# DELL POWERVAULT MD3200I / MD3600I DEPLOYMENT GUIDE FOR VMWARE ESX4.1 SERVER SOFTWARE

A Dell Technical White Paper
Version 1.7

*PowerVault MD3200i and MD3600i Storage Arrays*

www.dell.com/MD3200i

www.dell.com/MD3600i

www.dell.com/pvresources

**DISCLAIMER:**

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

For more information, contact Dell.
Information in this document is subject to change without notice.

http://www.dell.com

# TABLE OF CONTENTS

# TERMINOLOGY/GLOSSARY

VD == virtual disk

VM == virtual machine

NIC == network interface card

MPIO == Multi-Path I/O

SAN == Storage Area Network

RDM == Raw Device Map

DVS == Distributed Virtual Switch

HA == high availability

DRS == Distributed Resource Scheduler

MRU == Most Recently Used

IQN == iSCSI Qualified Name

# INTRODUCTION

These steps necessary to connect to an iSCSI SAN are documented in VMware's iSCSI SAN Configuration Guide which can be found on VMware's website. This whitepaper goes into depth with configuration examples for connecting to a Dell™ PowerVault™ iSCSI SAN utilizing the software iSCSI initiator inside the VMware® ESX4.1 Server™. Additionally, Appendix A gives details regarding configuring hardware initiators from within VMware® ESX4.1 Server™.

The Dell™ PowerVault™ MD3200i & MD3600i iSCSI SAN storage solution consists of either a standard or high availability configuration. The standard (simplex) configuration has a single controller with four 1GbE ports. It can be deployed to support up to 32 hosts non-redundantly. The high availability (duplex) configuration has dual controllers with four 1GbE ports per controller for a total of eight 1GbE ports. The dual controller option can also connect up to 32 fully redundant hosts

Provisioning of MD3200i storage in a VMware® ESX4.1 environment is a multi-step process starting with establishing the iSCSI session, defining the server topology for host access, and finally allocating storage to the individual virtual machines (VMs).

VMware® vSphere4™ offers many new and advanced enhancements over the iSCSI software initiator in conjunction with iSCSI SAN connectivity. Many of these new features require advanced configuration in order to work properly. Administrators who are familiar with ESX 3.5 iSCSI SAN configuration may find that their current configuration steps are not sufficient to enable all of the advanced features offered in vSphere4.

## *NEW FEATURES IN VSPHERE4 SOFTWARE ISCSI INITIATOR*

VMware vSphere4 ESX4.1 has new support for various new advanced capabilities that were not found in ESX 3.5. This whitepaper will cover the new features in the iSCSI software initiator as well as how to configure them to connect to the SAN.

**iSCSI Software Initiator** – With ESX4.1, the iSCSI software initiator was re-written from the ground up for better performance and functionality.

**Jumbo Frames** – With ESX 4.1 and vSphere4, Jumbo Frames can be enabled on the iSCSI software initiator. Jumbo Frame support allows for larger packets to be transferred between the ESX4.1 servers and the SAN for increased efficiency and performance. Jumbo Frame Support can be enabled via the vSphere vCLI.

**MPIO** – With ESX4.1 and vSphere4, customers can benefit from Multi-Path I/O(MPIO) from the ESX4.1 server and the SAN. This allows for multiple connections to be used to allow for greater bandwidth. This is especially important for the PowerVault SAN as each PowerVault member has multiple connections and now ESX4.1 can take full advantage of these connections.

**Third Party MPIO Support** – With ESX4.1 and vSphere4, VMware has provided an architecture that enables PowerVault platforms to provide new and advanced intelligent integration.

# SUPPORTED HARDWARE AND SOFTWARE

## HARDWARE REQUIREMENTS

Refer to the following VMware website for a complete up-to-date list of the prerequisites for installing VMware ESX server.
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_server_config.pdf

## SUPPORTED OPERATING SYSTEMS FOR MD3200I ARRAY

ESX4.1 and ESXi5.0 are two of the supported VMware OSs for MD3200i. Please see the MD Support Matrix at support.dell.com for a complete listing of supported VMware OSs.

# ARCHITECTURAL SETUP

As a best practice, Dell recommends using a separate Gigabit Ethernet network switch to handle iSCSI storage traffic, see PowerVault MD32x0i IP SAN Best Practices. An IP SAN consists of one or more hosts, connected to one or more storage arrays through an IP network, utilizing at least one switch in the network infrastructure. Each switch in the configuration has a path to the MD3200i via duplex quad-port controllers.

There are many ways to implement an IP SAN based on need, available resources and intended application. Some general rules can be used when designing your IP SAN. For example, in order to maximize the throughput of your storage arrays, all host ports should be utilized. Nothing requires you to use all four ports for a given server. Figure 1 below shows the high level connection details with all four controller paths utilized. However, there are no technical architectural or functional features in the MD32x0i/36x0i that requires you to use all four ports for a given server.

At the hypervisor level, there is no ESX-driven requirement, for either performance-related or pathing-related implementation that requires multiple IP-address subnets. More specifically, all iSCSI ports can reside on the same subnet. The ESX native multipathing (NMP) driver is able to use all paths on the same subnet, so if you choose a RR pathing policy, these ports will all be used for Round Robin. Depending upon the workload needs, as shown in Figure 1 below, you certainly could only use two of the ports if that meets your IO <throughtput> requirements.

Later in this whitepaper we include an in-depth explanation regarding why one might want to use multiple subnets.

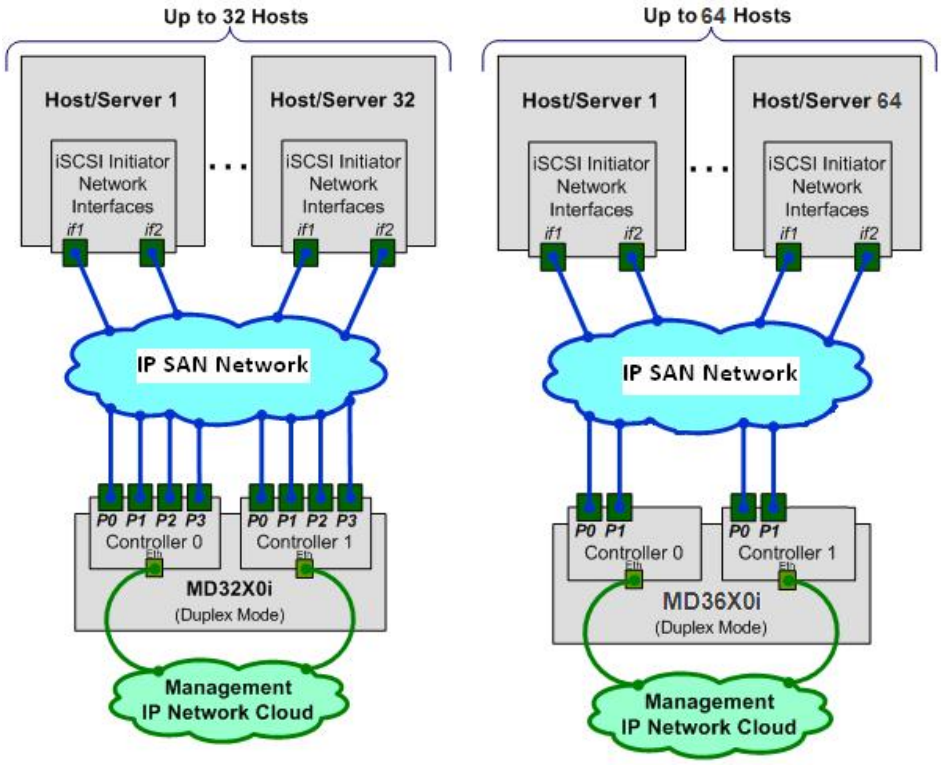


**FIGURE 1: HIGH-LEVEL CONNECTION DETAILS FOR BOTH MD3200I AND MD3600I**

## CONSIDERATIONS WHEN USING ISCSI SOFTWARE OR HARDWARE INITIATORS FOR ESX4.1 ON THE MD3200I STORAGE ARRAY

Taking advantage of new ESX4.1 features requires some new steps to be taken by ESX administrators. Configuration is done via either the vSphere GUI or using vCLI. The remainder of this whitepaper focuses on installation and configuration of an iSCSI software initiator connection to a PowerVault Series SAN. Each of these commands can be found inside the VMWARE ISCSI SAN CONFIGURATION Guide and where names and IP Addresses are used, they will be different for each environment. This serves as an example and demonstration of how to configure a new vSphere4 ESX4.1 server correctly and connect it to the PowerVault SAN. The screenshots displayed in this whitepaper are for reference only and may vary from site specific implementations.

## ESTABLISHING SESSIONS TO A SAN

Before continuing the examples, we first must discuss how VMware ESX4.1 establishes its connection to the SAN utilizing the vSphere4 iSCSI Software Adapter.

With previous versions of ESX, sessions were established using a single NIC path and any additional NICs were there for failover only. With the improvements to vSphere4 and MPIO, administrators can now take advantage of multiple paths to the SAN for greater bandwidth and performance. This does require some additional configuration which is discussed in detail in the sections below.

VMware uses VMkernel ports as the session initiators. Each VMkernel is bound to a separate physical adapter. Depending on the environment this can create a single session to an array or up to eight sessions (ESX4.1 maximum number of connections per volume). For a typical deployment, it is recommended to use a one-to-one (1:1) ratio of VMkernels to physical network cards. As shown in Figure 2 below, this means if there are four physical NICs, you would establish one VMkernel per physical NIC and associate a separate NIC with each VMkernel port.

This scheme can be expanded depending on the number of NICs you have in the system. As the environment grows, you can establish multiple sessions to the SAN by oversubscribing VMkernel ports to actual physical NICs. This establishes multiple sessions to an array but still utilizes the same physical NICs as the means to get to the storage. However, VMWare reports issues with oversubscribing VMkernel ports. Consequently we recommend maintaining a 1:1 port binding.

Once these sessions to the SAN are initiated, the VMware NMP will take care of load balancing and spreading the I/O across all available paths.

***Note: Port binding requires that all target ports of the storage array must reside on the same broadcast domain as the VMkernel ports because routing is not supported with port binding. See VMware KB #2017084 here.***

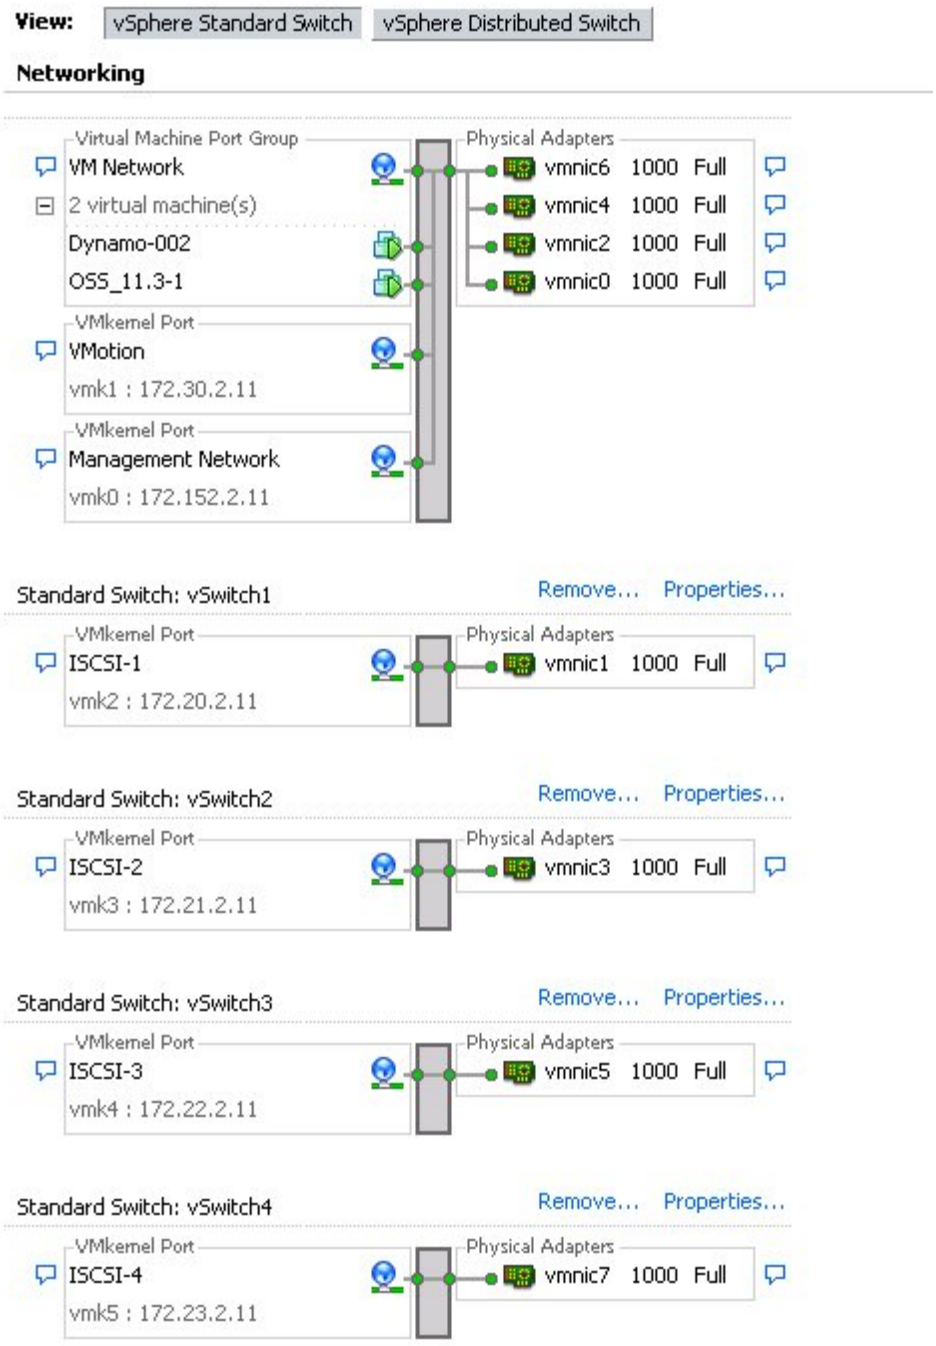


**Figure 2: Port Binding**

The following assumptions are made for these examples:

1. Running ESX4.1
2. Running latest Dell PowerVault MD3200i firmware
3. More than one Network Interface Card (NIC) set aside for iSCSI traffic.

4. ***Note: Port binding requires that all target ports of the storage array must reside on the same broadcast domain as the VMkernel ports because routing is not supported with port binding. See VMware KB #2017084 here.***

5. No Distributed Virtual Switch (DVS) for iSCSI traffic

6. The MD3200i IP address on the host ports will have already been configured, see PowerVault MD32x0i IP SAN Best Practices

7. The network topology is connected and both the ESX server and MD3200i are powered on.

8. The environment uses multiple NICs and attaching to a Dell PowerVault SAN utilizing Native Multipathing (NMP) from VMware.

# CONFIGURE iSCSI STORAGE ON ESX4.1 SERVER - INSTALLATION STEPS

Not every environment requires all of the steps detailed in this example.

Users connecting their vSphere4 environment using just iSCSI HBAs or users wishing to only assign a single iSCSI NIC with no Jumbo Frame support will not follow these steps and instead configure their environment as normal. Users who wish to only enable Jumbo Frame support for their environment will want to take note of steps 1 and 2 but only create a single VMkernel port through the vCenter GUI after that.

Users wishing to only enable Jumbo Frame support for the iSCSI connection need to follow 1 and 2 with the following changes:

Step 1: Configure vSwitch and Enable Jumbo Frames – No changes to the instructions

Step 2: Add iSCSI VMkernel Ports – Instead of assigning multiple VMkernel Ports, administrators will only assign a single VMkernel Port

Once these two steps are done, the rest of the configuration can be accomplished in the vCenter GUI by attaching NICs, assigning storage and then connecting to the storage.

## CONNECT TO THE ESX SERVER/VCENTER USING VI CLIENT AND FOLLOW THE STEPS BELOW.

### STEP1: CONFIGURE VSWITCH & ENABLE JUMBO FRAMES

This step will create a new vSwitch and enable Jumbo Frame support for this switch. Currently there is no option to enable Jumbo Frames on a vSwitch from the VMware vCenter GUI so these commands must be run via vCLI. Be sure to check the environment to make sure that Jumbo Frames are supported at the networking layer before enabling it on the ESX host.

The following command will create a new vSwitch called vSwitch2:

esxcfg-vswitch –a vSwitch2

Next, enable Jumbo Frames on the vSwitch:

esxcfg-vswitch –m 9000 vSwitch2

To verify that the switch was configured properly run the following command:

esxcfg-vswitch –l

Your output will look similar to this:

| Switch Name | Num Ports | Used Ports | Configured Ports | MTU Uplinks |
|---|---|---|---|---|
| vSwitch2 | 64 | 1 | 64 | 9000 |

As shown in Figure 3 below, you can note the new vSwitch2 with the MTU of 9000 to verify that the switch was created correctly. You can also see it displayed in the GUI of vCenter. Throughout these procedures some of the verification can be done via command line or seen in the vCenter GUI. The polling rate of vCenter is not instant so changes will not show up immediately after it is typed.


**Figure 3: Creating a vSwitch**

### STEP2: ADD ISCSI VMKERNEL PORTS

This next step will assign VMkernel Ports to the new vSwitch2. It will also configure Jumbo Frame support as well as assign the IP Addresses.

