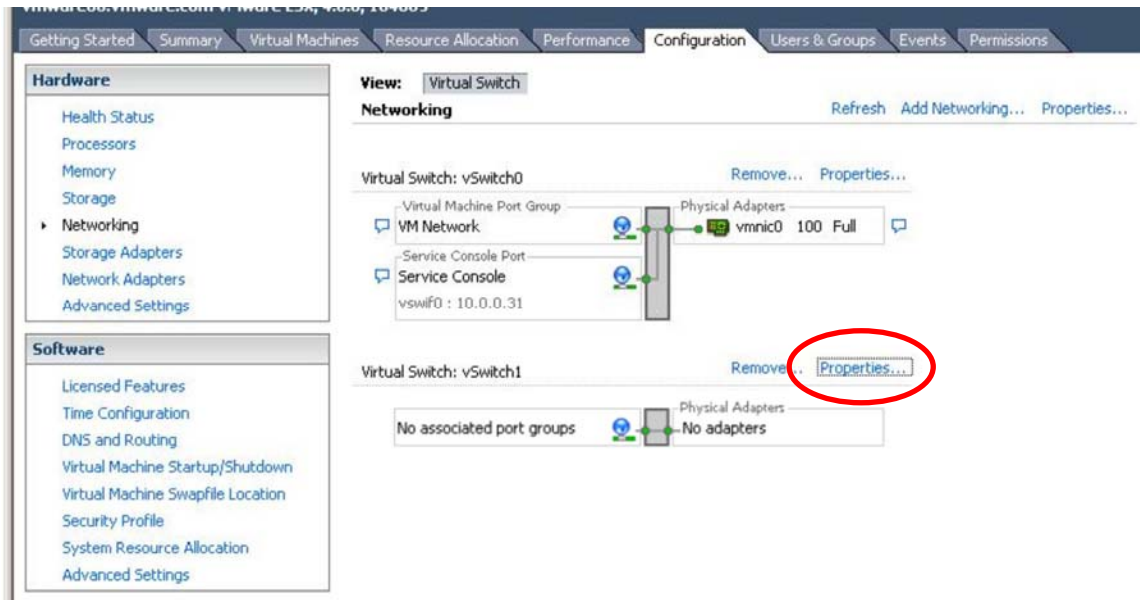


Figure 3. Accessing vSwitch Properties

In the *vSwitch Properties* window, click *Add...*

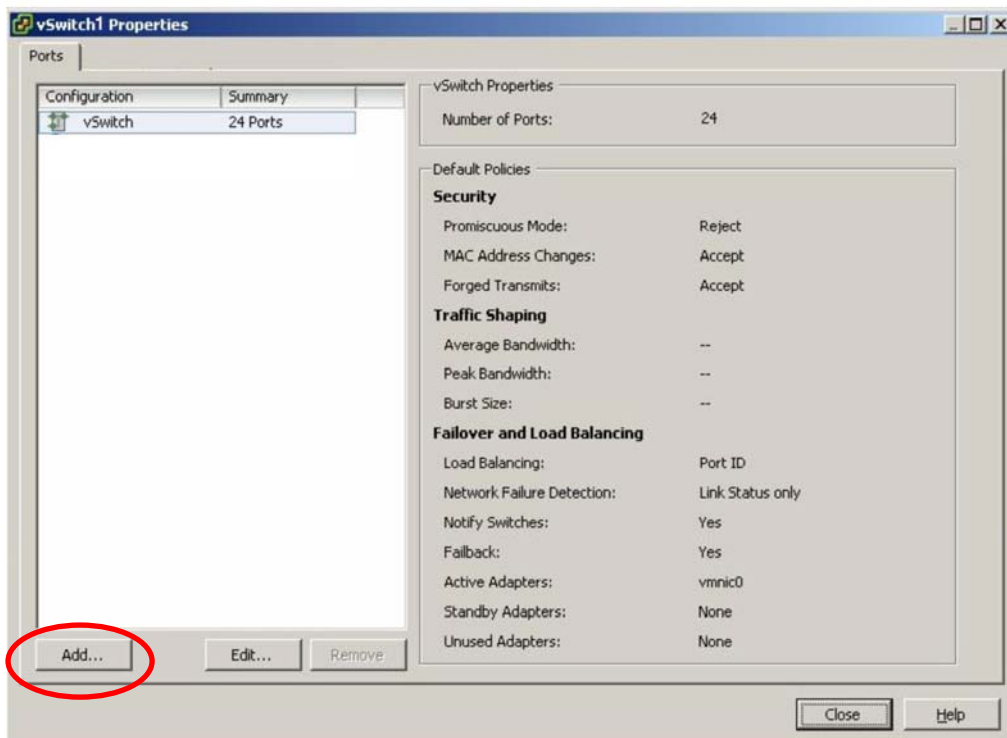


Figure 4. vSwitch Properties

In the *Add Network Wizard*, select *VMkernel* as the connection type.

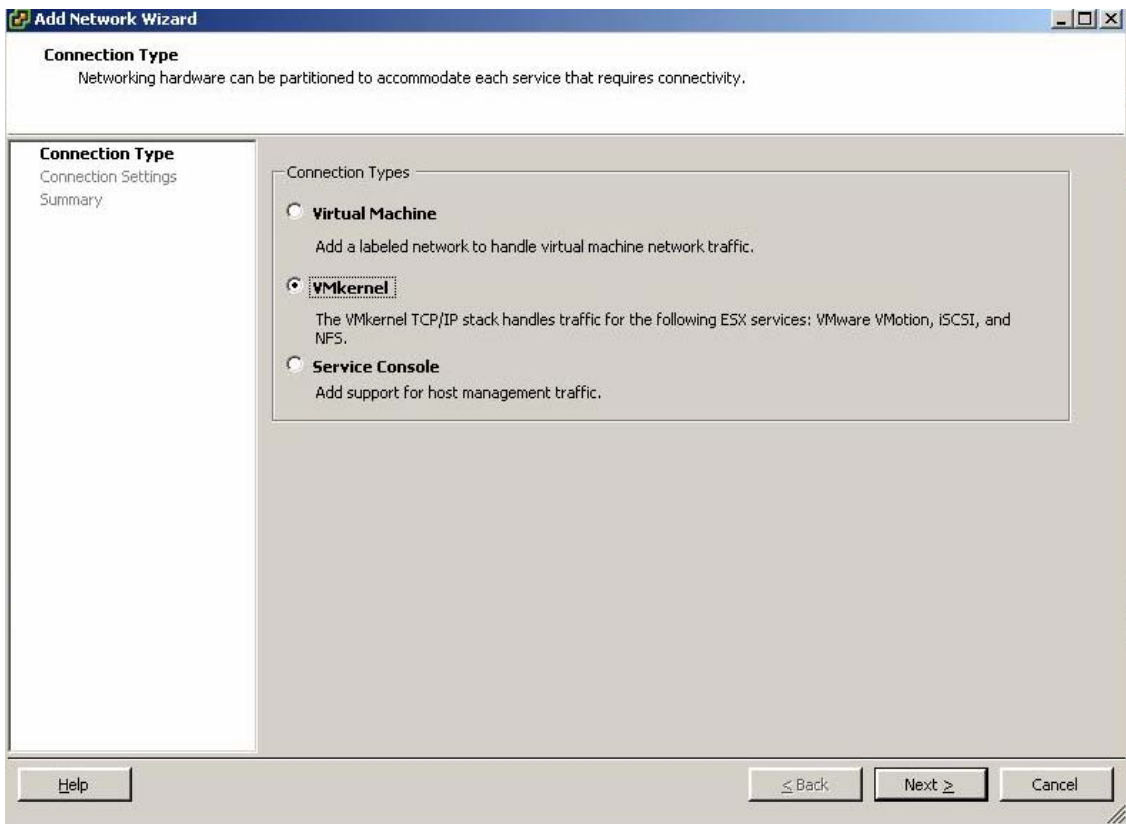


Figure 5. Selecting VMkernel as the Connection Type

After setting Network Label and configuring IPs, a VMkernel port will be successfully added to the vSwitch. Repeat the process if you need to create more VMkernel ports.