Administrators familiar with iSCSI connectivity in ESX3.5 will find in 4.1 that it is no longer required to configure a Service Console port for the iSCSI connection. So the iSCSI switch environment can be on a different subnet from the public environment or existing service console. If multiple controller ports are used, each iSCSI VMkernel Port will need its own IP Address and they must all be on the same subnet as the MD array's controller or target IP Addresses.

Note that, If you choose to use all four *target* ports and you choose a Round Robin pathing policy, then the target ports need to be on separate subnets for RR to work correctly.

Alternatively, if you are using VLANs you'll need separate VLANs for RR to work correctly. As shown in Figure 1, you certainly can only use two of the ports if that meets your throughput requirements.

In a default configuration assign one VMkernel port for each physical NIC in the system. So if there are three NICs, assign three VMkernel Ports. This is referred to in VMware's iSCSI SAN Configuration Guide as 1:1 port binding.

VMware vCenter has a maximum of eight connections to a single volume. In this whitepaper we choose three connections in the 1:1 scenario, which results in twelve connections per ESXi host given four iSCSI software initiators mapped to the four controller ports. This provides scalability and performance as the SAN environment grows without having to make changes on each ESX host.

***Note: Port binding requires that all target ports of the storage array must reside on the same broadcast domain as the VMkernel ports because routing is not supported with port binding. See VMware KB #2017084 here.***

Always keep the entire virtual datacenter in mind when deciding on path and volume count. View the Release Notes of the PowerVault Firmware for the current connection limits for the Dell PowerVault.

All of these configurations are done for the vSwitch itself. This means that once it is done, the ESX4.1 host will create multiple connections to the PowerVault SAN. Once this is configured there only need to be changes made if more NICs are being added or if more or less paths are needed.

*Note: Host profiles do not keep information on Jumbo Frames or Port Bindings.*

For the rest of this example the configuration steps and commands will be given for the 1:1 binding.

The following command will add a new iSCSI VMkernel Port named iSCSI1 on the vSwitch created in the previous step.

esxcfg-vswitch –A iSCSI1 vSwitch2

This next command will configure the IP Address, Subnet Mask and enable Jumbo Frame support for the new VMkernel Port iSCSI1

esxcfg-vmknic –a –i 10.10.6.206 –n 255.255.255.0 –m 9000 iSCSI1

For our example with a 1:1 relationship with 3 NICs we need to create 2 more VMkernel Ports named iSCSI2 and iSCSI3

esxcfg-vswitch –A iSCSI2 vSwitch2

esxcfg-vmknic –a –i 10.10.7.207 –n 255.255.255.0 –m 9000 iSCSI2

esxcfg-vswitch –A iSCSI3 vSwitch2

esxcfg-vmknic –a –i 10.10.8.208 –n 255.255.255.0 –m 9000 iSCSI3

To verify the configuration enter the following command:

esxcfg-vswitch –l

The output will look similar to this:

| Switch Name | Num Ports | Used Ports | Configured Ports | MTU | Uplinks |
|---|---|---|---|---|---|
| vSwitch2 | 64 | 7 | 64 | 9000 | |

| PortGroup Name | VLAN ID | Used Ports | Uplinks |
|---|---|---|---|
| iSCSI3 | 0 | 1 | |
| iSCSI2 | 0 | 1 | |
| iSCSI1 | 0 | 1 | |

This will show the VMkernel ports that are assigned to the vSwitch. To verify the IP addresses enter the following command:

esxcfg-vmknic –l

The output will look similar to the graphic below.

You can also verify the IP Addresses via the vCenter GUI. In vCenter, on the ESX Host, navigate to ***Configuration -> Networking***.

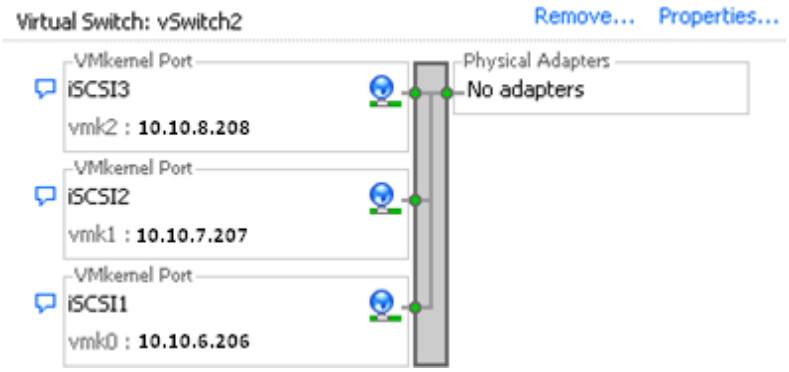


**Figure 4: Adding iSCSI VMKernel ports**

## STEP3: ASSIGN NETWORK ADAPTERS

The next step in the process is to assign the network adapters (NICs) that will be attached to the iSCSI network and used for iSCSI traffic. These will be attached to the vSwitch2 that we created earlier. This can be done two ways, in the vCenter GUI or by vCLI.

To list all of the adapters in the system run the following command:

esxcfg-nics –l

The output will look similar to this:

Name PCI Driver Link Speed Duplex MAC Address MTU

vmnic0 03:00.00 bnx2 Up 1000Mbps Full 00:21:9b:8b:4b:b0 1500

This will list all of the adapters in the system. Assign the NICs that are physically connected to the SAN infrastructure and to the vSwitch. The following command assumes that we are assigning vmnic1, vmnic2, and vmnic3 to the vSwitch.

esxcfg-vswitch –L vmnic1 vSwitch2

esxcfg-vswitch –L vmnic2 vSwitch2

esxcfg-vswitch –L vmnic3 vSwitch2

Once again, to verify the configuration type the following command to list the vSwitch information:

esxcfg-vswitch –l

Your output will look similar to the following. Note the new vmnics that were assigned to the vSwitch under uplinks.

| Switch Name | Num Ports | Used Ports | Configured Ports | MTU | Uplinks |
|---|---|---|---|---|---|
| vSwitch2 | 64 | 9 | 64 | 9000 | |

| PortGroup Name | VLAN ID | Used Ports | Uplinks |
|---|---|---|---|
| iSCSI3 | 0 | 1 | vmnic1,vmnic2,vmnic3 |
| iSCSI2 | 0 | 1 | vmnic1,vmnic2,vmnic3 |
| iSCSI1 | 0 | 1 | vmnic1,vmnic2,vmnic3 |

Adding a NIC can also be configured and verified in the vCenter GUI. Remember that the polling of vCenter is not instant so a refresh might need to occur to see the latest changes.

To configure this same process from the GUI, first navigate to the Networking section on the ESX host you are configuring. ***Configuration -> Networking***.

From here, click ***Properties*** on the vSwitch2.

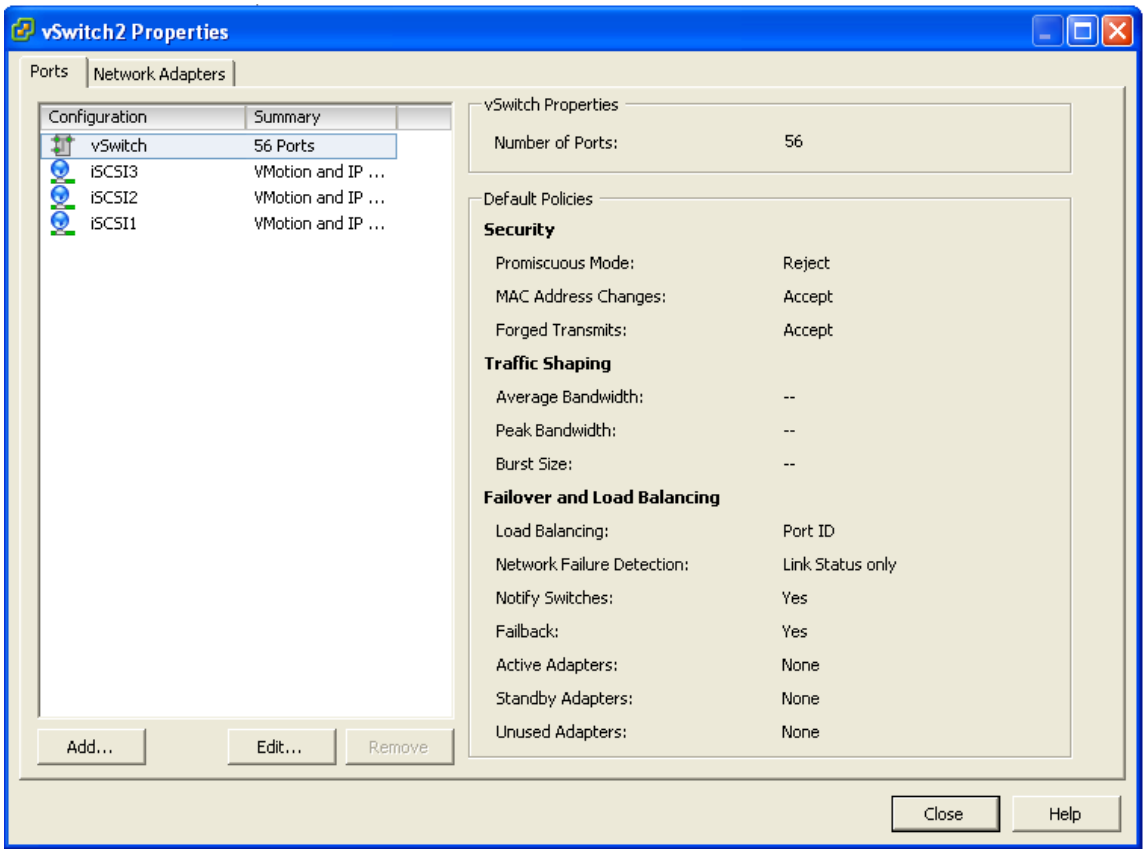


**Figure 5: Assigning Network Adapters**

Click the ***Network Adapters*** tab. Then click ***Add***. This will open up the Add Adapter Wizard. From here select the vmnics that you want to add to the vSwitch. In our example it will be vmnic1, vmnic2 and vmnic3.

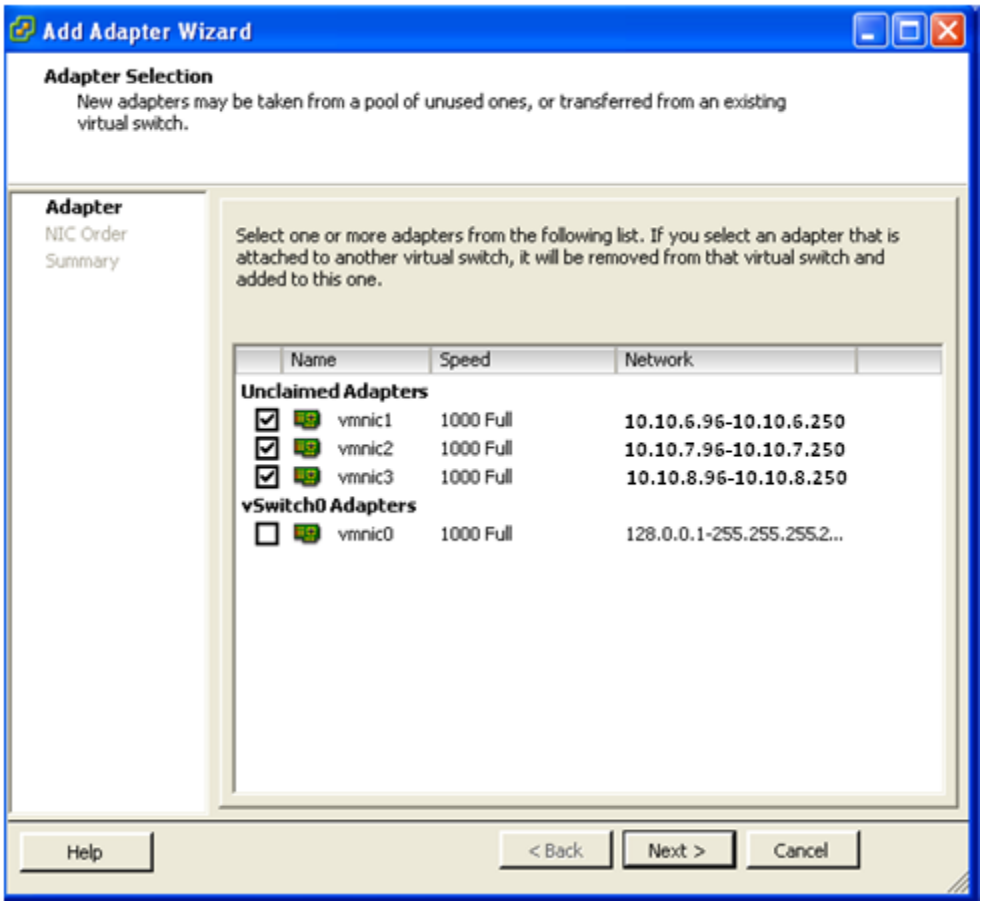


**Figure 6: Assigning Network Adapters (continued)**

Click ***Next*** after you have selected the chosen adapters. For now keep the defaults listed in the Failover Order screen and click ***Next***. Review the adapters listed and click ***Finish*** completing the process.

These adapters will now show up in the GUI under the Network Adapters tab.

## STEP4: ASSOCIATE VMKERNEL PORTS TO PHYSICAL ADAPTERS

The next step is used to create the individual path bindings for each VMkernel to a NIC. This is required in order to take advantage of features such as Most Recently Used(MRU) and Round Robin(RR) multipathing policies or 3rd party MPIO plug-ins currently available from Dell.

From our previous step there are 3 VMkernel ports and 3 NICs. This means that each NIC will have 1 VMkernel ports assigned to it. Again, each environment will differ and these numbers can change based on the number of NICs and the number of paths assigned.

This process can be done either via vCLI or through the vCenter GUI.

By default, all the vmnics are assigned to each VMkernel port. We need to remove all but one vmnic from each VMkernel port so that each VMkernel port has only one uplink.

Before running these commands the switch information looks like the following (obtained using *esxcfg-vswitch –l* again):

| Switch Name | Num Ports | Used Ports | Configured Ports | MTU | Uplinks |
|---|---|---|---|---|---|
| vSwitch2 | 64 | 7 | 64 | 9000 | |

| PortGroup Name | VLAN ID | Used Ports | Uplinks |
|---|---|---|---|
| iSCSI3 | 0 | 1 | vmnic1,vmnic2,vmnic3 |
| iSCSI2 | 0 | 1 | vmnic1,vmnic2,vmnic3 |
| iSCSI1 | 0 | 1 | vmnic1,vmnic2,vmnic3 |

You can see that there are three vmnics in each uplink for each VMkernel Port. This is what we need to change so that only a single vmnic is in each uplink and that we manually load balance them across all available VMkernel Ports.

To configure this process via vCLI first note the vmnic number of a NIC you want to remove and type the following command:

esxcfg-vswitch –p iSCSI1 –N vmnic3 vSwitch2

What this will do is remove vmnic3 from VMkernel port iSCSI1 so that now vmnic1 and vmnic2 are left on iSCSI1. We then need to remove vmnic2 so that only vmnic1 is associated with the iSCSI1. To do this type the following command:

esxcfg-vswitch –p iSCSI1 –N vmnic2 vSwitch2

Now that we have just one vmnic associated with one VMkernel port we need to remove the excess NICs on the other ports.

esxcfg-vswitch –p iSCSI2 –N vmnic1 vSwitch2

esxcfg-vswitch –p iSCSI2 –N vmnic3 vSwitch2

esxcfg-vswitch –p iSCSI3 –N vmnic1 vSwitch2

esxcfg-vswitch –p iSCSI3 –N vmnic2 vSwitch2

To verify that this was done correctly type the following command:

esxcfg-vswitch –l

The output will look similar to this:

| Switch Name | Num Ports | Used Ports | Configured Ports | MTU | Uplinks |
|---|---|---|---|---|---|
| vSwitch2 | 64 | 7 | 64 | 9000 | |

| PortGroup Name | VLAN ID | Used Ports | Uplinks |
|---|---|---|---|
| iSCSI3 | 0 | 1 | vmnic3 |
| iSCSI2 | 0 | 1 | vmnic2 |
| iSCSI1 | 0 | 1 | vmnic1 |

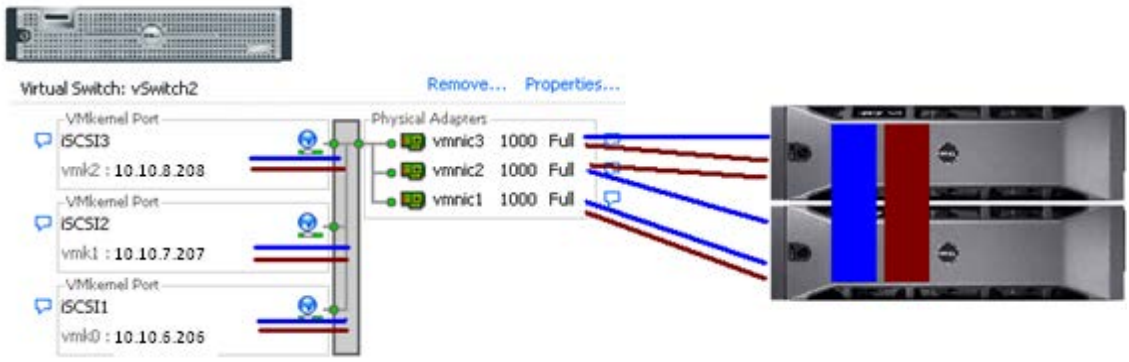


**Figure 7: Associating VMKernel ports**

The important thing to note is that under the Uplinks section there is only one vmnic assigned to each iSCSI VMkernel port and that they are evenly distributed across them.

This can also be done through the vCenter GUI. To configure this from the GUI first navigate to the Networking section on the ESX host you are configuring. ***Configuration -> Networking***.

From here, click ***Properties*** on the vSwitch2.

Select one of the VMkernel Ports, in this example iSCSI1 and click ***Edit***.

From here select the ***NIC Teaming*** tab.

Here you are going to select the check box for ***Override vSwitch failover order***.

Just like in the vCLI example we will assign vmnic1 to iSCSI1. This is done by selecting the adapters that are not going to be assigned to the VMkernel (vmnic2 and vmnic3 in this case) and clicking the ***Move Down*** button until it is listed under Unused Adapters. The

following figure shows the completed result. Click ***Ok*** to complete the process. Do this same thing for each of the iSCSI VMkernel ports so that each VMkernel port is mapped to only one adapter and they are all balanced. In this example we assigned iSCSI1to vmnic1, iSCSI2 to vmnic2 and iSCSI3 to vmnic3. See also VMware's KB discussing reported issues with oversubscribing VMkernel ports.

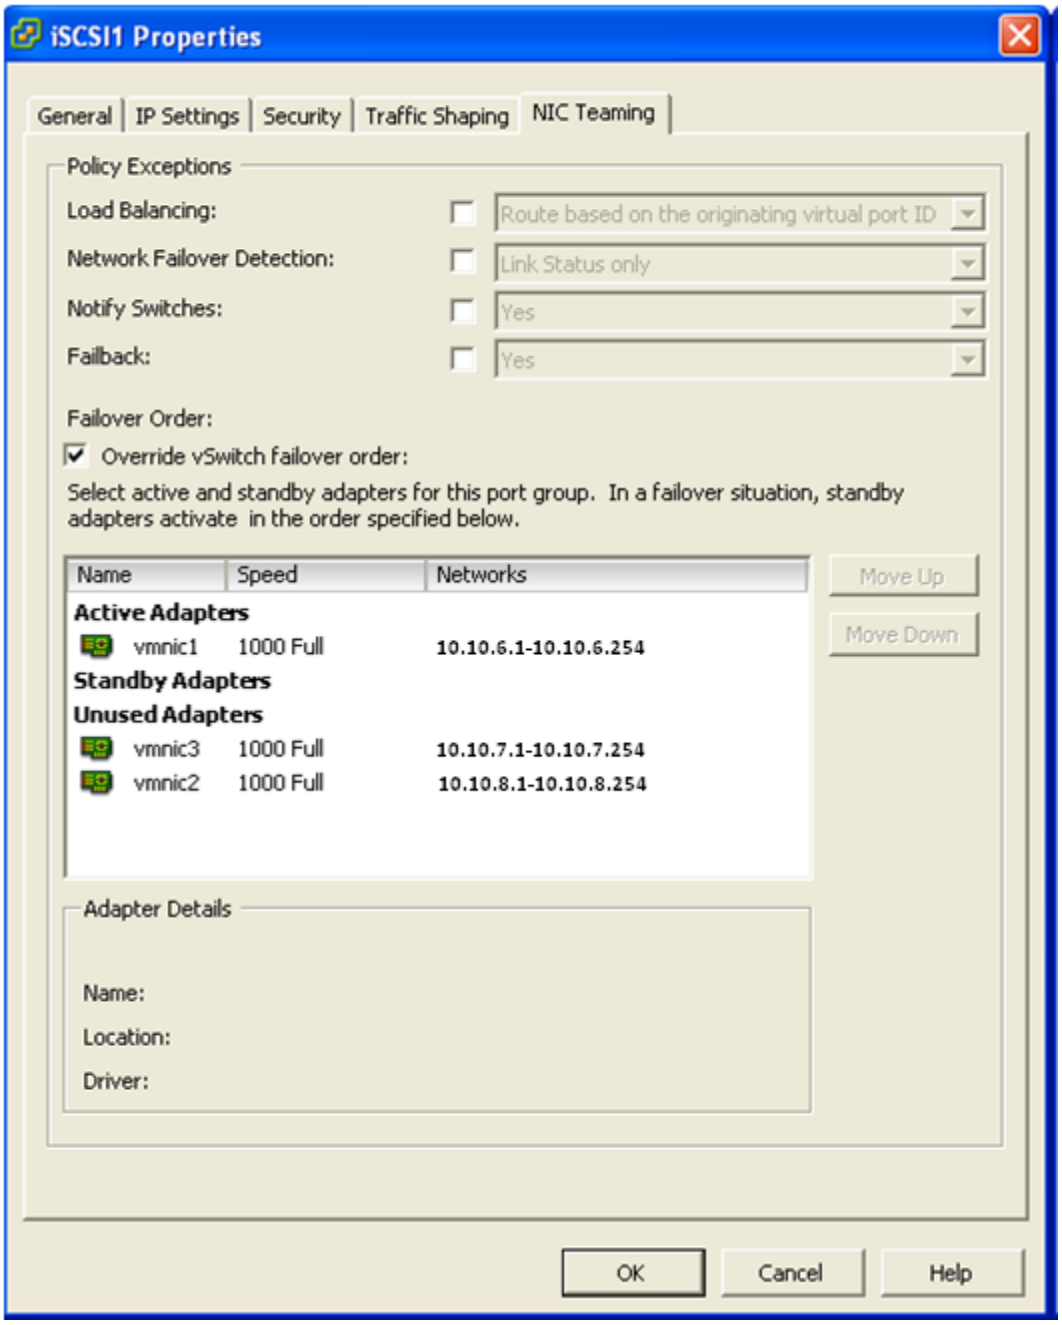


**Figure 8: Associating VMKernel ports (continued)**

## STEP5: ADD AN MD3600I INTO THE ESX RULE SET

To use MD3600i with ESX 4.1 and ESX 4.1 Update 1, the MD3600i PID has to be added to the manageable device list for VMW_SATP_LSI module. Otherwise, the guest host will not be able to see the mapped volumes and perform I/O's. The following command needs to be run from the command line on each ESX host before mapping virtual disks to ESX guest OS's. The setup of MD3600i vs an MD3200i is identical through Step4 of the procedure above.

At Step6 below, and prior to enabling the iSCSI initiator to prepare the ESX host to connect to the PowerVault SAN, on the ESX host you need to run

#/*example to add MD 3600i VID/PID (following the MD3200i format)*/

#esxcli nmp satp addrule -V DELL -M MD3600i -s VMW_SATP_LSI

This command adds MD3600i devices into the device list for VMW_SATP_LSI module.

Then proceed with Step6 and enable the iSCSI Software Initiator on ESX hosts.

## STEP6: ENABLE VMWARE ISCSI SOFTWARE INITIATOR

This step enables the iSCSI initiator to prepare the ESX host to connect to the PowerVault SAN. This can be done either through a vCLI command or through the vCenter GUI.

To enable the iSCSI initiator through the vCLI type the following command:
esxcfg-swiscsi –e

This will enable the software iSCSI initiator. To verify that it is enabled type the following command:

esxcfg-swiscsi –q

This can also be accomplished by using the vCenter GUI.

From the vCenter GUI on the ESX host navigate to ***Configuration -> Storage Adapters***. Select the iSCSI Software Adapter and click ***Properties***.

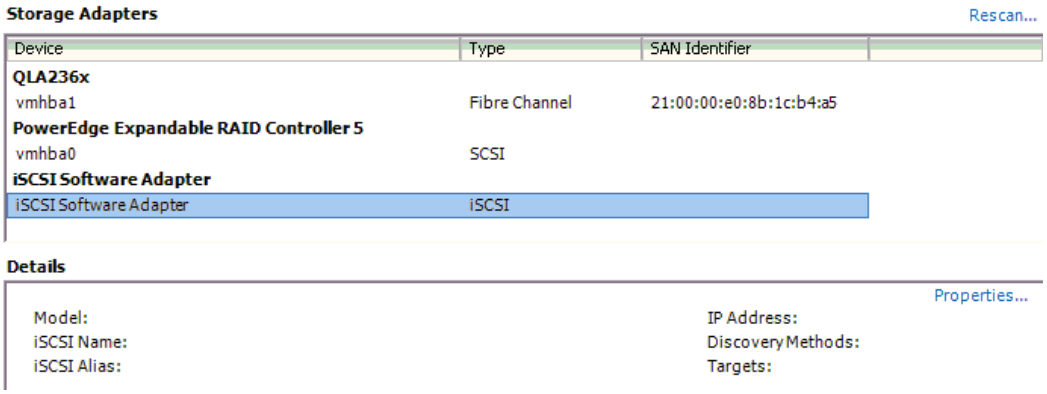


When the iSCSI properties window opens under the General tab click the ***Configure*** button. Place a check mark in ***Enabled*** and hit ***Ok***. This will enable the iSCSI initiator and assign a unique iqn to the ESX host. Administrators familiar with enabling iSCSI in ESX 3.5 will notice that the firewall policies are automatically set when you enable it in vSphere4.

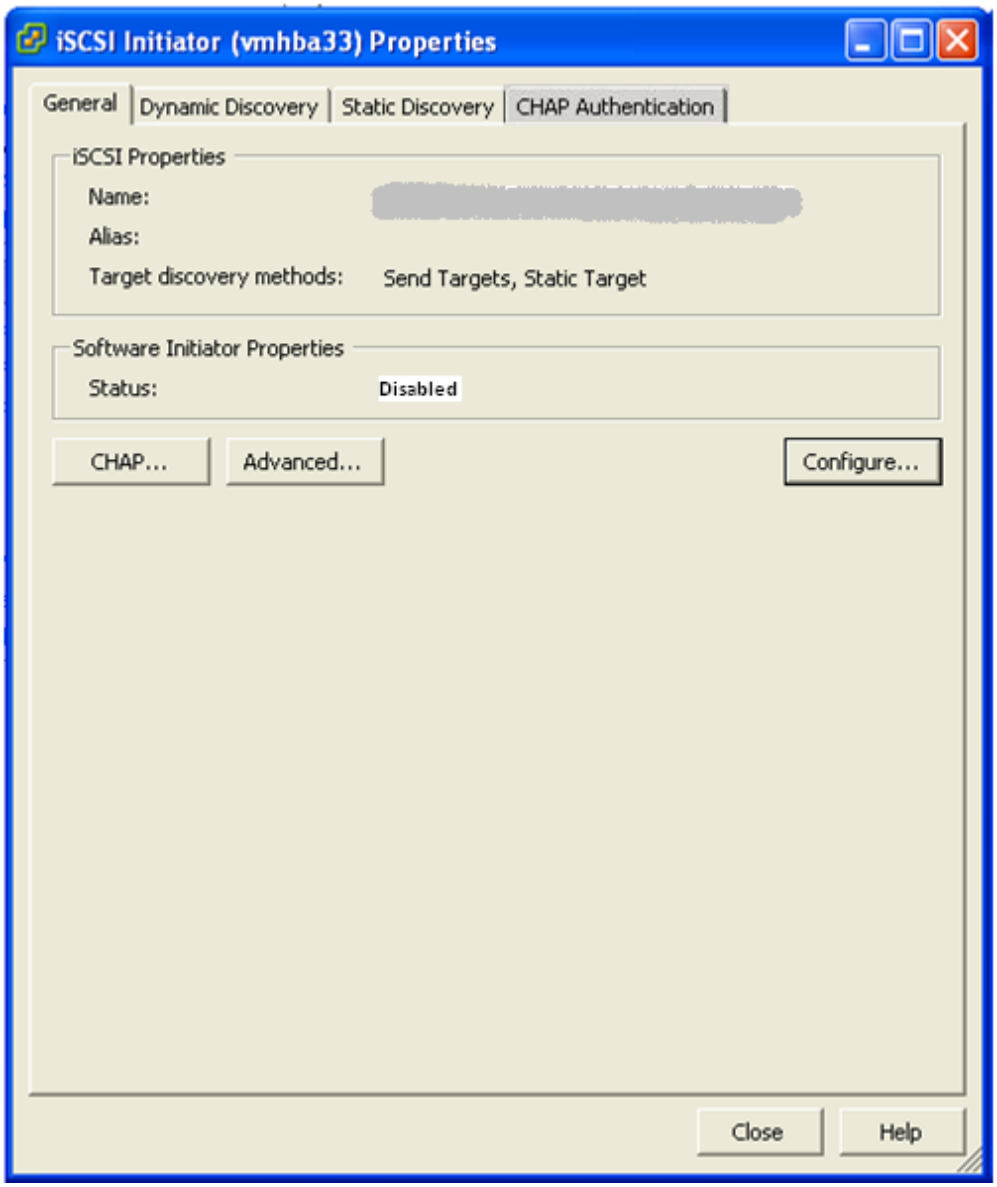


**Figure 9: Enabling iSCSI Software Initiator**

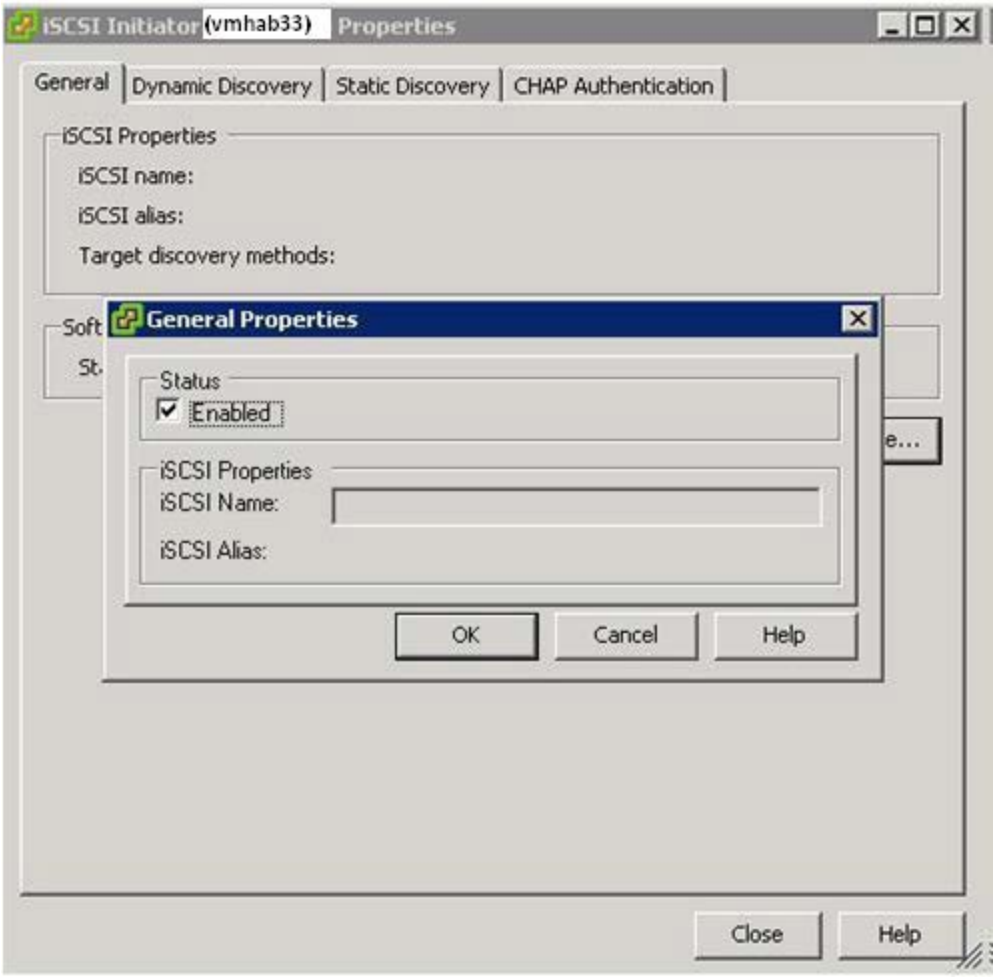


**Figure 10: Enabling iSCSI Software Initiator (continued)**

## STEP7: BINDING VMKERNEL PORTS TO ISCSI SOFTWARE INITIATOR

This next step will bind the VMkernel ports, which were configured in Step 4 earlier, to the iSCSI Software Initiator. If this step is skipped there will only ever be a single connection to the PowerVault SAN. This step must be done via vCLI.

The first thing to do is to note the vmhba# of the iSCSI Software Initiator. This can be seen in the vCenter GUI on the ESX host under ***Configuration -> Storage Adapters***.

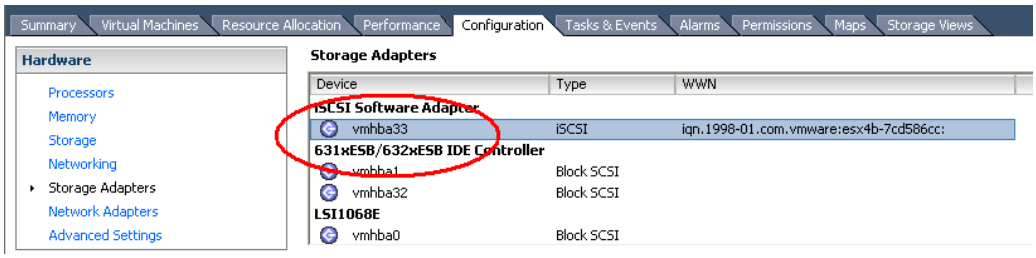


**Figure 11: Binding VMKernel Ports**

This can also be found by running the following vCLI command to discover all SCSI devices including the iSCSI software adapter:

esxcfg-scsidevs –a

The output will look something like the following:

vmhba0 mptsas link-n/a sas.5001ec90e0ba7c00
(1:0.0) LSI Logic / Symbios Logic LSI1068E
vmhba1 ata_piix link-n/a ide.vmhba1
(0:31.1) Intel Corporation 631xESB/632xESB IDE Controller
vmhba32 ata_piix link-n/a ide.vmhba32
(0:31.1) Intel Corporation 631xESB/632xESB IDE Controller
vmhba33 iscsi_vmk link-n/a iscsi.vmhba33
() Software iSCSI

In this example from both the GUI and vCLI we can determine that the vmhba# for the iSCSI Software Initiator is vmhba33. This will be used in the next part. This will differ on various systems based on the devices installed.