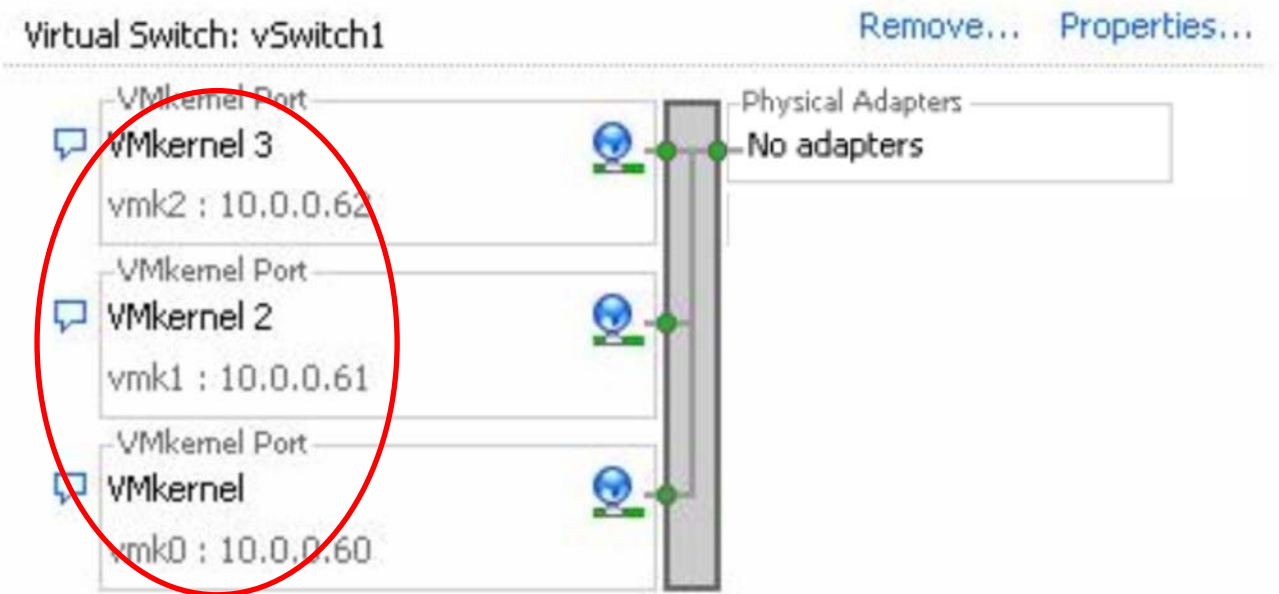


Figure 6. VMkernel Ports Created

Note: VMware vCenter allows the maximum of eight connections to a single data volume and a single network adapter can be associated with more than one VMkernel port. The number of VMkernel ports you create in this step may vary with your applications needs. For example, if you have two network adapters and would like to associate two VMkernel ports for each of them, you have to create four VMkernel ports in this step.

Step 2. Assign Network Adapters

In vCenter GUI, Click *Networking* under the *Hardware* panel. Click *Properties* for the vSwitch you would like to use for iSCSI traffic. In the *vSwitch Properties* window, click the *Network Adapters* tab, and click *Add...* In the *Add Adapter Wizard*, select the vmins you would like to add to the vSwitch and use for iSCSI traffic; then click *Next*.

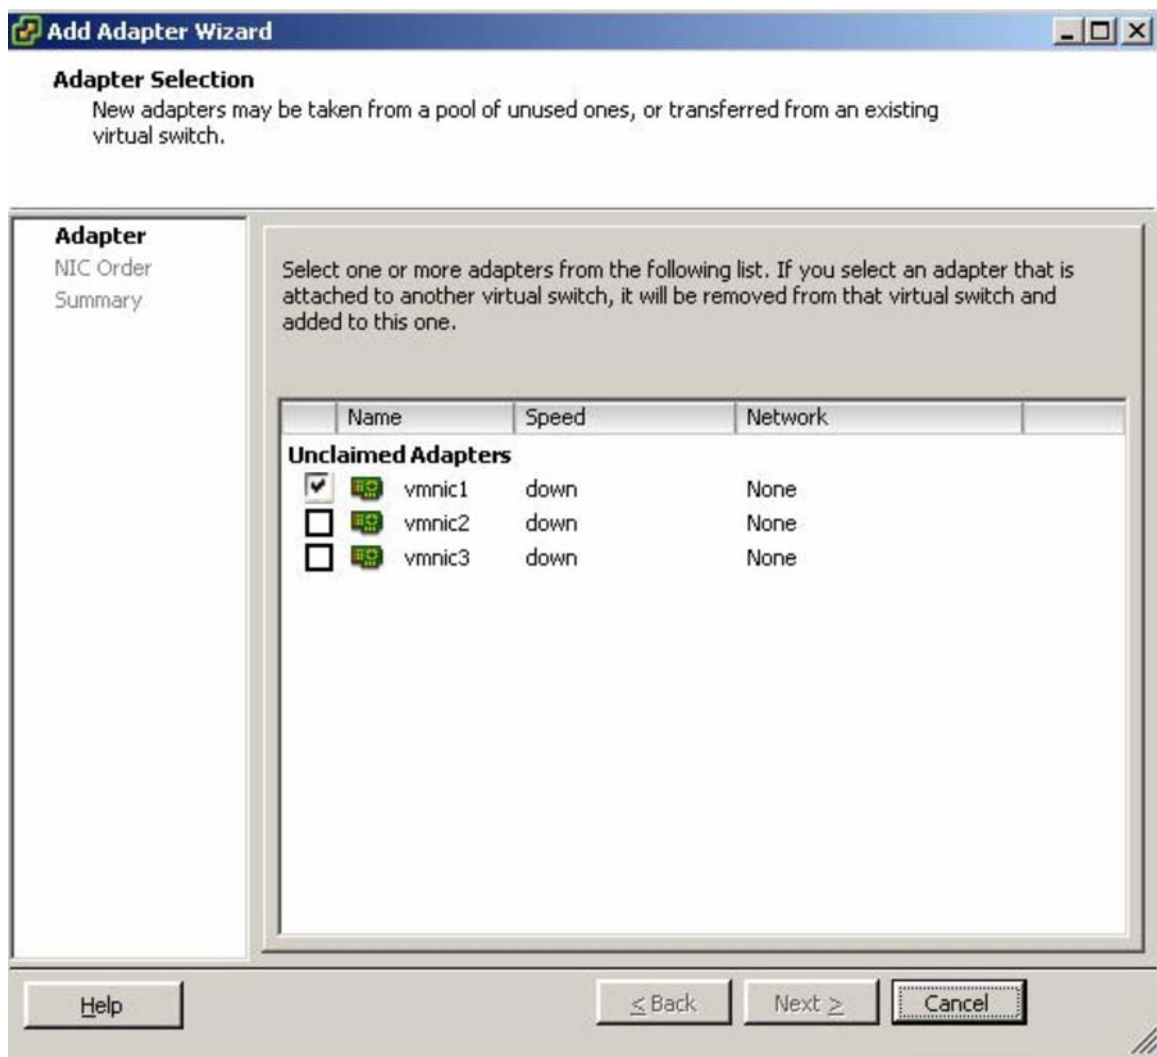


Figure 7. Selecting the Network Adapter

After configuring the policy failover order, click *Finish*.