The next piece of information to gather is the vmk# of each of the VMkernel ports. This can be done via the GUI or vCLI.

To determine the vmk# of each VMkernel port from the vCenter GUI navigate to ***Configuration -> Networking***. From the vSwitch that was created earlier under each VMkernel port, the vmk# will be listed.

*NOTE: It is important to recognize that they may not start with vmk0, vMotion ports and other VMkernels will utilize the same numbers based on the order they are created.*

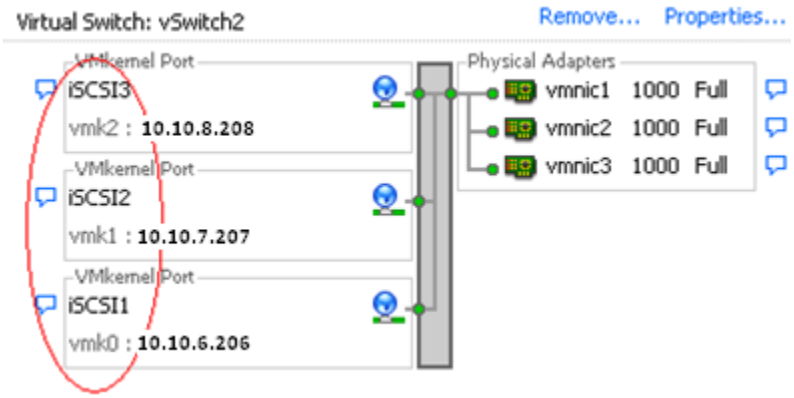


**Figure 12: Binding VMKernel Ports (continued)**

In this example we see that iSCSI1 is vmk0, iSCSI2 is vmk1, and iSCSI3 is vmk2. This is also information that we need to note.

We can also see this in the vCLI by using the following command:

esxcfg-vmknic –l

The output will look similar to this:

Interface Port Group/DVPort IP Family IP Address
Netmask Broadcast MAC Address MTU TSO MSS
Enabled Type
vmk0 iSCSI1 IPv4 10.10.6.206
255.255.255.0 10.10.5.255 00:50:56:7b:d8:3e 9000 65535 true
STATIC
vmk1 iSCSI2 IPv4 10.10.7.207
255.255.255.0 10.10.5.255 00:50:56:7e:ae:80 9000 65535 true
STATIC
vmk2 iSCSI3 IPv4 10.10.8.208
255.255.255.0 10.10.5.255 00:50:56:74:a4:e0 9000 65535 true
STATIC

We can determine this same information from the GUI.

Now that we know the vmhba# and the vmk# we can map each VMkernel Port to the iSCSI Software Initiator. This is done through the vCLI by typing the following command:

esxcli swiscsi nic add –n vmk0 –d vmhba33

This will bind the vmk0 VMkernel port to the iSCSI Software Adapter vmhba33. We then proceed to bind all of the other vmk# to the same vmhba.

esxcli swiscsi nic add –n vmk1 –d vmhba33

esxcli swiscsi nic add –n vmk2 –d vmhba33

To verify that all of the vmk# are bound properly to the vmhba run the following command:

esxcli swiscsi nic list –d vmhba33

This will list all of the information for VMkernel ports that are assigned to the iSCSI Software Adapter.

## STEP8: CONNECT TO POWERVAULT MD3200I STORAGE

Now that the advanced configuration for the vSphere4 iSCSI Software Initiator has been completed, the next stage is to connect to the Dell PowerVault SAN.

More information for complete administration of the Dell PowerVault SAN can be found in the PowerVault User's Guide. In this example we will attach the iSCSI Software

Initiator to the SAN. We will skip configuring CHAP but both one way and bi-directional CHAP is supported by the PowerVault SAN.

Navigate in the vCenter GUI to ***Configuration -> Storage Adapters***.
You should now see your iSCSI Target name listed.

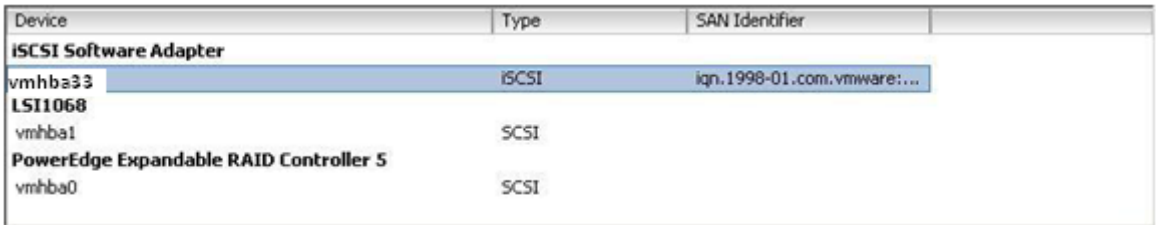


**Figure 13: Connecting to PowerVault Storage**

Click on the iSCSI Software Adapter and click ***Properties***.

In the iSCSI Server box type an IP Address for one of the iSCSI host ports of the PowerVault SAN and hit ***Ok***. There may be a slight delay before the process completes.

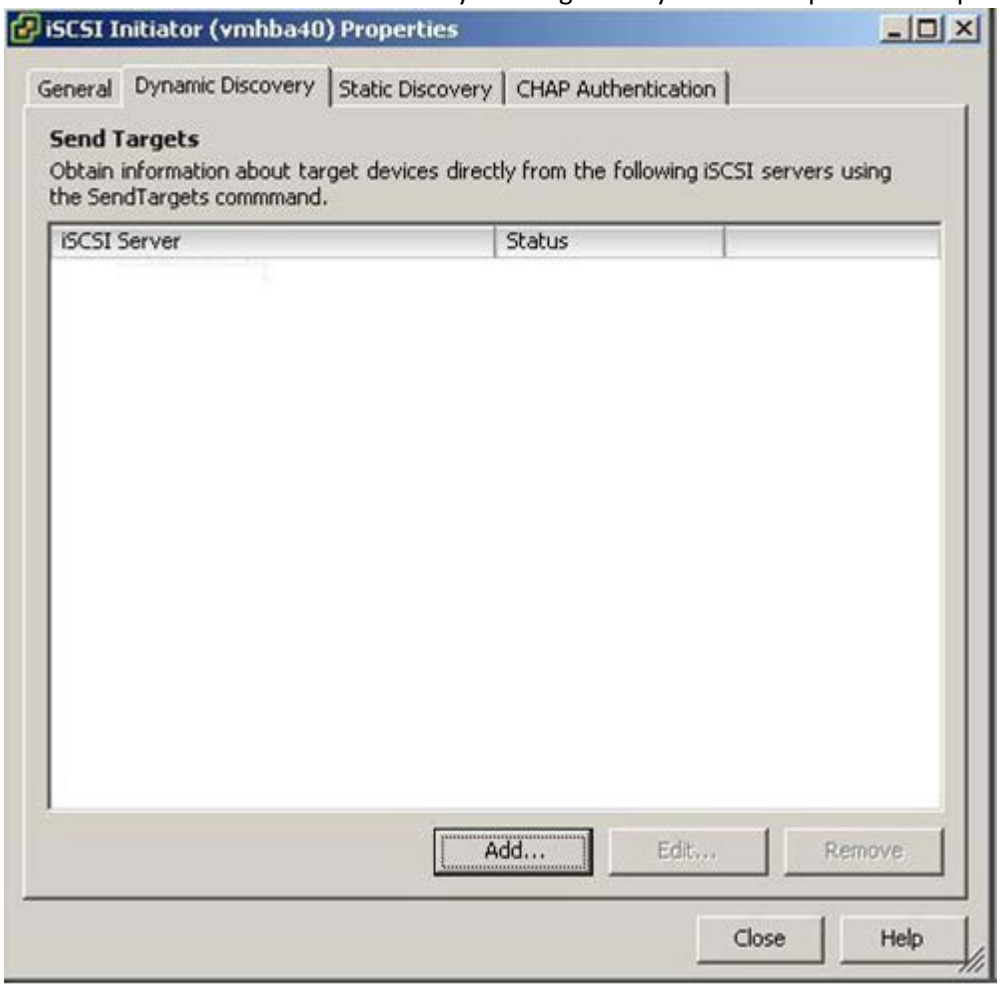

**Figure 14: Dynamic Discovery**

Click the ***Dynamic Discovery*** tab.

Click ***Add***.

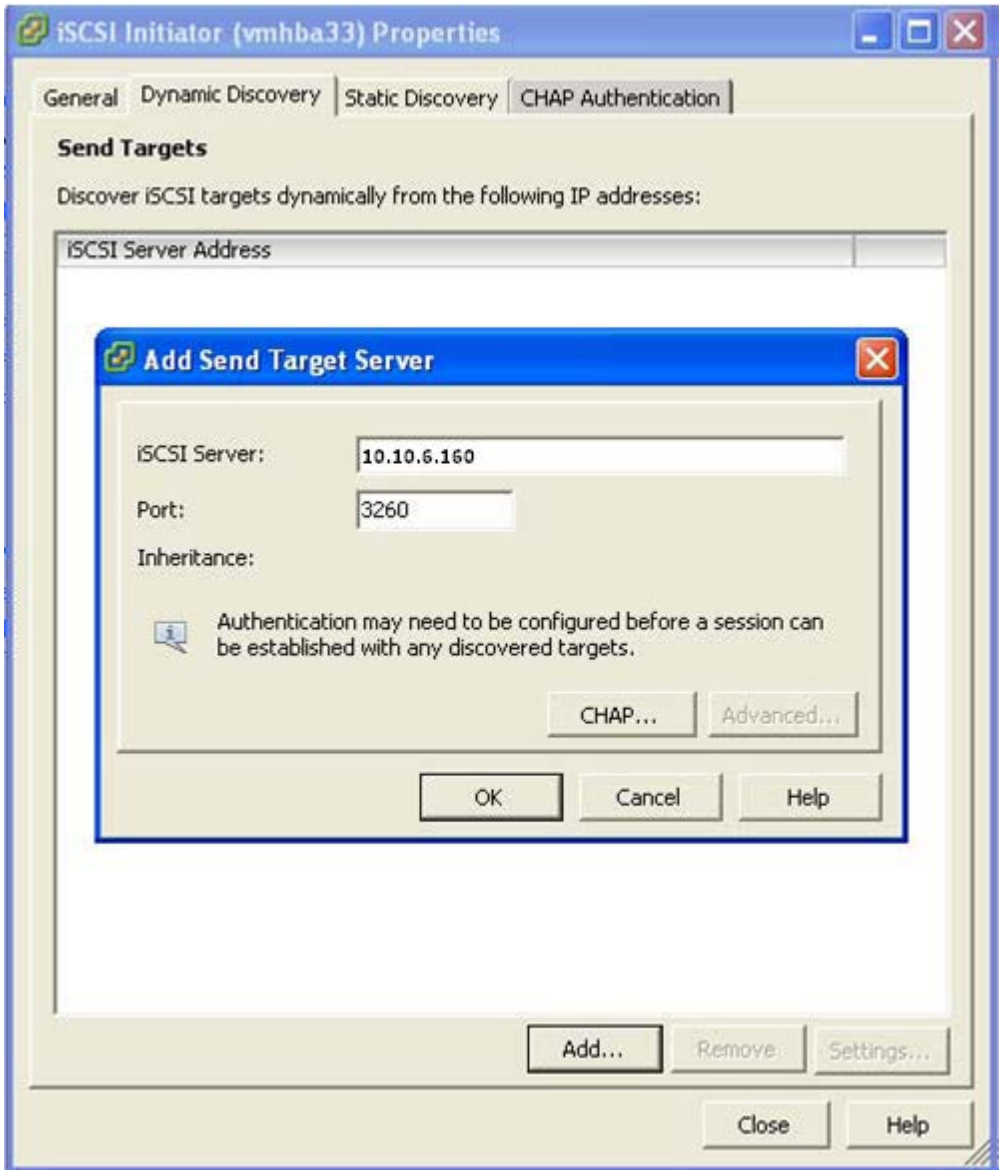


**Figure 15: Dynamic Discovery (continued)**

When this is done click ***Close*** or enter in another IP Address if there are multiple SANs in your environment.

***Note: Port binding requires that all target ports of the storage array must reside on the same broadcast domain as the VMkernel ports because routing is not supported with port binding. See VMware KB #2017084 here.***

You will be prompted for a Rescan of the HBAs but at this time as there are no virtual disks assigned it is unnecessary.

The next step will be to create a new volume and assign it to the ESX server. This can be done multiple ways so refer to the PowerVault User's Guide for more information. In this example we will create a 100GB virtual disk and assign it to this ESX host using the iqn name.

## STEP9: POWERVAULT MD3200I STORAGE SETUP AND CONFIGURATION

Create virtual disks on MD3200i using steps described in:

http://support.dell.com/support/edocs/systems/md3200i/multlang/GSG/DAO_BCC/GSG.pdf.

Note: in this example the Storage array is an MD3200i with virtual disks already configured using the *Configure Storage* Array selection under the Setup Tab. The new ESX host being added is named "VMware_host1".

After opening the Modular Disk Storage Manager and selecting the MD3200i storage array to be configured, select the *Mappings* tab.

Manually Define Hosts *by* highlighting the Storage Array Name and right clicking. Select Define -> Host.

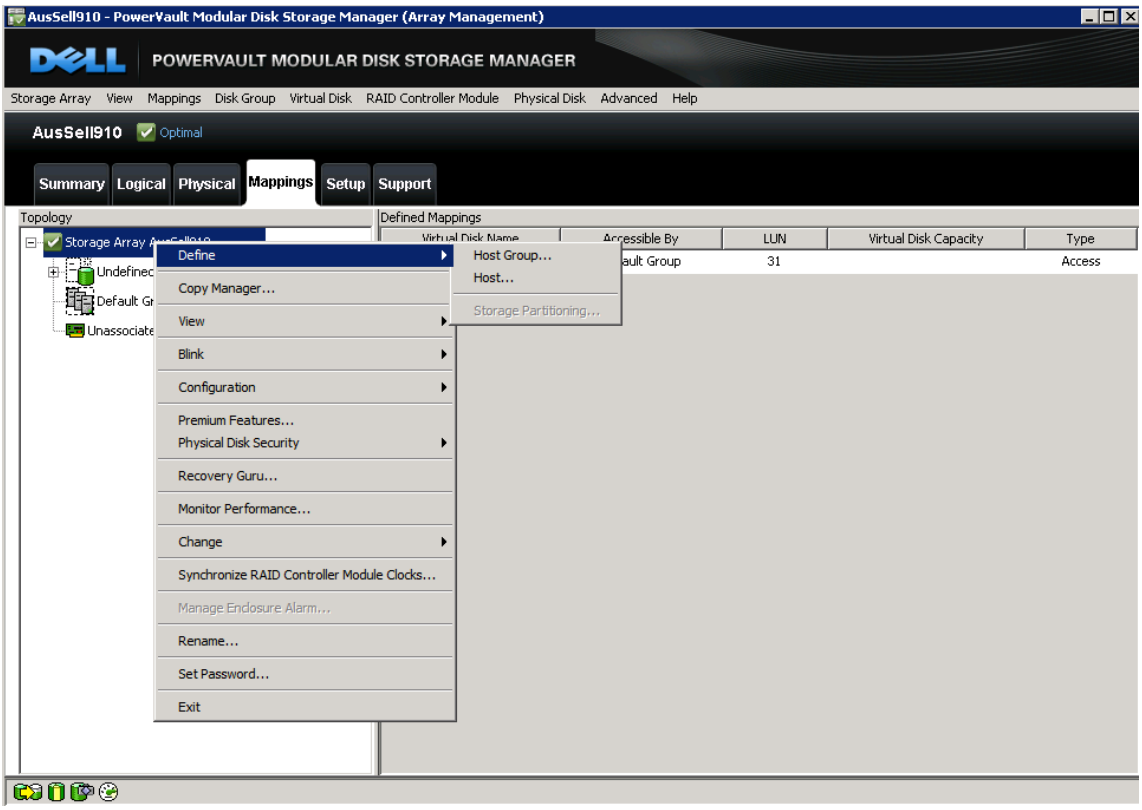


**Figure 16: Defining the Storage Array Host Topology**

Naming the host

Select a name that matches the naming convention used for the environment that you are configuring, for example, VMware_host1. Leave partitions enabled and select Next.

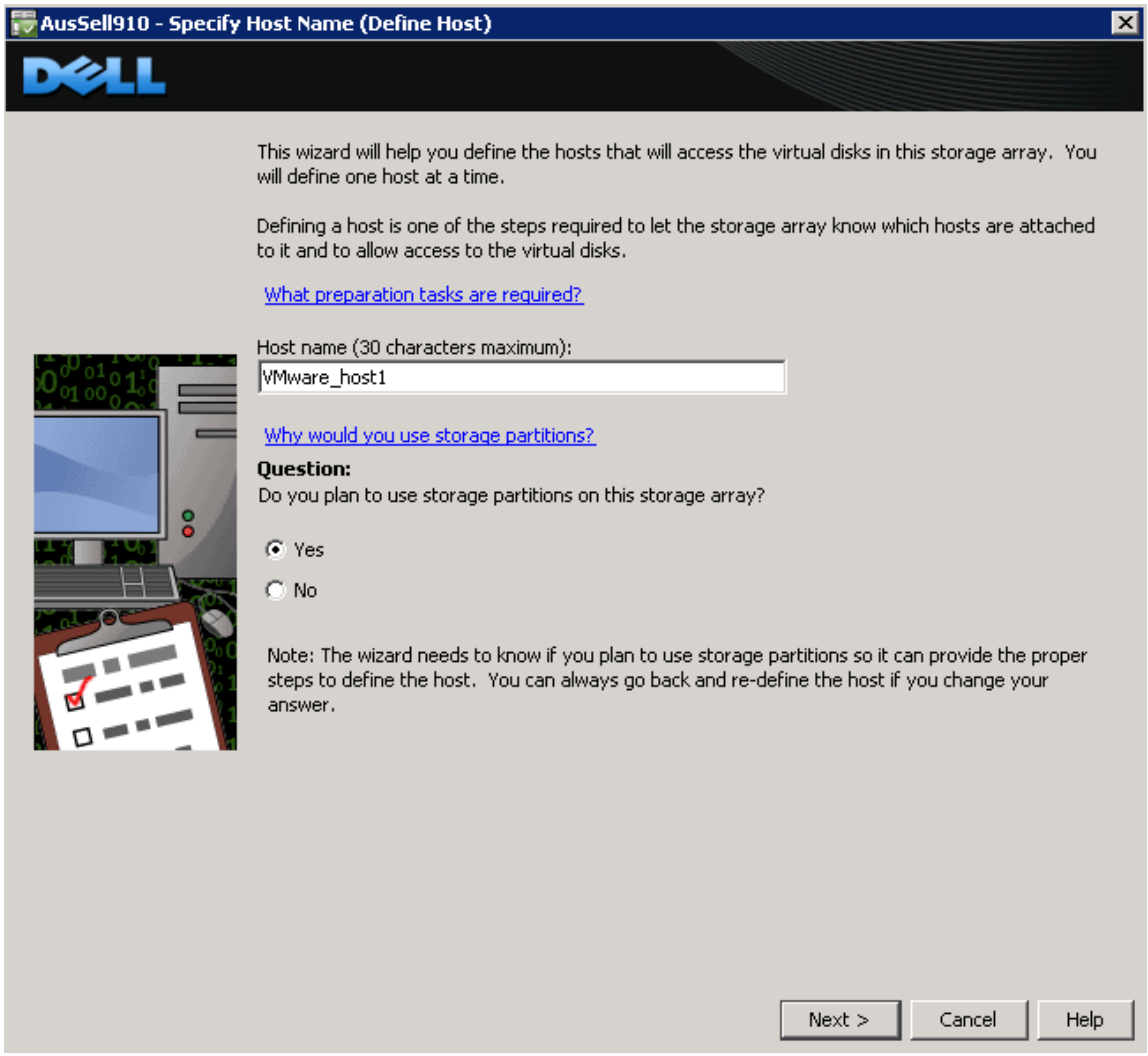


**Figure 17: Naming the Host**

Selecting host port identifiers

Using the pull down select the host port identifier. In this example because the iSCSI session has already been established the iqn is already in the list.

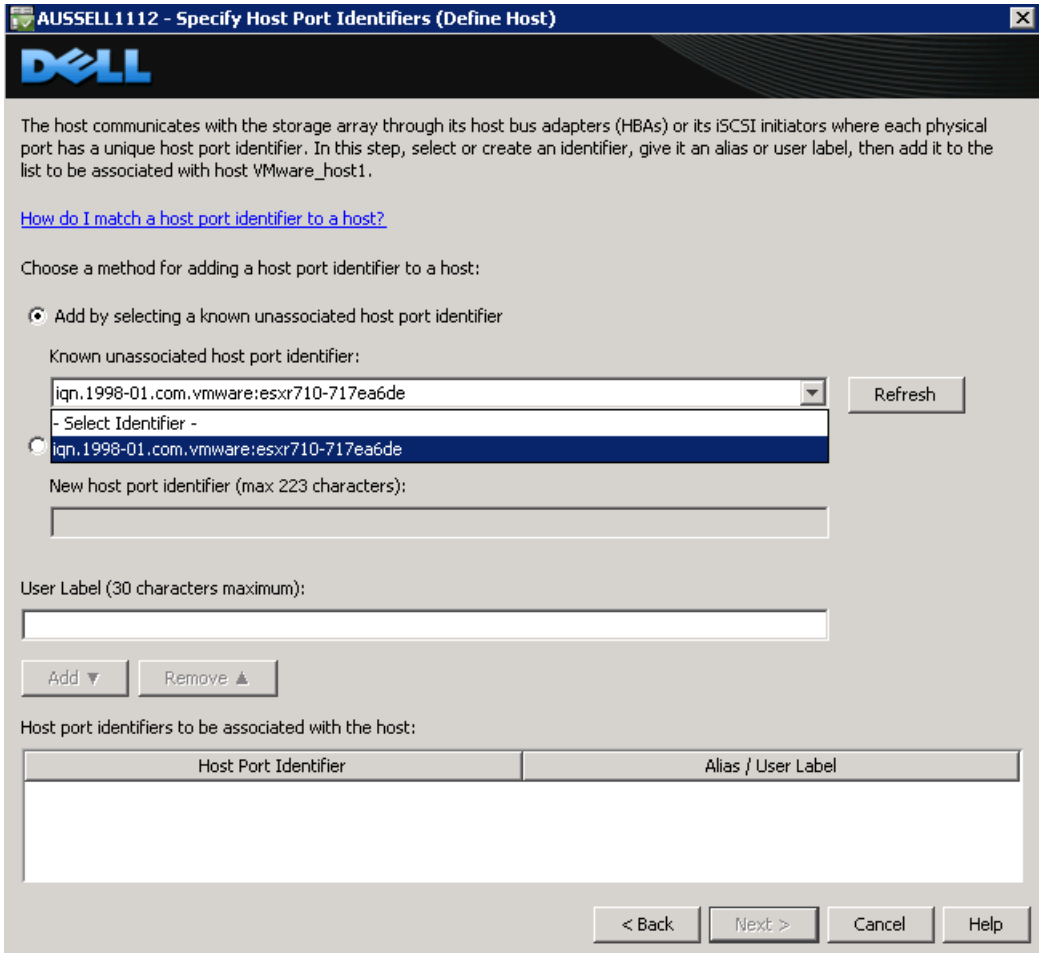


**Figure 18: Selecting Host Port Identifiers**

Note: You can also add the iqn manually. To find the iSCSI Initiator Name from the vCenter GUI go to ***Configuration -> Storage Adapters***. Click on the iSCSI Software Adapter. The iqn can be copied and pasted into the add by creating a host port identifier field.

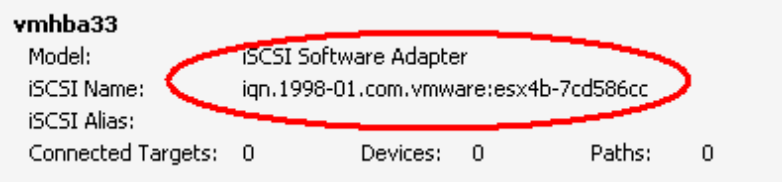

Entering host port alias

An alias is used in the topology tree to identify the port. Add a unique alias for this host port such as the one below, then select ***Add***

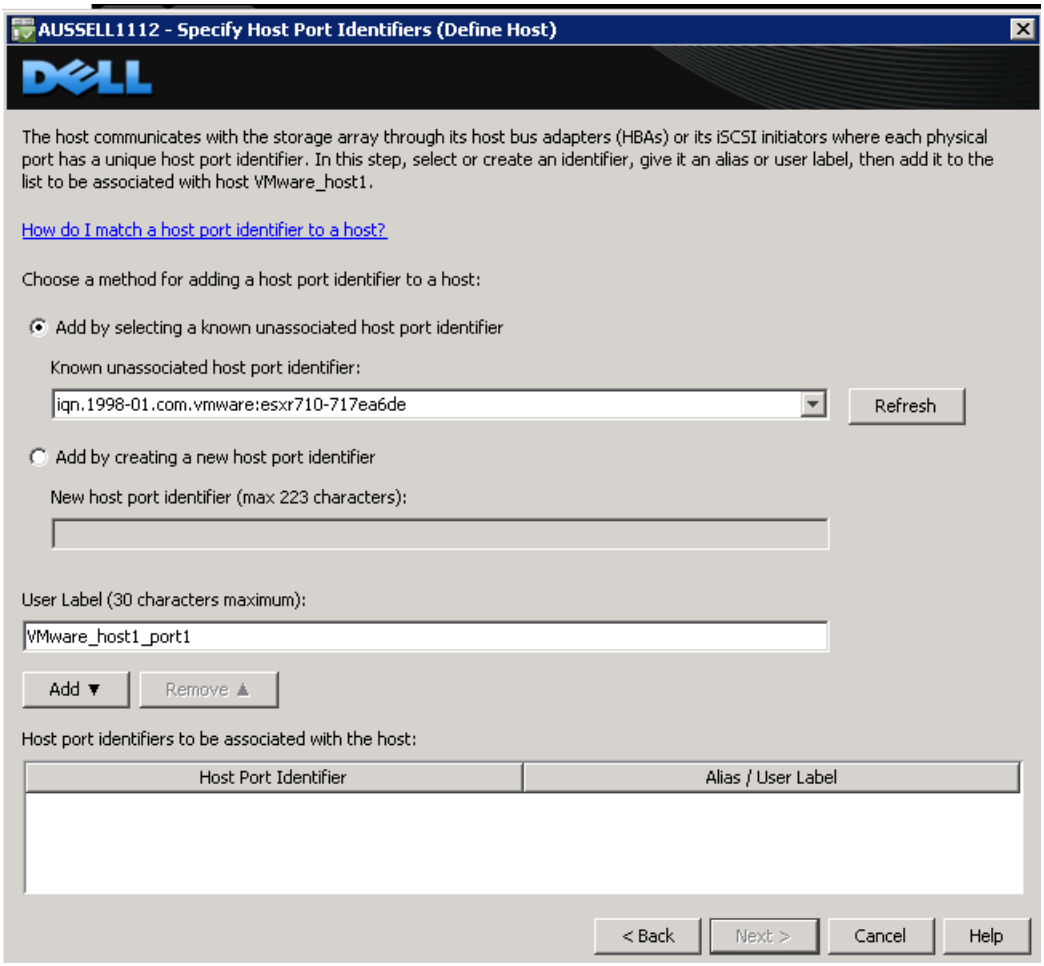


**Figure 19: Adding Host Port Alias**

The completed host port screen will be similar to the one below. Select ***Next*** to continue.

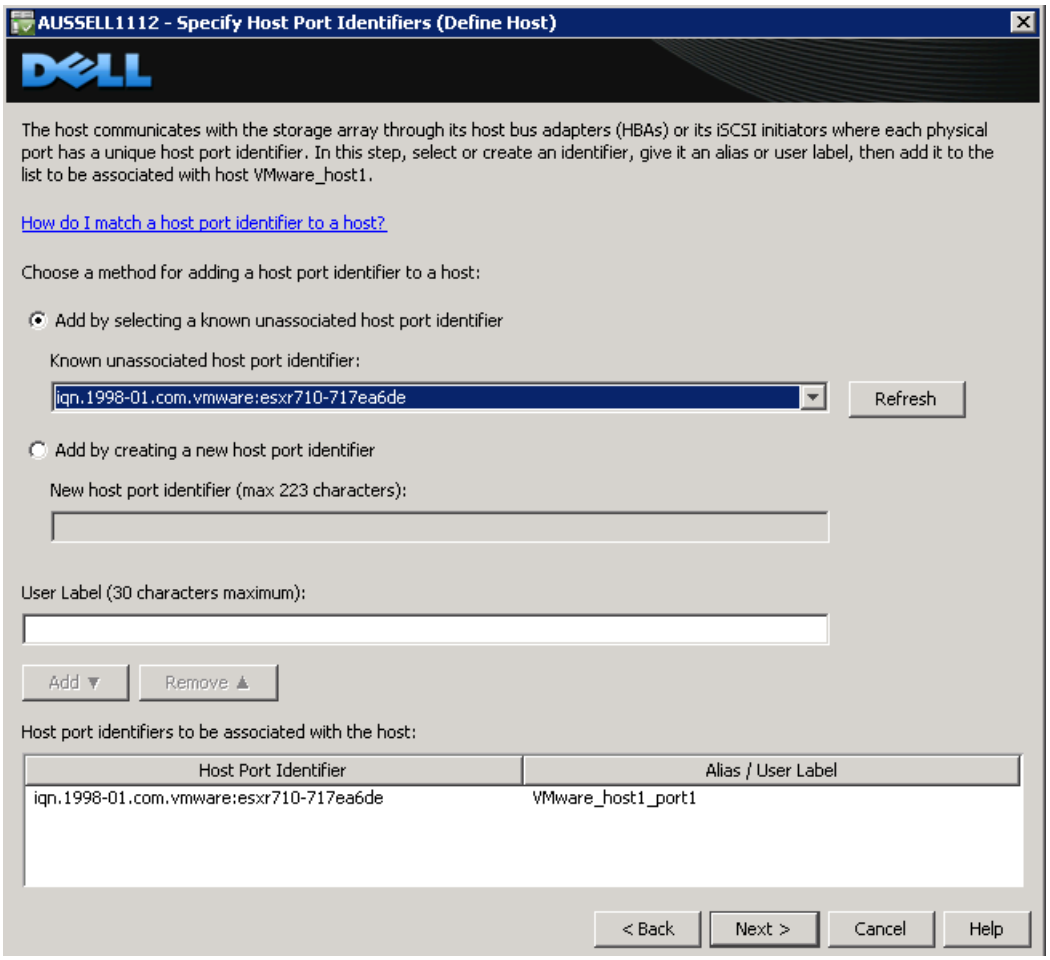


**Figure 20: Completed Host Port Identifier**

Select ***VMware*** as the host type

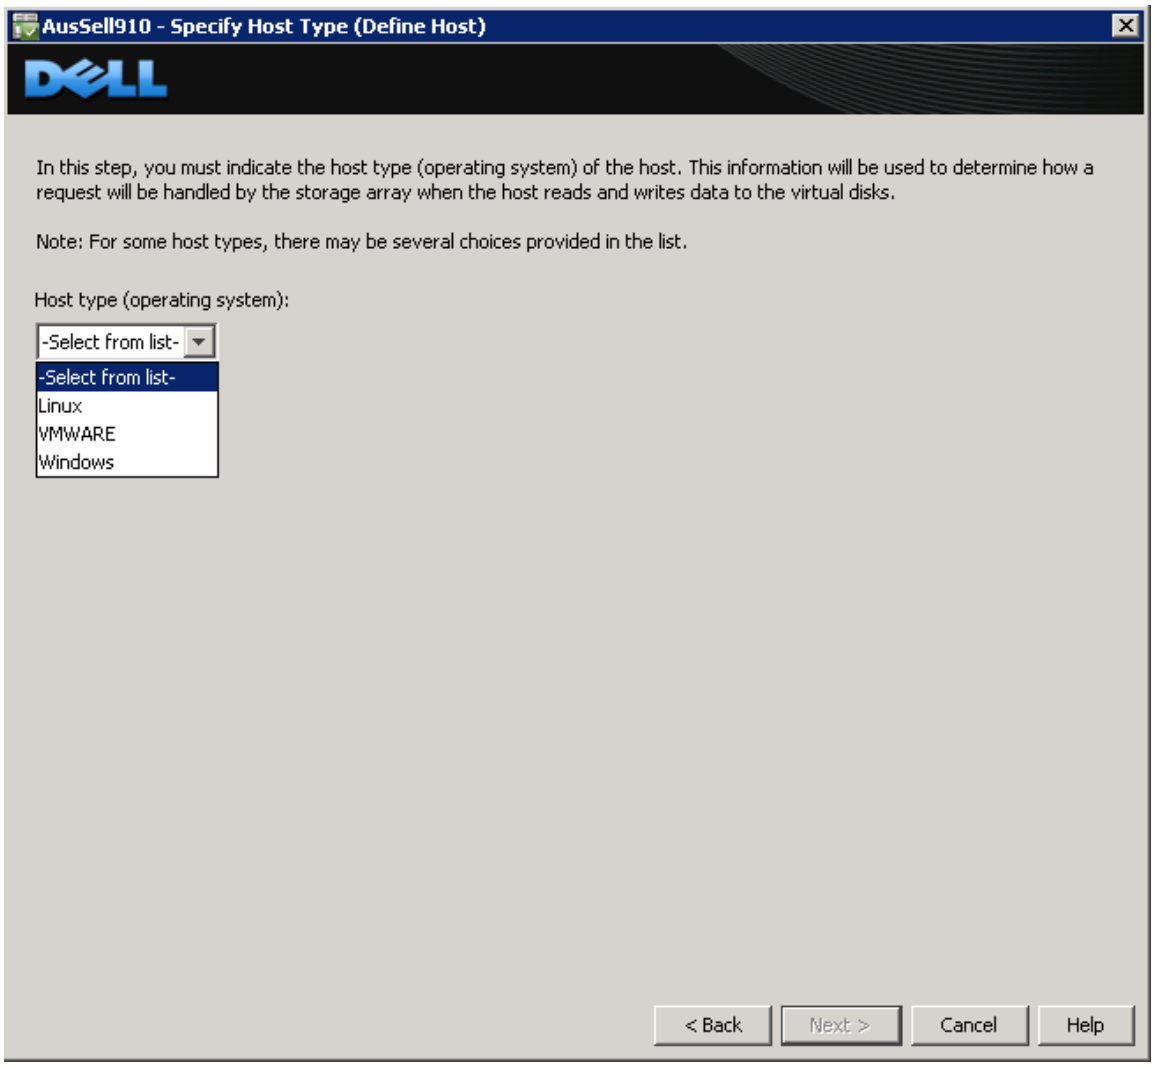


**Figure 21: Selecting Host Type**

If you intend to use advanced VMware features such as VMotion then this host will share access with other ESX servers and you will have to create a Host Group. We will create a host group for this example.