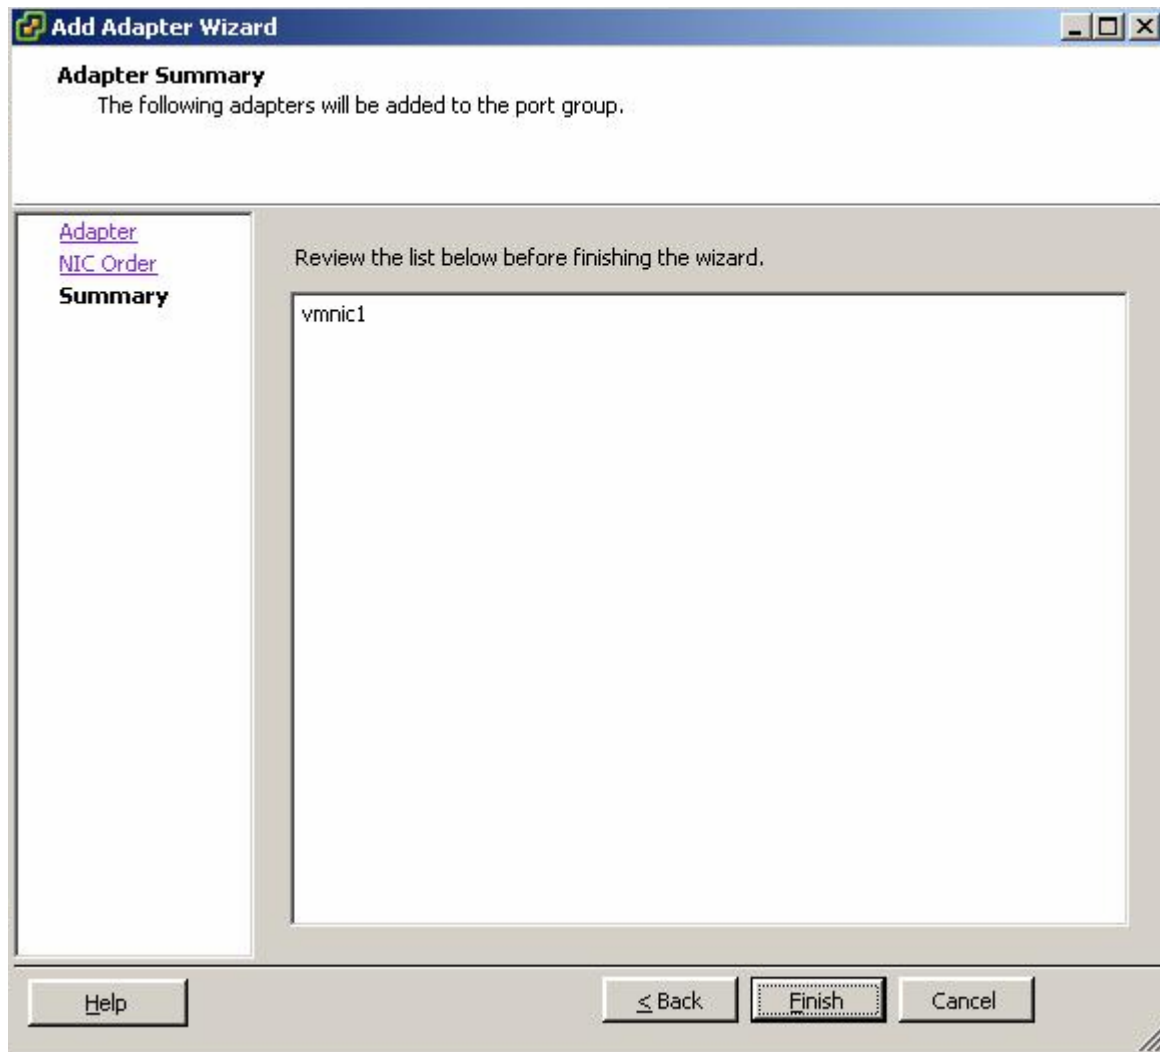


Figure 8. Finishing Adding the Adapter

Now the adapter(s) you choose will appear in the vCenter GUI under the *Network Adapters* tab.

Step 3. Associate VMkernel Ports with Network Adapters

In vCenter GUI, click *Networking* under the *Hardware* panel. Click *Properties* for the vSwitch. Select one of the VMkernel ports, and then click *Edit...* In the *VMkernel Properties* window, click the *NIC Teaming* tab. Then select the checkbox of *Override vSwitch Failover Order*. By default, the VMkernel port will be assigned to all of the NICs you install. To make each VMkernel port mapped to only one adapter and balanced across them, you should select the adapter you do not want the VMkernel port is assigned to, and then click the *Move Down* button until it is listed under *Unused Adapters*.

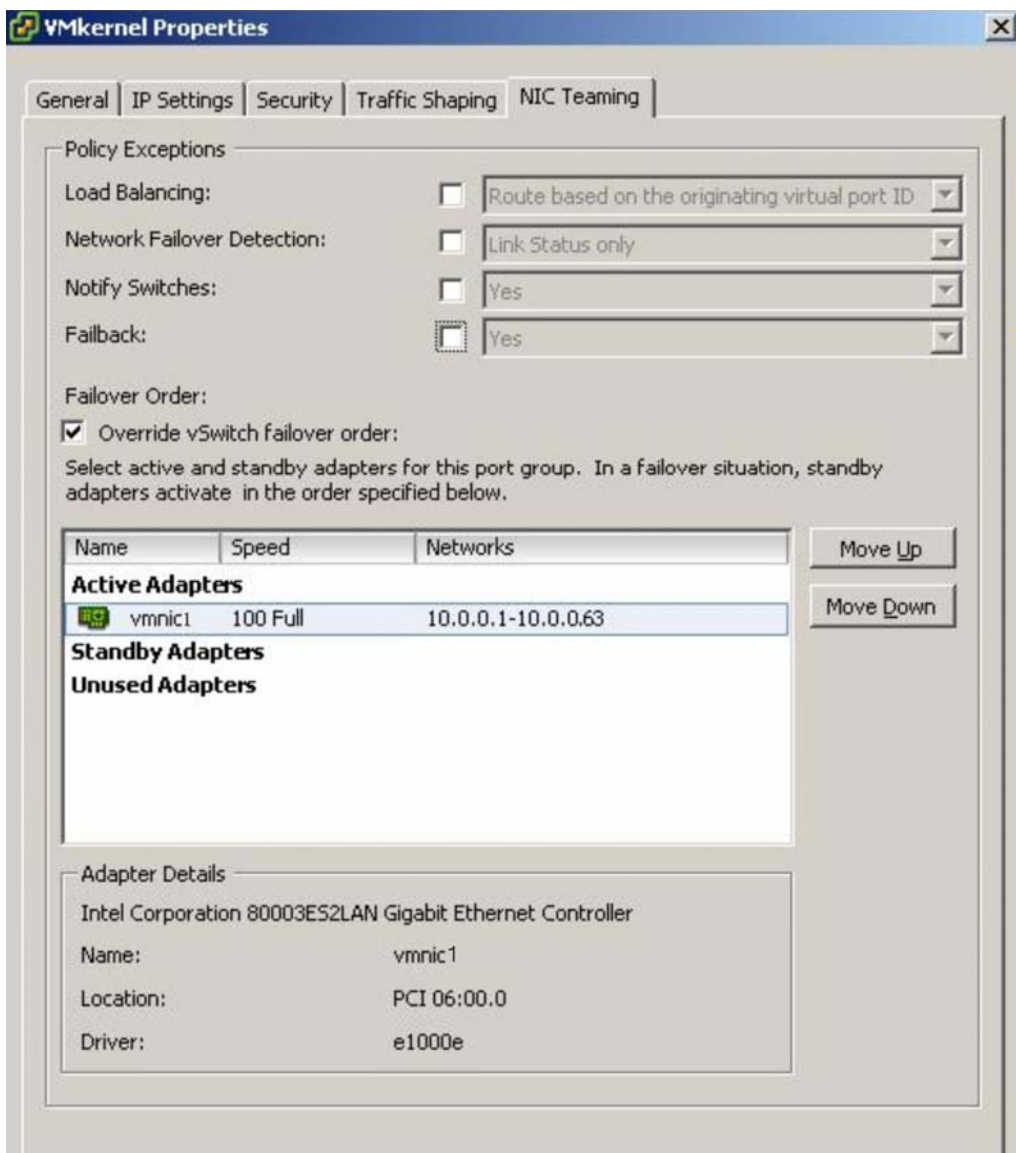


Figure 9. VMkernel Properties

Step 4. Enable VMware iSCSI Initiator and Set up CHAP Authentication

In vCenter GUI, click *Storage Adapters* under the *Hardware* panel. Select the iSCSI software adapter, and then click *Properties*. In the *iSCSI Initiator Properties* window, select the *General* tab. Click the *Configure...* button, and a *General Properties* window will pop out. Select the check box *Enabled* in the window. Then click *OK*.

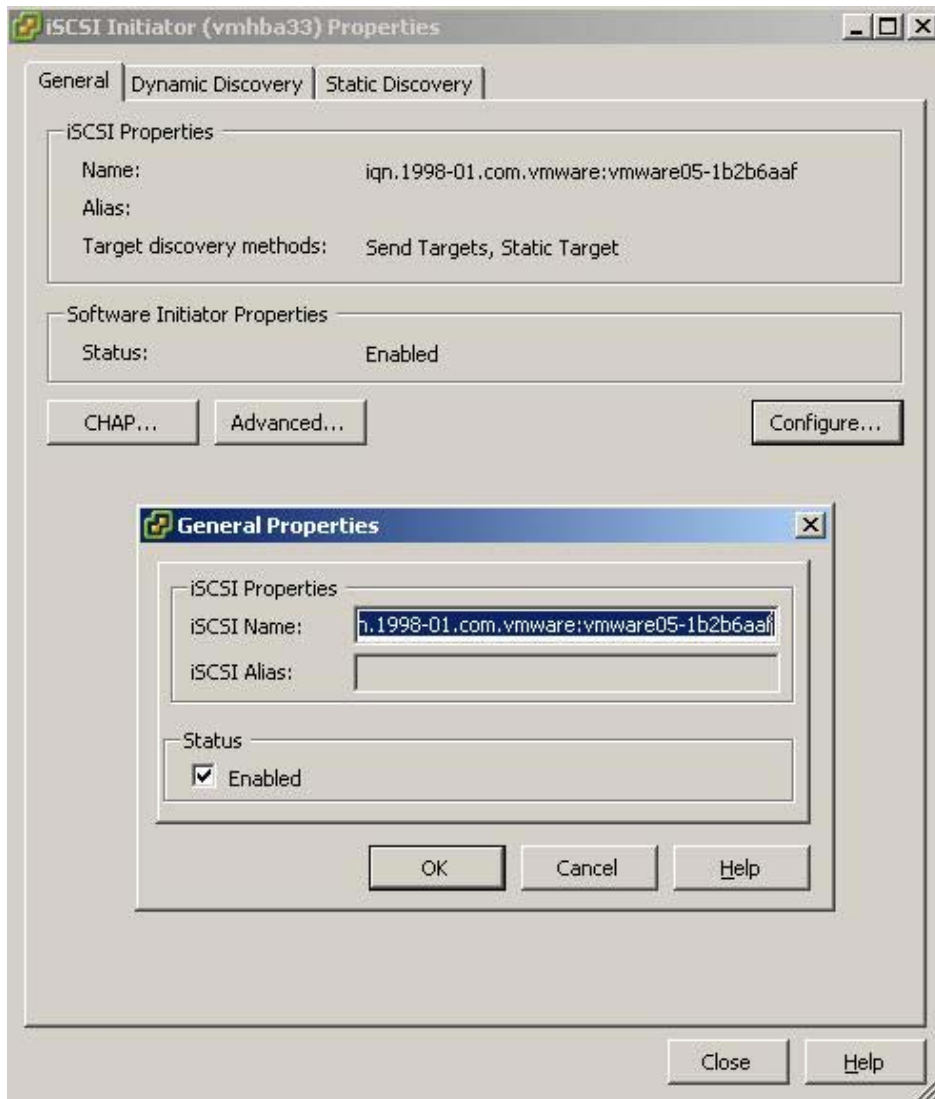


Figure 10. Enabling iSCSI Initiator

Click the *CHAP...* button and the *CHAP Credentials* window will pop out. VMware ESX4.0 supports both one-way and two-way CHAP authentication. If you would like the iSCSI targets to authenticate host only, select *Use CHAP* from the dropdown menu after *Select Option* in the *CHAP* configuration section. If you would like host also to authenticate iSCSI targets, select *Use CHAP* from the dropdown menu after *Select Option* in the *Mutual CHAP* configuration section. , Then enter the required information of name and secret. With the setting, successful connections can be established only when they pass CHAP authentication

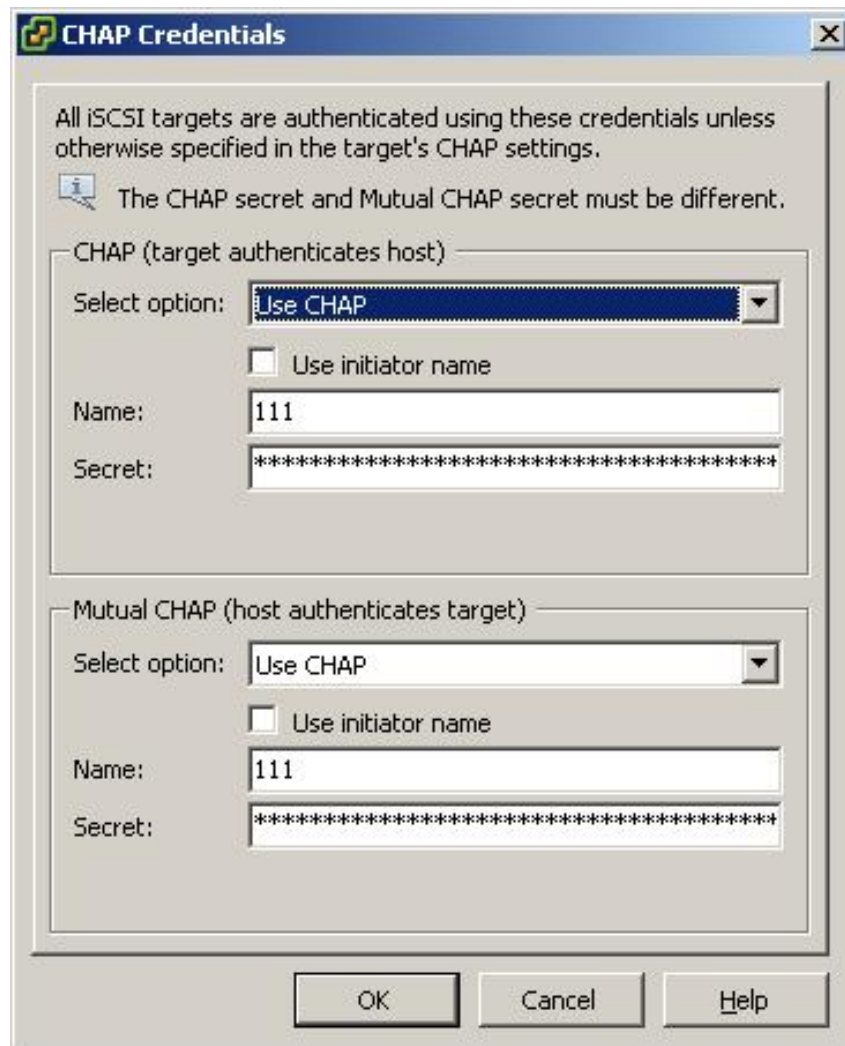


Figure 11. Enabling CHAP Authentication

Note: Infortrend storage does not support configuring different credentials for different targets. Here the CHAP authentication is configured at the root level of the iSCSI initiator and the credentials will be inherited by all iSCSI targets.

Step 5. Binding VMkernel Ports to iSCSI Software Initiator

This step of configuration has to be done through CLI commands. Before configuring, record the vmhba# of the iSCSI software initiator and the vmk# of each VMkernel port. In vCenter GUI, the former can be seen by clicking *Storage Adapters* under the *Hardware* panel.




Figure 12. Viewing the vmhba# of iSCSI Software Adapter

The latter can be seen by clicking *Networking* under the *Hardware* panel. Under the vSwitch for iSCSI traffic, you can see the vmk# listed under each VMkernel port.

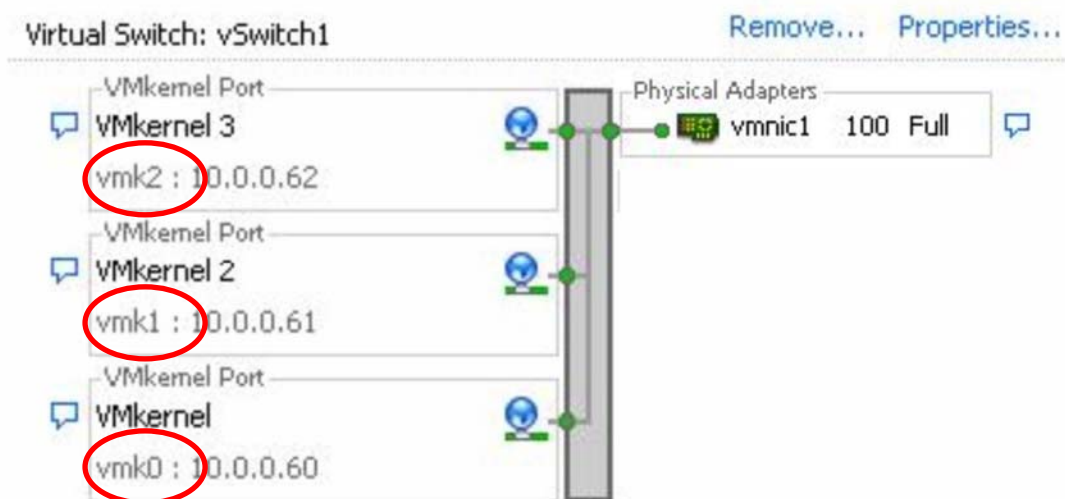


Figure 13. Viewing the vmk# of VMkernel Ports

If you would like to bind the vmk0 VMkernel port to the iSCSI Software Adapter vmhba33, type the following command:

```
esxcli swiscsi nic add -n vmk0 -d vmhba33
```

Repeat the commands until all VMkernel ports are bound to the iSCSI Adapter.

Step 6. Add ESVA iSCSI Channel Port IPs to the iSCSI Software Adapter

In vCenter GUI, click *Storage Adapters* under the *Hardware* panel. Select the iSCSI software adapter, and then click *Properties*. In the *iSCSI Initiator Properties* window, select the *Dynamic Discovery* tab. Click *Add...* and the *Add Send Target Server* window will pop out. In the window, type in the IP addresses of the iSCSI channels ports on your ESVA system.