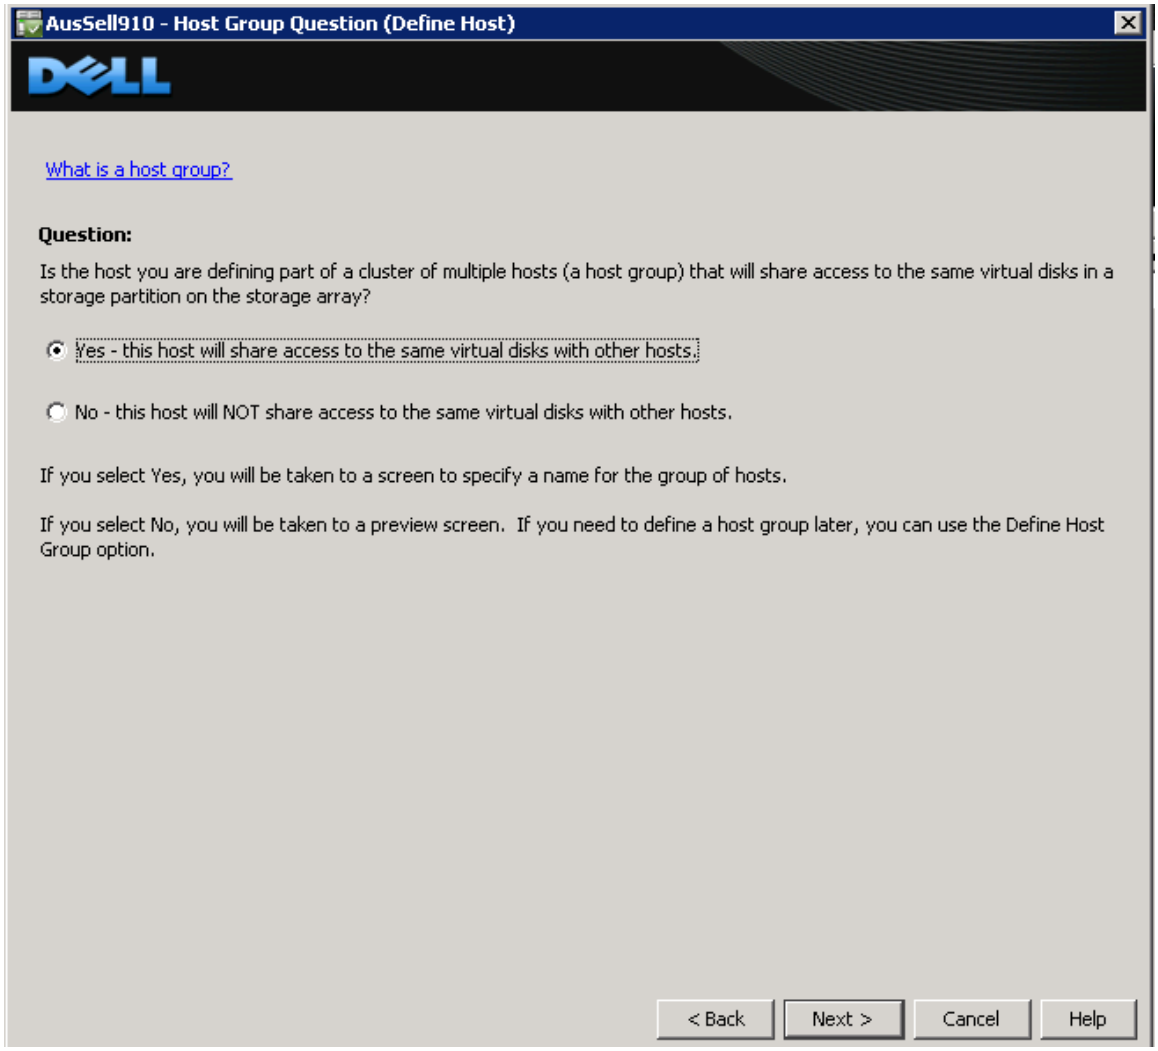


**Figure 22: Host Group Question**

Host Group name

Enter a host group name that is appropriate for your environment. For this example we used VMware_Group1

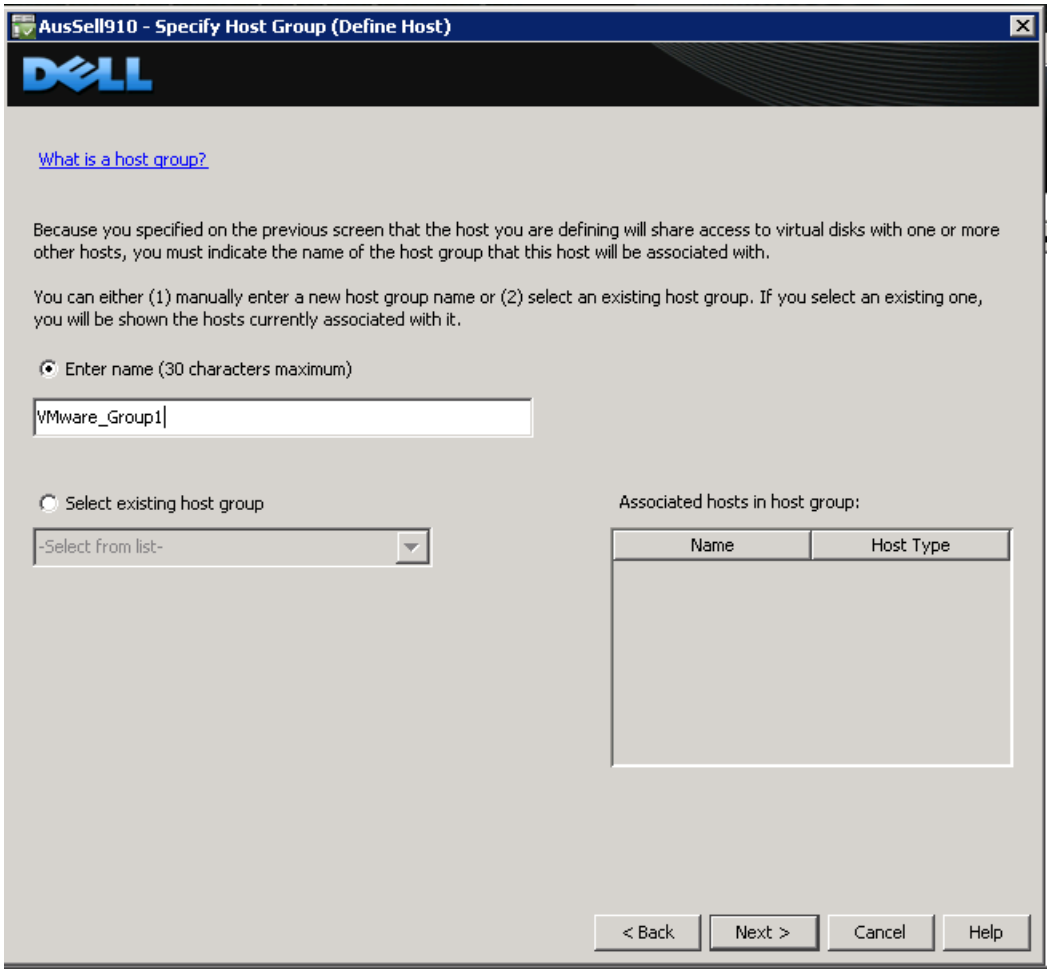


**Figure 23: Host Group Name**

Preview Define host

If all of the information is correct for your environment select Finish.

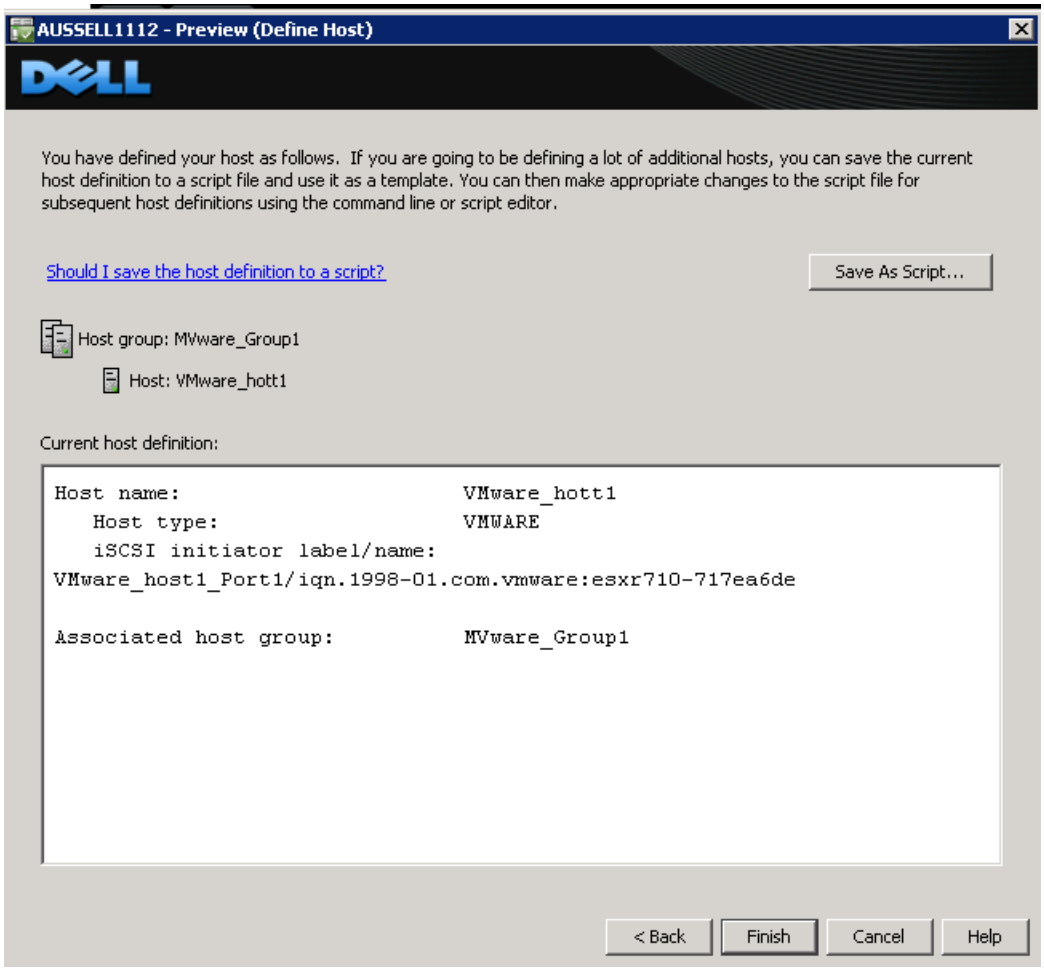


**Figure 1 Preview (Define Host)**

Creation successful

Select No at this time. You can add additional Hosts after you have finished configuring the current host.

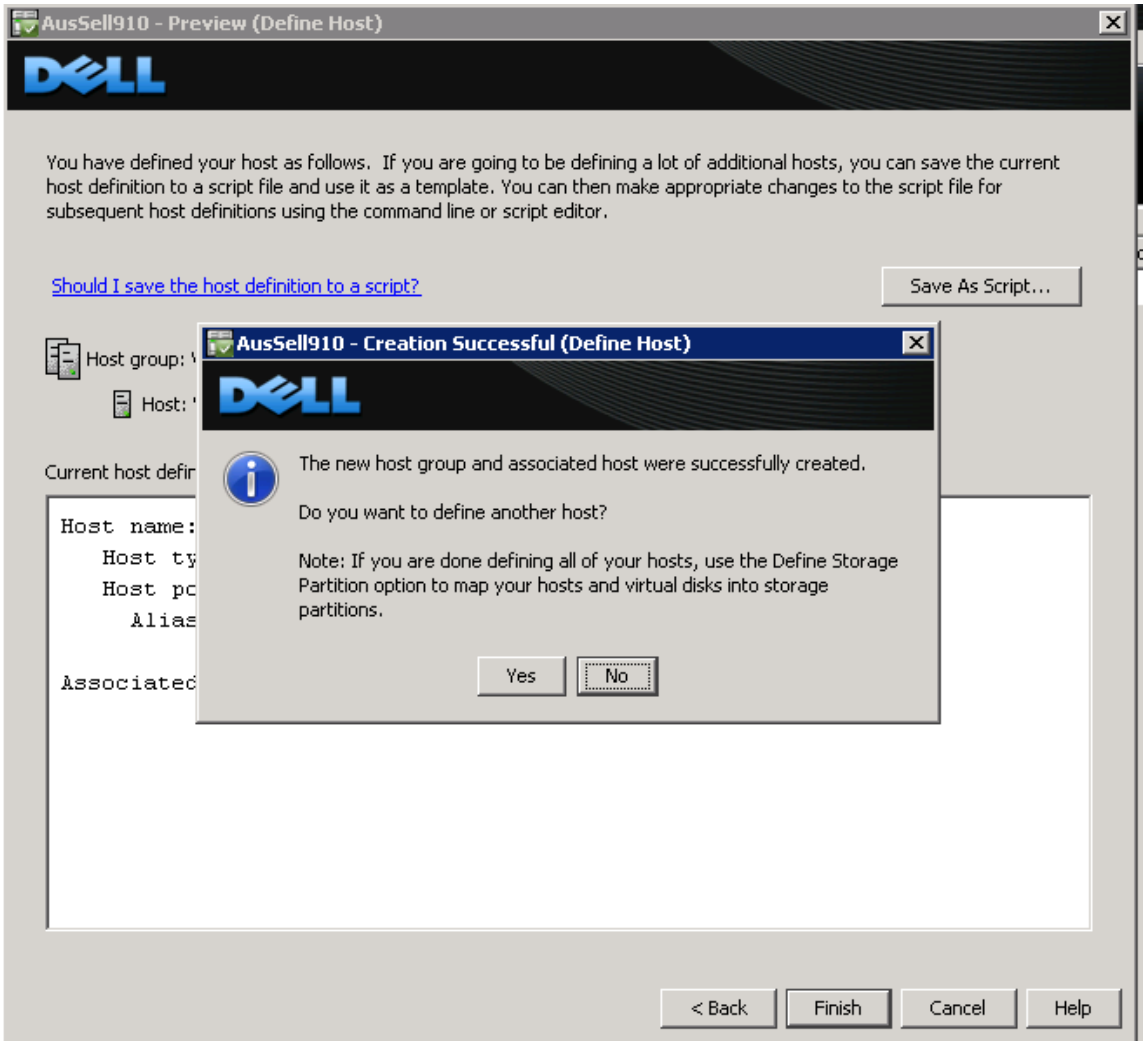


**Figure 2 Topology Creation Successful**

Define Mappings for LUNs

Note: In this example the Disk Groups and Virtual Disks have already been created using the wizard under the Setup Tab.

In the topology tree expand the Undefined Mappings and highlight one of the Virtual Disks. Right Click and select *Define Additional Mappings*.

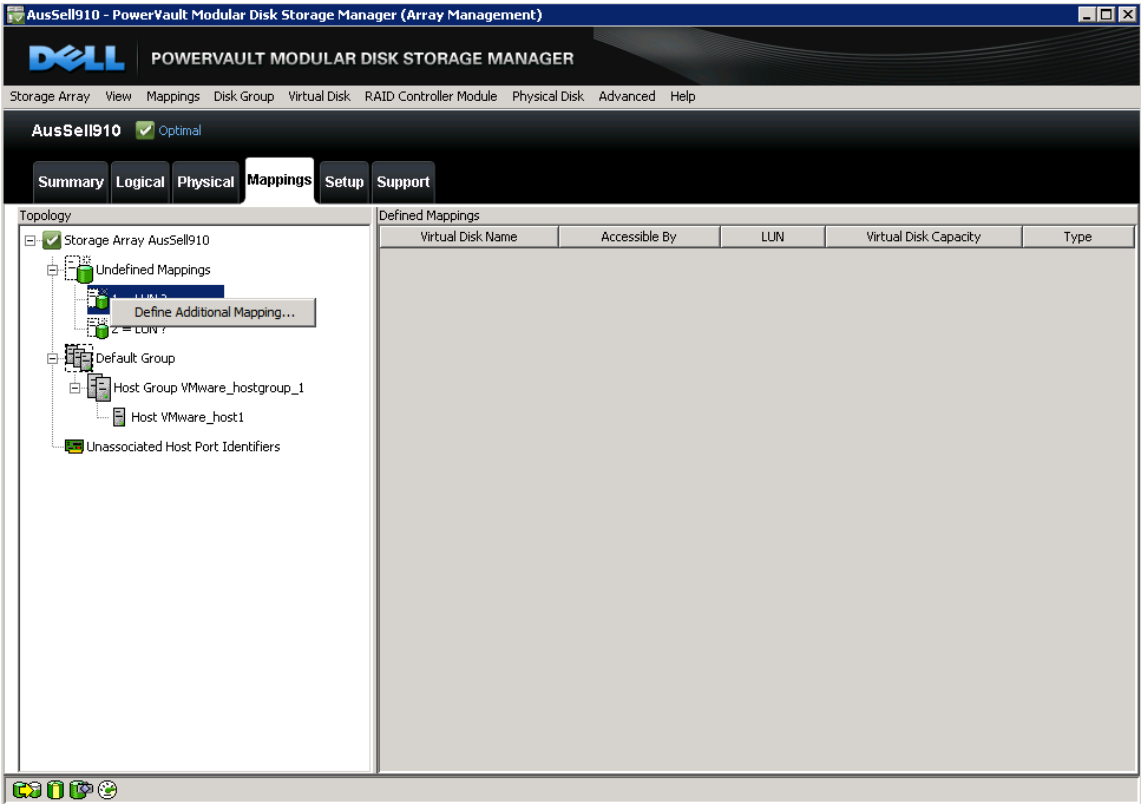


**Figure 3: Selecting a Virtual Disk for mapping**

Remember that the virtual disk is assigned to the host group and not the host. For this example we selected the host group that was defined in the previous steps.