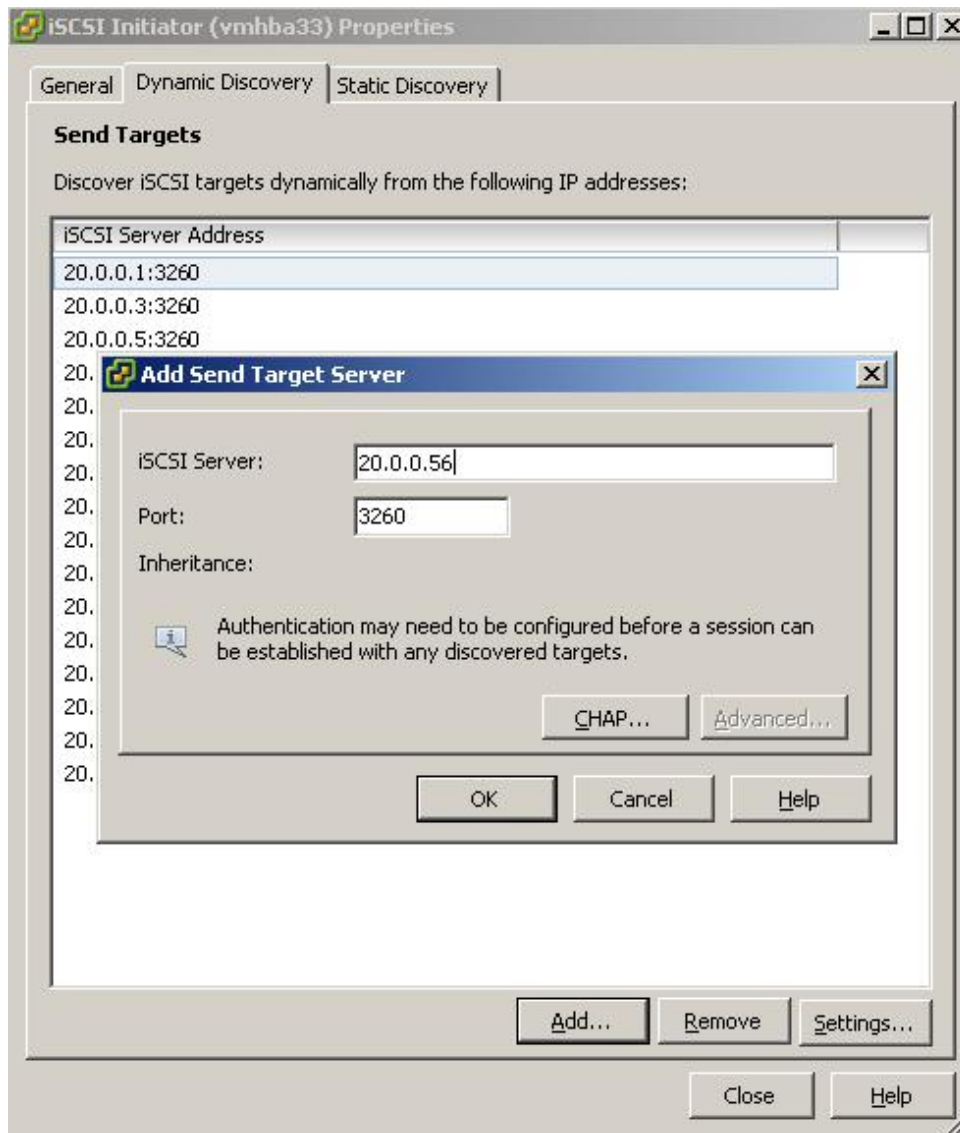


Figure 14. Adding iSCSI Channel Port IP Addresses

Repeat the process until all IP addresses are added. Then click *Close*.

Step 7. Create a Virtual Volume on ESVA

This configuration step should be done in the Virtualization Manager of Infortrend's proprietary storage management suite – SANWatch. Please refer to your Virtualization Manager User's Guide for details.

Step 8. Add a VMFS Datastore to the Virtual Machine

In vCenter GUI, click *Storage Adapters* under the *Hardware* panel. Select the iSCSI Software Adapter, and then click *Rescan*. In the pop-out window, select both *Scan for New Storage Devices* and *Scan for New VMFS Volumes*, and click *OK*.

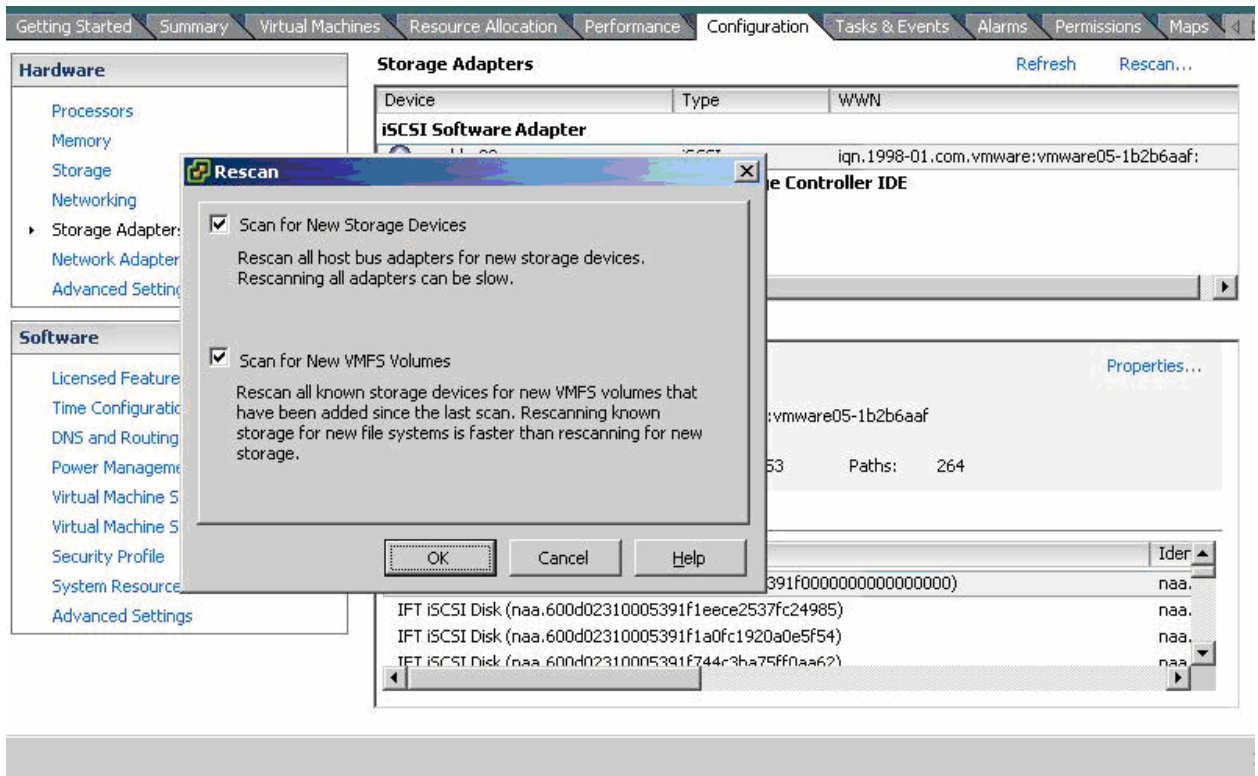


Figure 15. Rescanning Storage

In vCenter GUI, click *Storage* under the *Hardware* panel. Then click *Add Storage....*