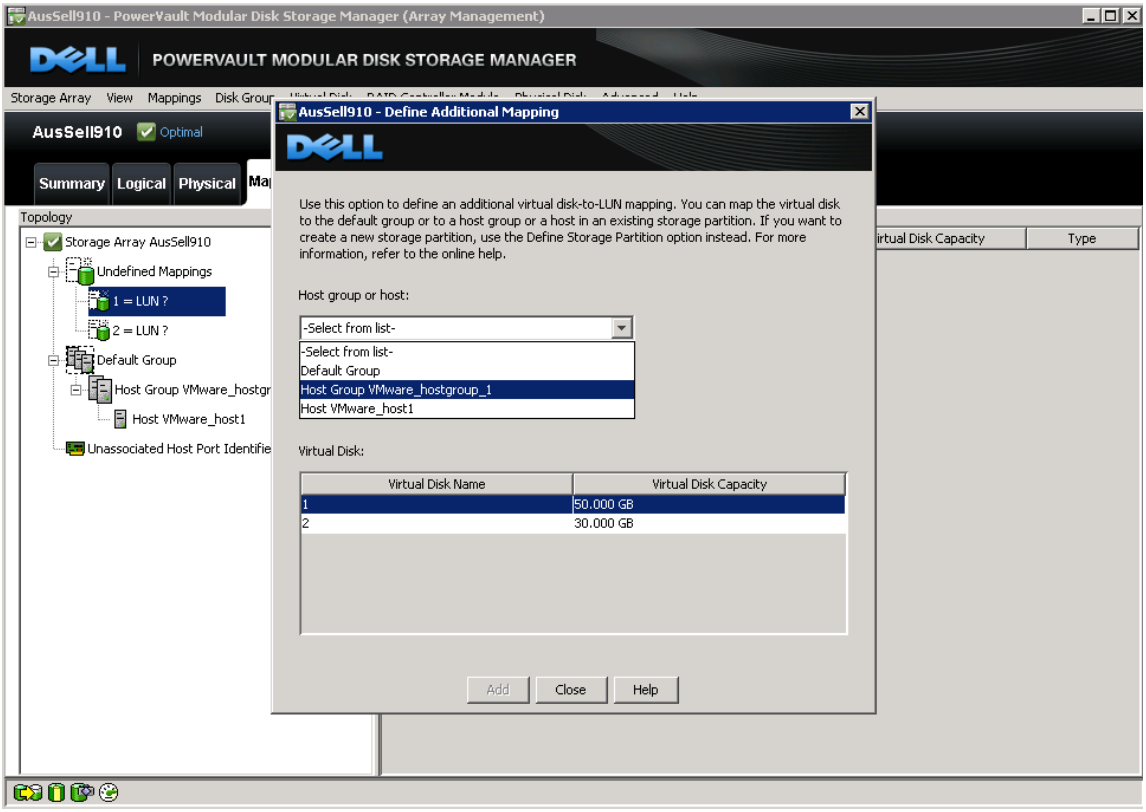


**Figure 4: Selecting the Host Group**

Assign the other virtual disk to the same host group.

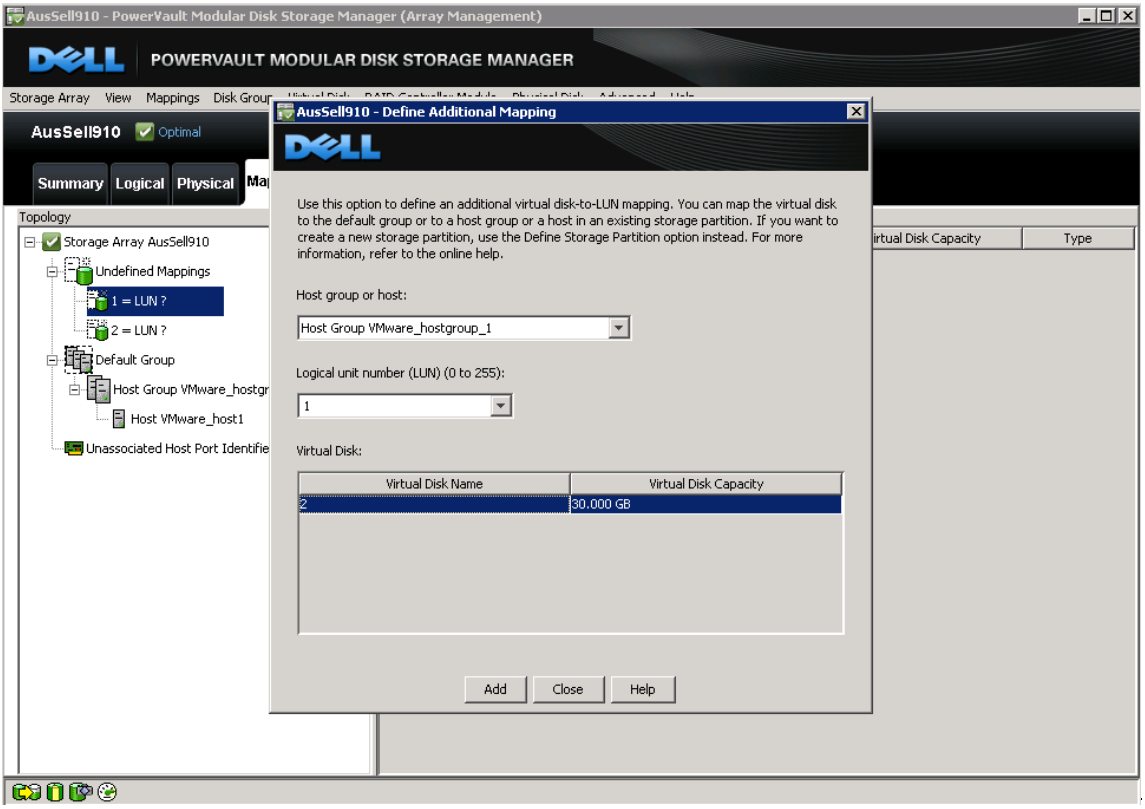


**Figure 5: Assigning additional virtual disks**

After the virtual disks are assigned notice that the host group and its associated hosts are no longer under the default group in the topology. This completes the configuration.

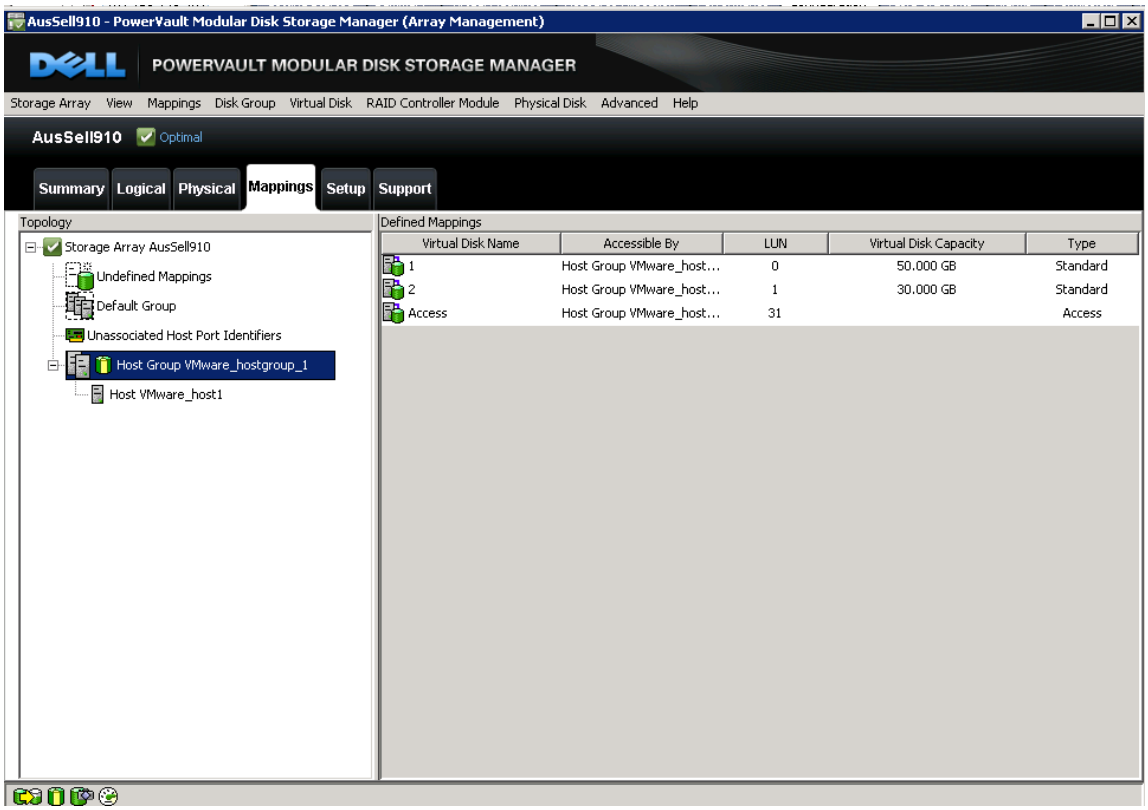


**Figure 6: Completed Topology with Assigned Virtual Disks**

## STEP10: PERFORM RESCAN SO VMWARE SEES THE VIRTUAL DISK

The next step is to connect to the volume on the SAN and verify the connection status. Since the iSCSI access and configuration was configured in the last step, the only thing to do now is to rescan the iSCSI Adapter and make sure the volume appears correctly.

In the vSphere4 GUI click on ***Configuration -> Storage Adapters*** and select the iSCSI Software Adapter.

Click ***Rescan All*** and choose to Scan for New Storage Devices and select ***Ok.***

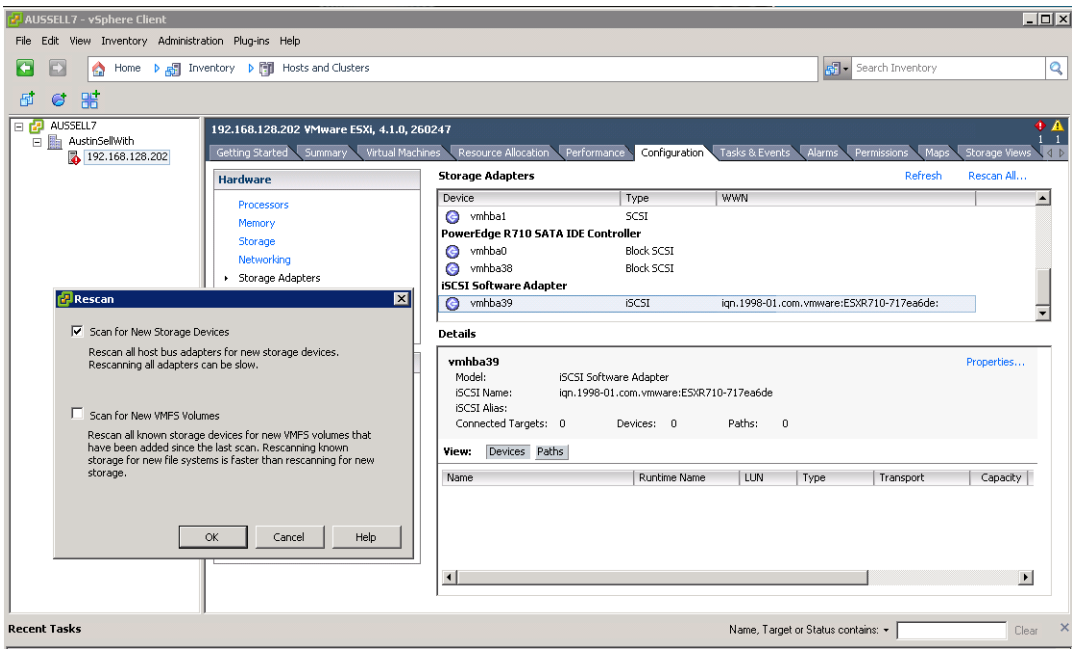


**Figure 7: Perform Rescan**

If everything has been configured properly under Devices there will be new PowerVault iSCSI Disks with the correct LUNs, similar to what's shown below.

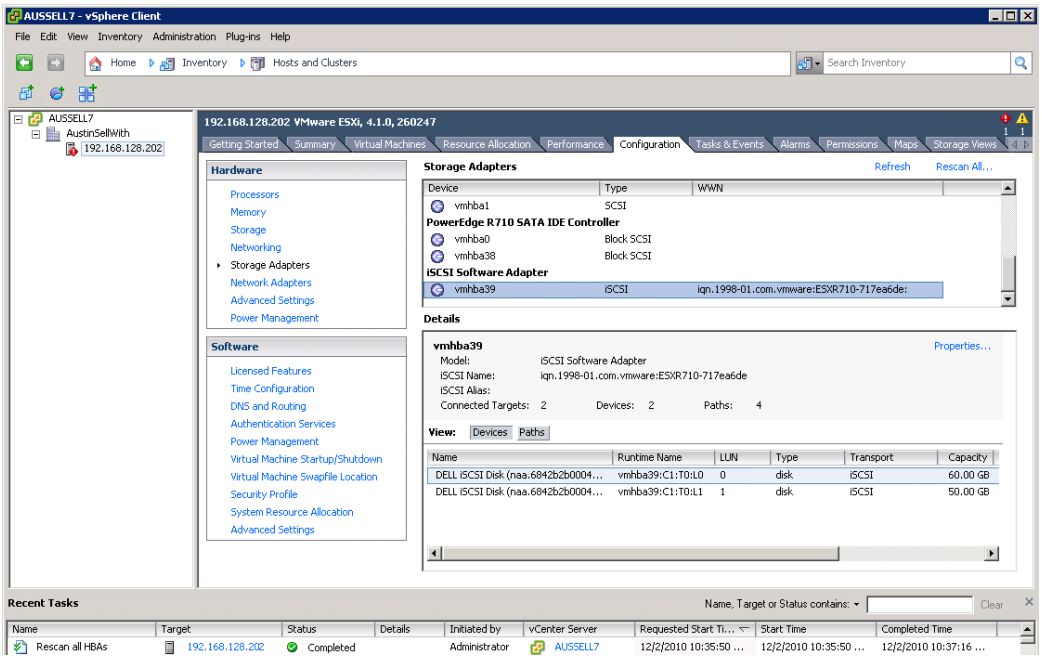


**Figure 8: Configured Devices**

If the device information is correct select the Path Tab and verify that there are active and standby paths for each device (LUN)

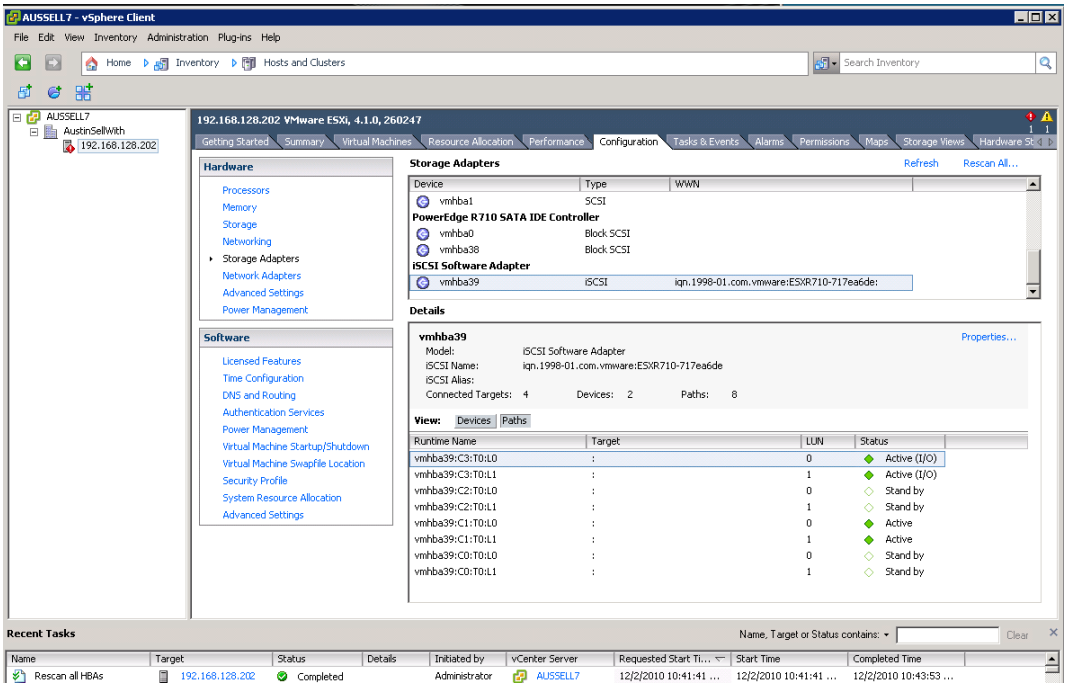


**Figure 9: Configured Paths**

Repeat this step for each additional new Dell iSCSI disk that is created.