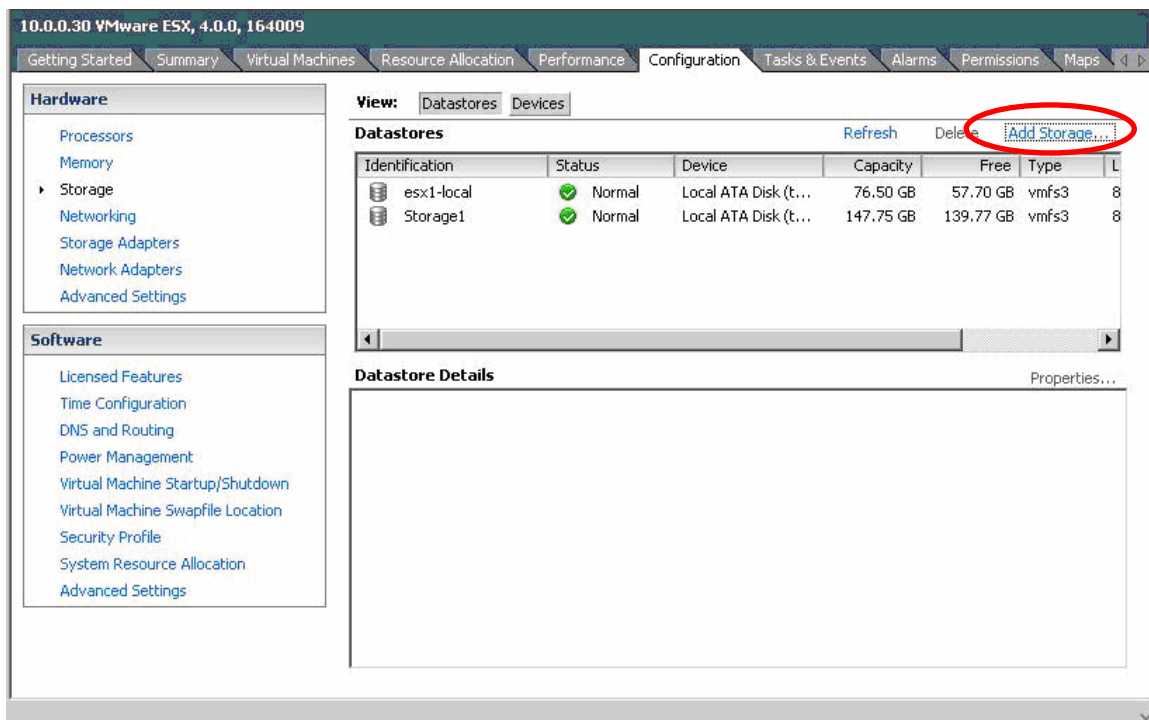


Figure 16. Adding Storage

The *Add Storage* window will pop out. First select Disk/LUN as the storage type, and click *Next*.

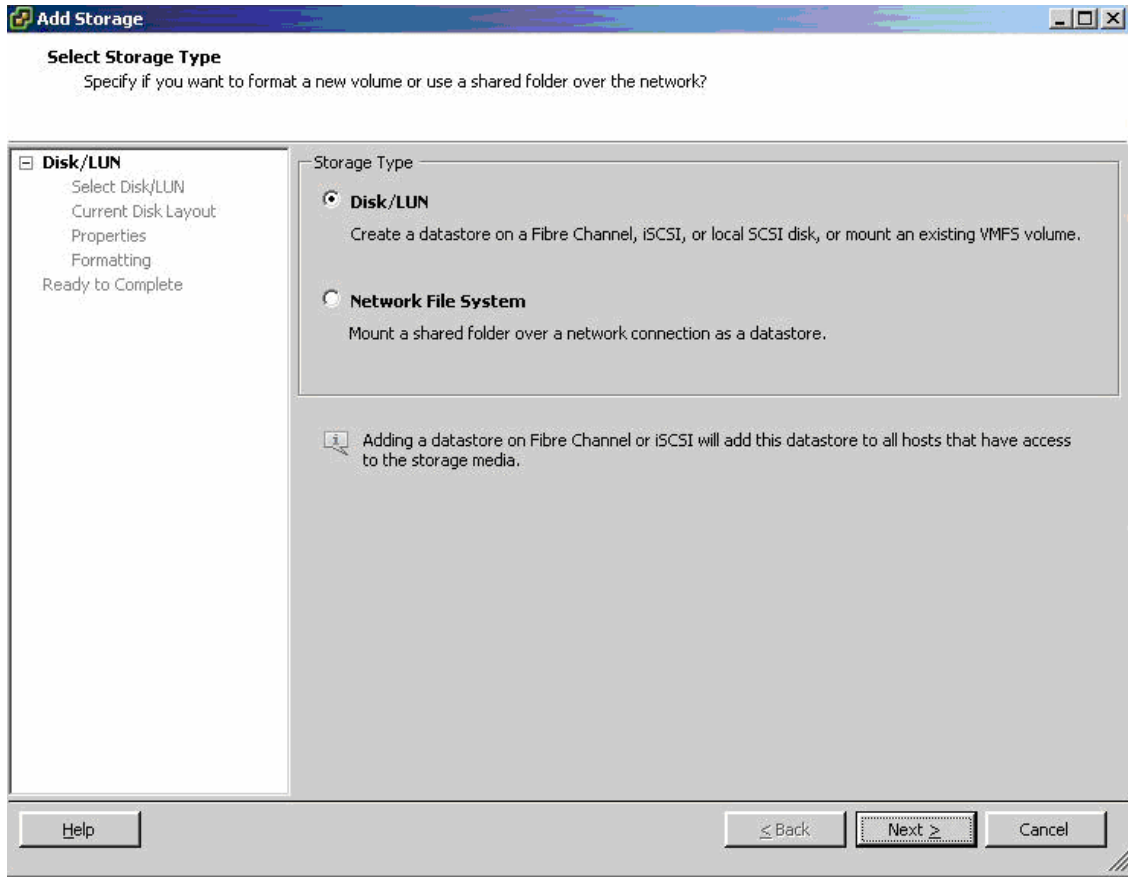


Figure 17. Selecting Disk/LUN as the Storage Type

Then in the list of all available iSCSI disks, select the one you would like to add as the new VMFS datastore. Click *Next*.

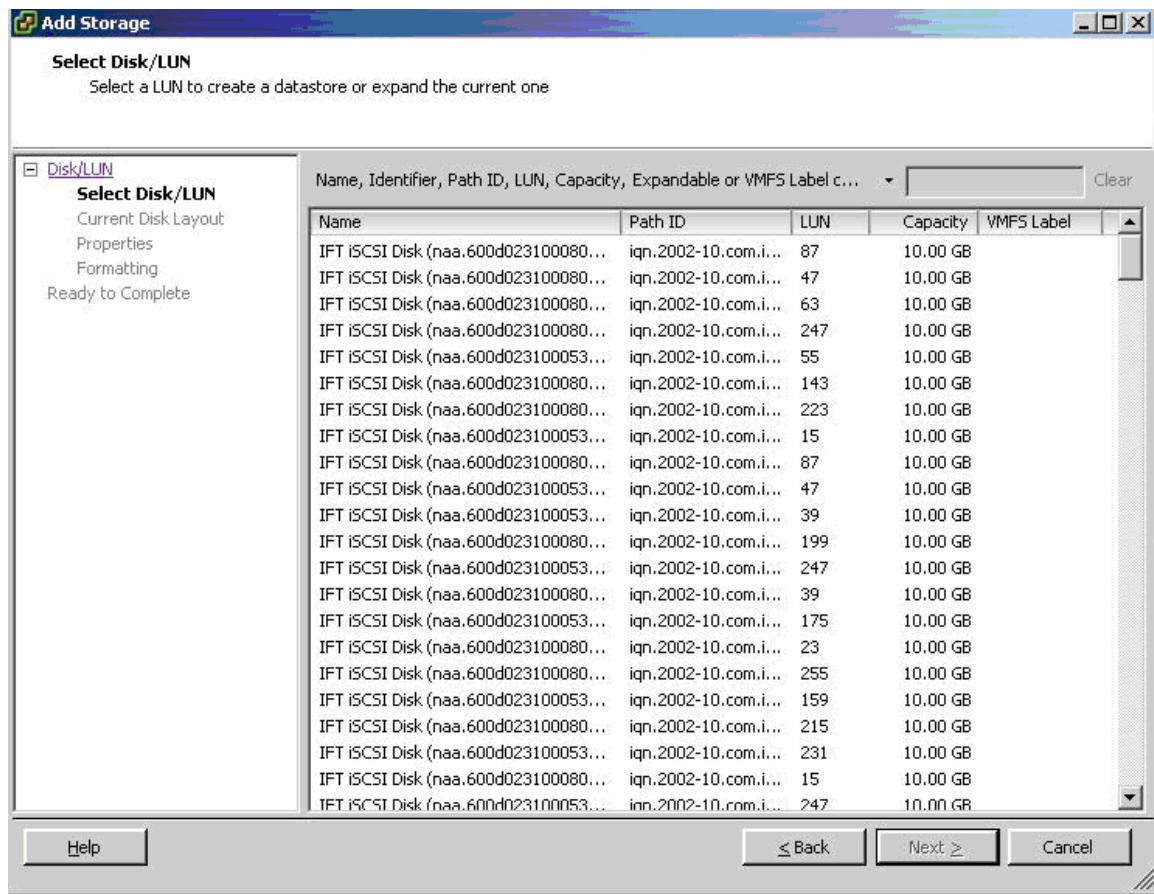


Figure 18. Selecting an iSCSI Disk

After going through the configuration process of entering the datastore name and maximum file size, click *Finish*.

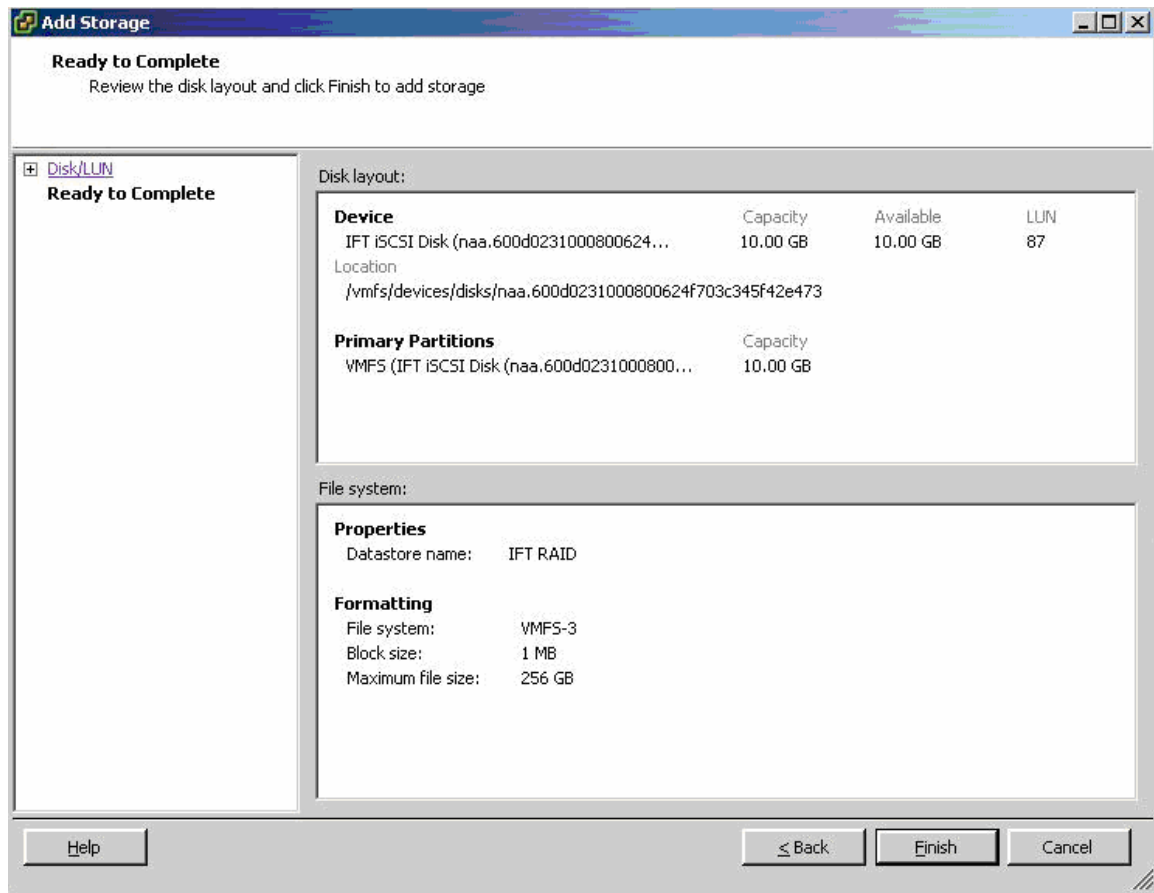


Figure 19. Finishing Datastore Creation Process

Then a new VMFS datastore will be created.

Step 9. Enable VMware Native Multipathing

In vCenter GUI, click *Storage* under the *Hardware Panel* and select the volume you would like to configure multipathing on. Then right click and select *Manage Paths*. In the pop-out window, select the path selection policy from the drop-down menu next to *Path Selection*.





Figure 20. Selecting a Path Selection Policy

Note: The selection can vary with application needs but Round Robin is suggested since it allows the volume to distribute workloads across all available paths.

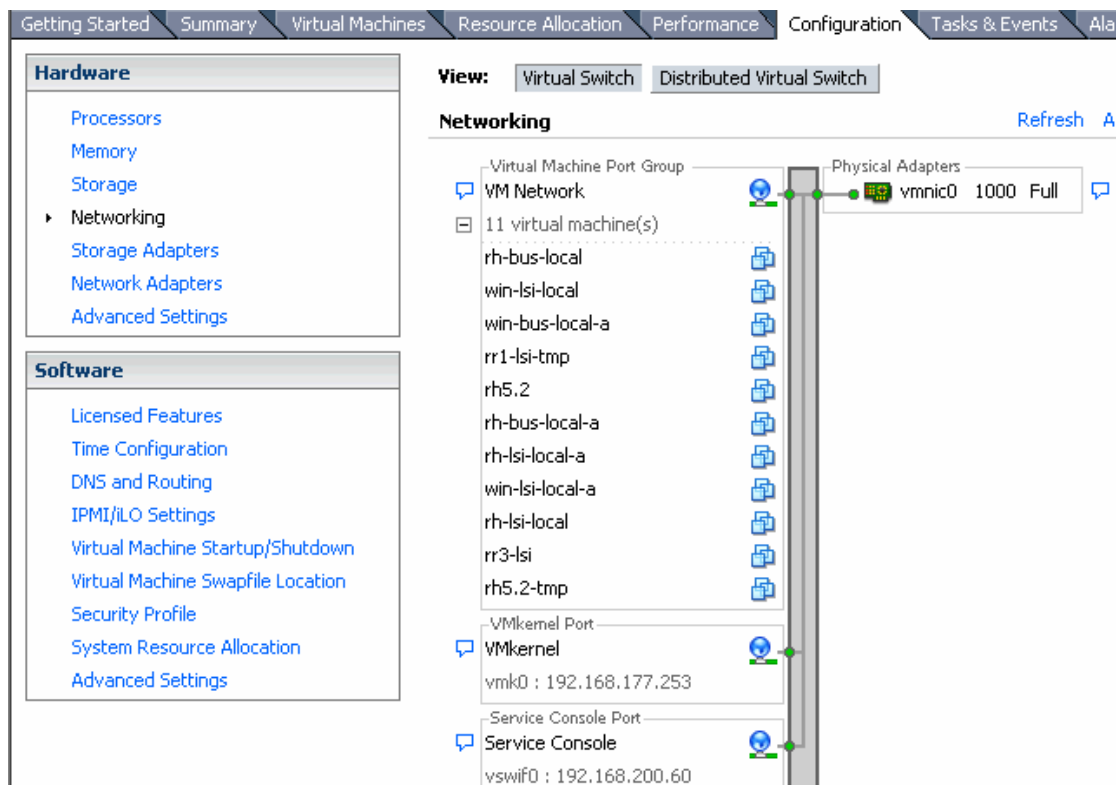
Basic Troubleshooting and FAQ

1. What information should I prepare when I need your help to do troubleshooting?
Please provide the following information
 - ESX server version (for example, ESX4.0, ESX3.5)
 - Storage model, its firmware version and event logs
 - Topology of your ESX server, switch and storage as shown below



- Storage configuration, including LDs, Virtual Pools, Virtual Volumes partitions, and LUN mapping; you can get them via SANWatch
- Descriptions of the behaviors making you run into the problem
- ESX server configuration by screenshots of *Networking*, *Storage Adapters* and *Maps* in vCenter/VirtualCenter GUI

Networking: example





Storage Adapters: example

The screenshot shows the vSphere Configuration tab for Storage Adapters. The left sidebar contains a tree view with 'Storage Adapters' selected. The main area is divided into 'Storage Adapters' and 'Details' sections.

Storage Adapters

Device	Type	WWN
iSCSI Software Adapter		
vmhba33	iSCSI	iqn.1998-01.com.vmware:esx4-1-686cdb64
82801EB (ICH5) SATA Controller		
vmhba2	Block SCSI	
vmhba32	Block SCSI	
AIC-8902 U320 OEM		

Details

vmhba33

Model: iSCSI Software Adapter
iSCSI Name: iqn.1998-01.com.vmware:esx4-1-686cdb64
iSCSI Alias:
Connected Targets: 2 Devices: 2 Paths: 4

View: **Devices** Paths

Name	Identifier
IFT iSCSI Disk (naa.600d0231000e4261000000001227ac0d)	naa.600d0231000e426
IFT iSCSI Disk (naa.600d0231000e4261000000000de00b7)	naa.600d0231000e426

Maps: example

The screenshot shows the vSphere Maps tab. The main area displays a network diagram with nodes for 'VM Network', 'private', 'c-w2k3-b', and 'esx15'. The 'private' node is connected to 'VM Network' and 'c-w2k3-b'. 'c-w2k3-b' is connected to 'esx15'. The IP address '192.168.140.15' is shown below the 'c-w2k3-b' node.

Overview

Map Relationships:

Virtual Machine Resources

Host Options

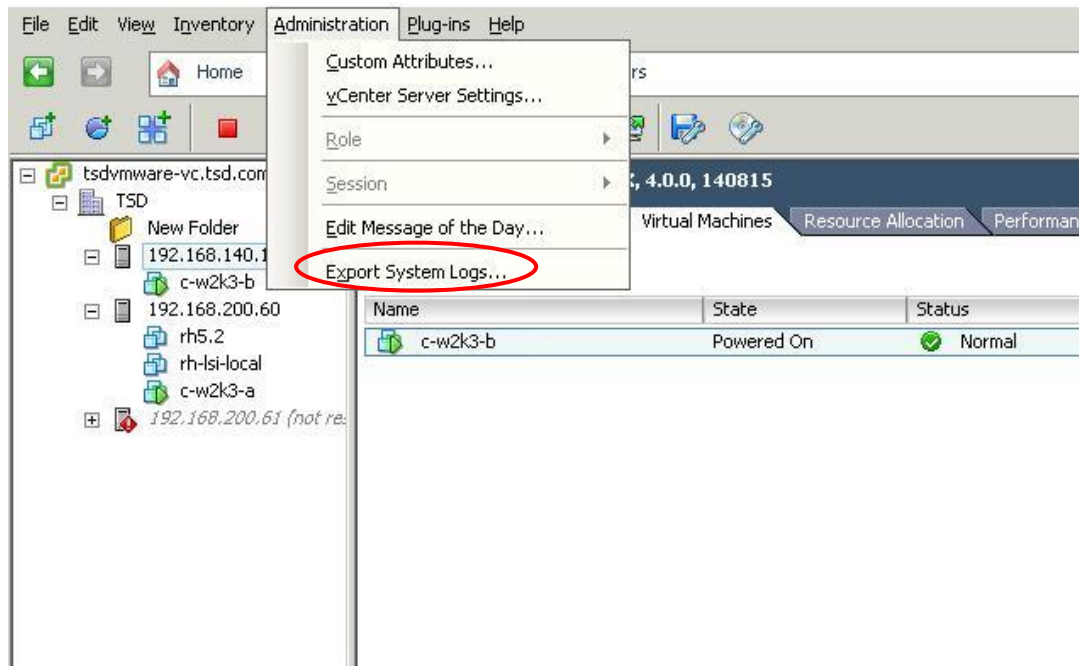
- Host to VM
- Host to Network
- Host to Datastore

VM Options

- VM to Network
- VM to Datastore
- Show only powered on VMs

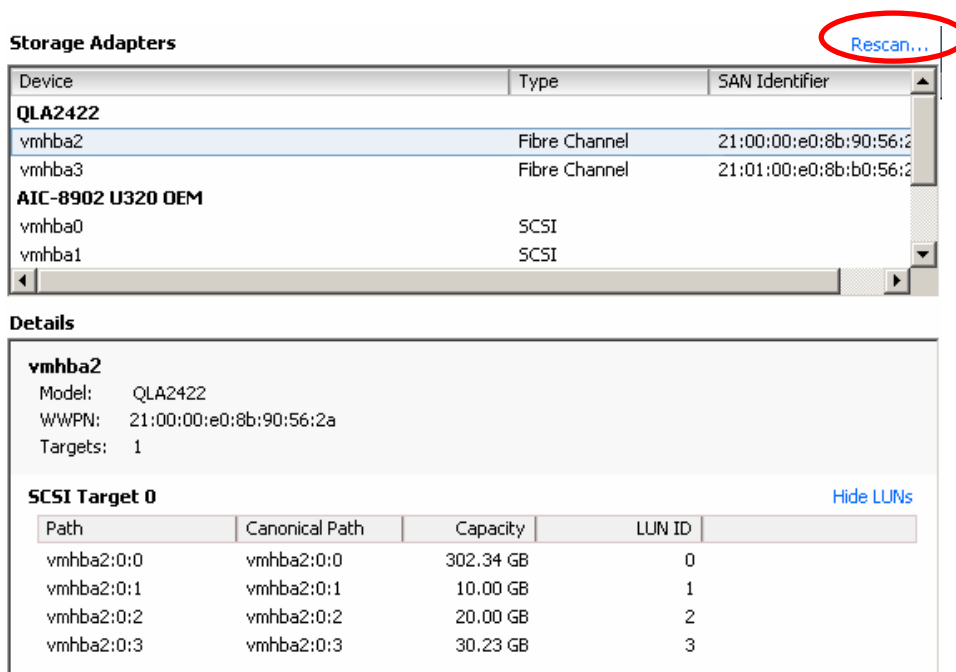
Apply Relationships

- ESX server event logs; you can get them from the folder `/var/log/vmkernel` or from vCenter/VirtualCenter GUI. Click *Administration* on the top menu and select *Export System Logs* from the drop-down menu.



2. Why can't I see the VMFS data volume I created before?

Please first check whether you can see the physical storage device you used to create the VMFS data volume in vCenter/Virtual Center GUI. If not, please check whether your cables are correctly connected and storage LUN mapping is properly configured. Then execute *Rescan* in the *Storage Adapters* screen.



If the problem is still not solved, please prepare the information mentioned in question 1 and contact us for further troubleshooting.

3. What are the storage configuration limitations in a VMware virtualized environment? Please check http://www.vmware.com/pdf/vsphere4/r40/vsp_40_config_max.pdf for details.

4. How can I make data paths successfully failover when redundant controllers failover?
If you are using ESX 4.0 with EonStor storage arrays installed with firmware ver.3.64 or later, system would automatically handle this without any manual configuration. However, if you are using ESX 3.x, please add a footnote for redundant controller storage following the steps in the application note: http://www.infortrend.com/doc/appNote/APP_VMware_footnote_1117.pdf.

5. Can virtual machines be migrated to a different data volume without interruptions? Yes, virtual machines can be migrated online to another data volume. Please check the following links for details:
http://www.vmware.com/products/vi/storage_vmotion.html
<http://blogs.vmware.com/vi/2008/06/storage-vmotion.html>

6. If I would like to implement multipathing, is there any special settings I should do on my storage? Should I install EonPath?
No, there is no special settings for storage. You can just follow general multipathing configurations. Moreover, since VMware supports native multipathing functions, you need not install other drivers, including EonPath, for multipathing implementations.

Appendix: Enable Jumbo Frames

Jumbo Frames is one of the new advanced capabilities supported by VMware vSphere ESX4.0. It allows larger packets to be transferred between host and the SAN so that system efficiency and performance can be increased. Currently, users can enable Jumbo Frames only through CLI commands. Below are the example configuration steps..

Note: Please ensure that Jumbo Frames are supported by your networking devices before doing the configuration.

Step 1. Enable Jumbo Frames on the vSwitch

If you want to enable Jumbo Frames for vSwitch1, run the following command:

```
esxcfg-vswitch -a vSwitch1
```

To verify this is successfully done, run the command:

```
esxcfg-vswitch -l
```

Then you will see output similar to this:

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch1	24	1	24	9000	

PortGroup Name	VLAN ID	Used Ports	Uplinks

Step 2. Add iSCSI VMkernel Ports to the vSwitch and Enable Jumbo Frames on the Ports

To enable Jumbo Frames, you should add iSCSI VMkernel ports via CLI commands, instead of vCenter GUI. If you would like to add a VMkernel port named VMkernel1 to vSwitch1, run the following command:

```
esxcfg-vswitch -A VMkernel1 vSwitch1
```

Then configure the IP Address, Subnet Mask, and enable Jumbo Frame support for the new VMkernel port with the following command:

```
esxcfg-vmknic -a -i 10.10.0.61 -n 255.255.255.0 -m 9000 VMkernel1
```

Repeat the above steps to create more VMkernel ports. To verify this is successfully done, run the command:



esxcfg-vswitch -1

Then you will see output similar to this:

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch1	24	4	24	9000	

PortGroup Name	VLAN ID	Used Ports	Uplinks
VMkernel1	0	1	
VMkernel2	0	1	
VMkernel3	0	1	

Jumbo Frames is now successfully enabled. The remaining configuration steps for using ESVA in an vSphere 4 environment are the same as those mentioned in the main contents.