## STEP11: CREATE VMFS DATASTORES

Now that the iSCSI Software adapter is set up and configured and the Dell iSCSI disks are visible to the adapter, they can also be formatted as VMFS Datastores and utilized as normal.

Creating a Data Store from the MD3200i LUNS is the same as creating a Data Store with any local disk. Begin by selecting Storage under Hardware and then select *Add Storage*

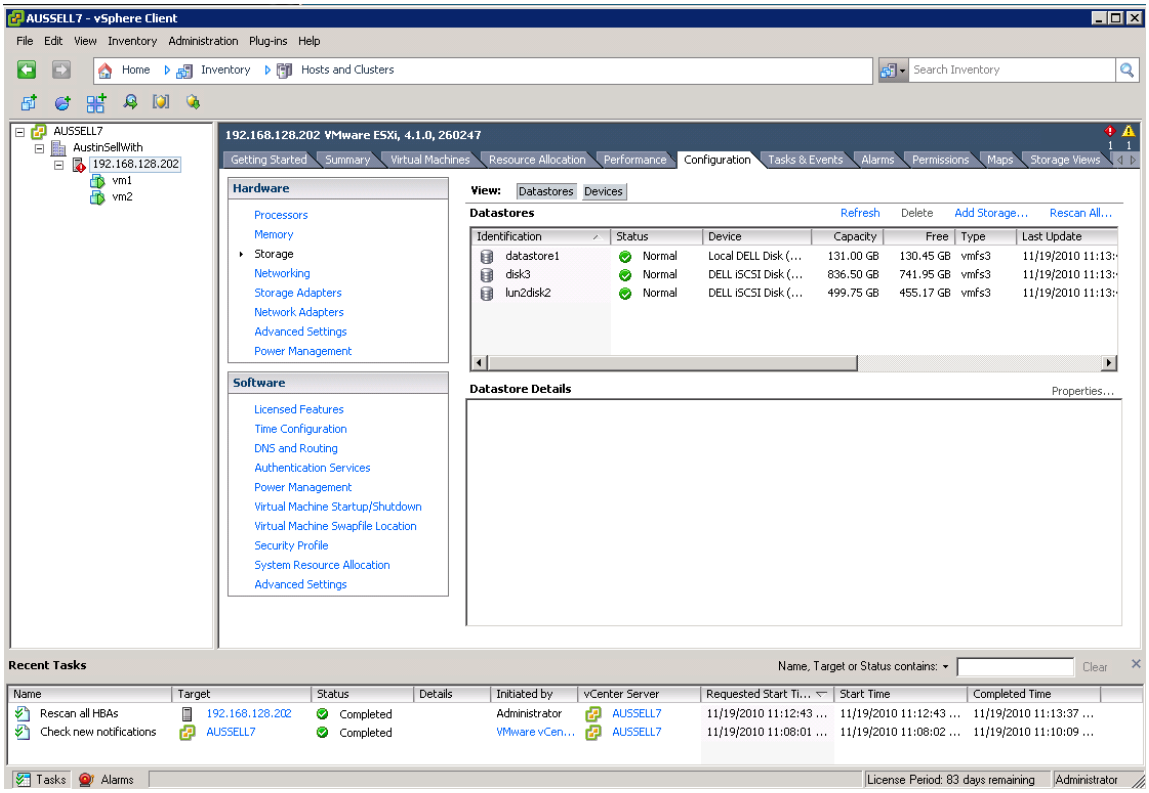


**Figure 10: Creating VMFS Datastores**

Because an iSCSI disk is considered a local SCSI disk the storage type is Disk/LUN

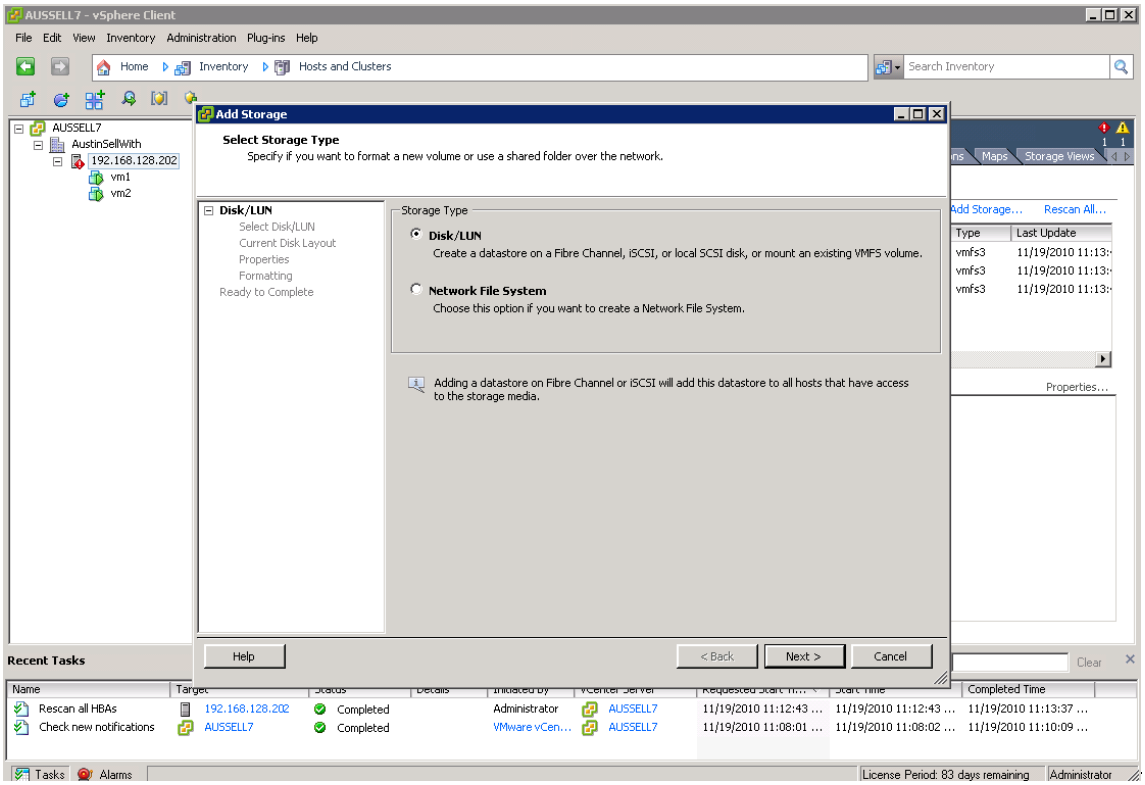


**Figure 11: Creating VMFS Datastores (continued)**

Select one of the LUNS from the MD3200i to create a Datastore from.

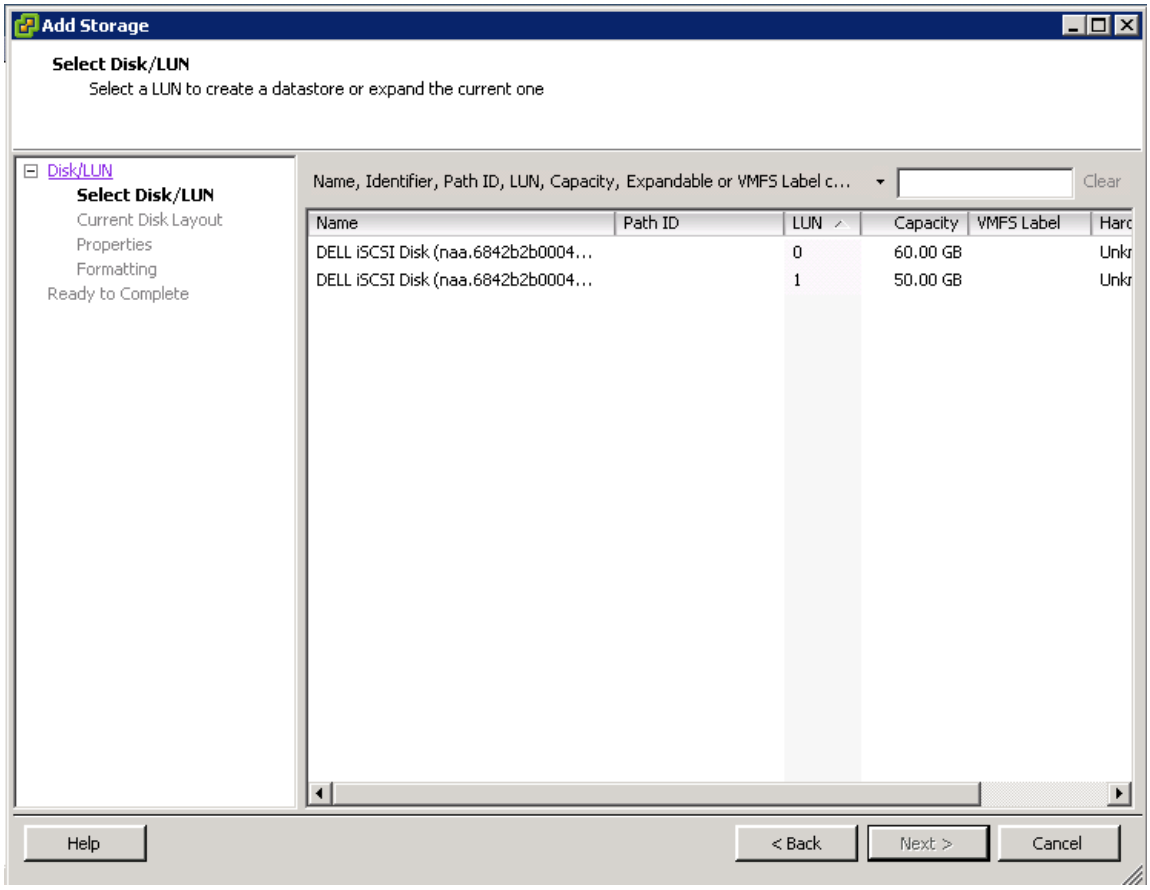


**Figure 35: Selecting a LUN to create a Datastore**

This screen displays the information about the disk layout. Select next to create a VMFS partition.

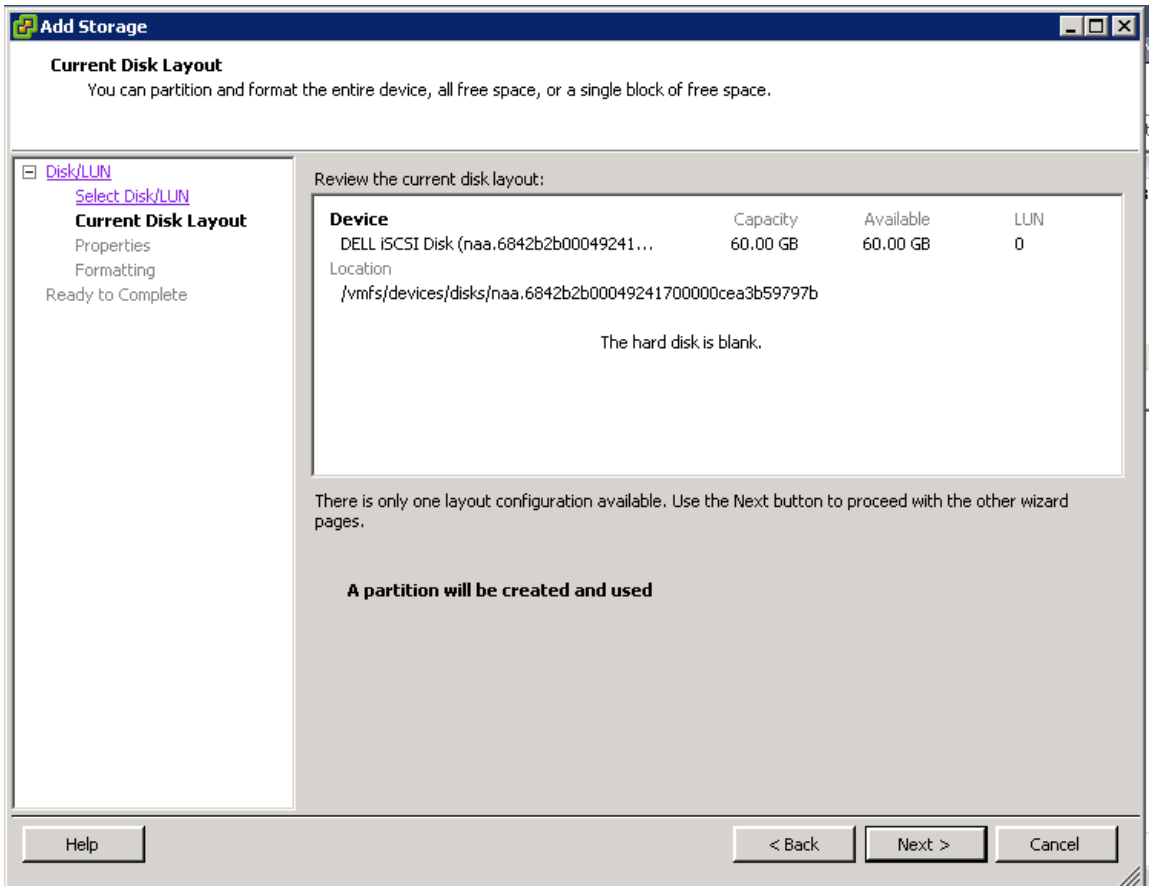


**Figure 36: Disk Layout**

Enter a Datastore name and select next.

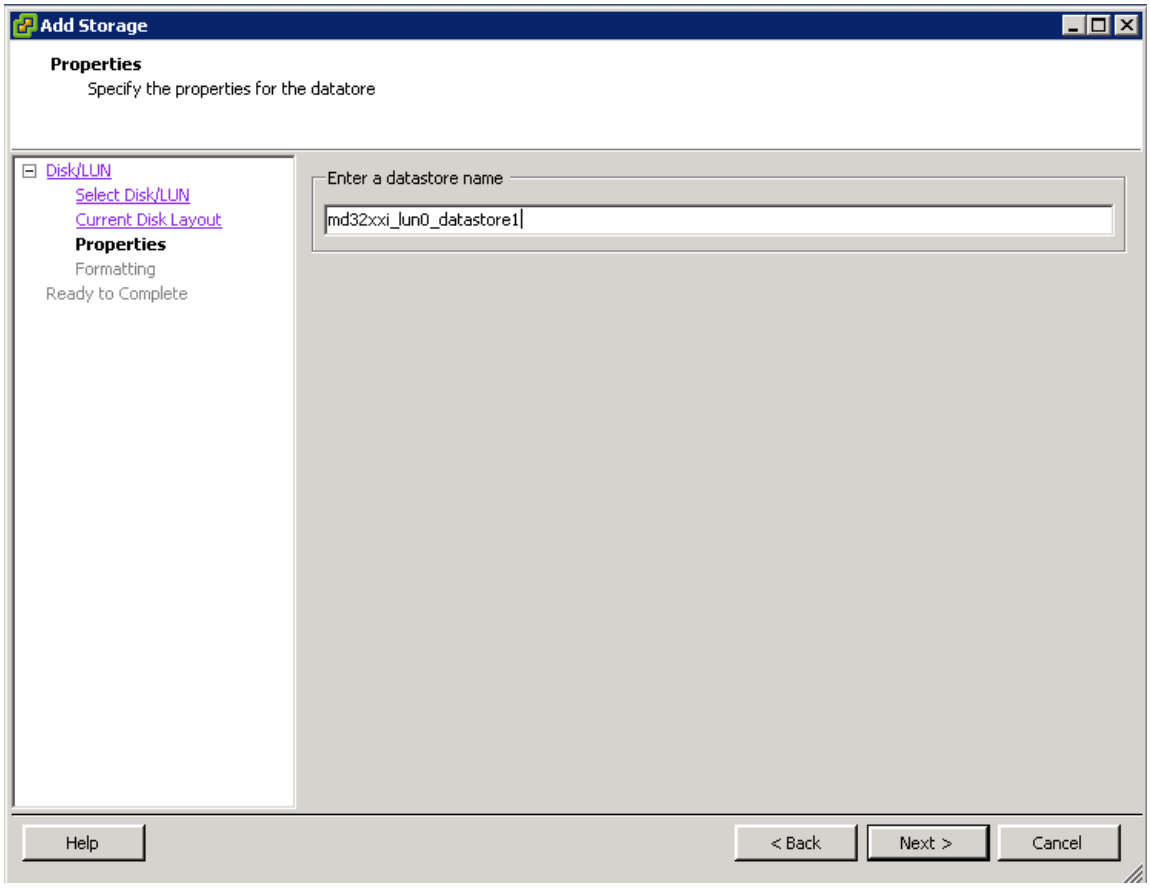


**Figure 12: Datastore naming**

Adjust the Maximum file size as needed. For this example we used the maximum capacity. Select ***Next*** when finished.

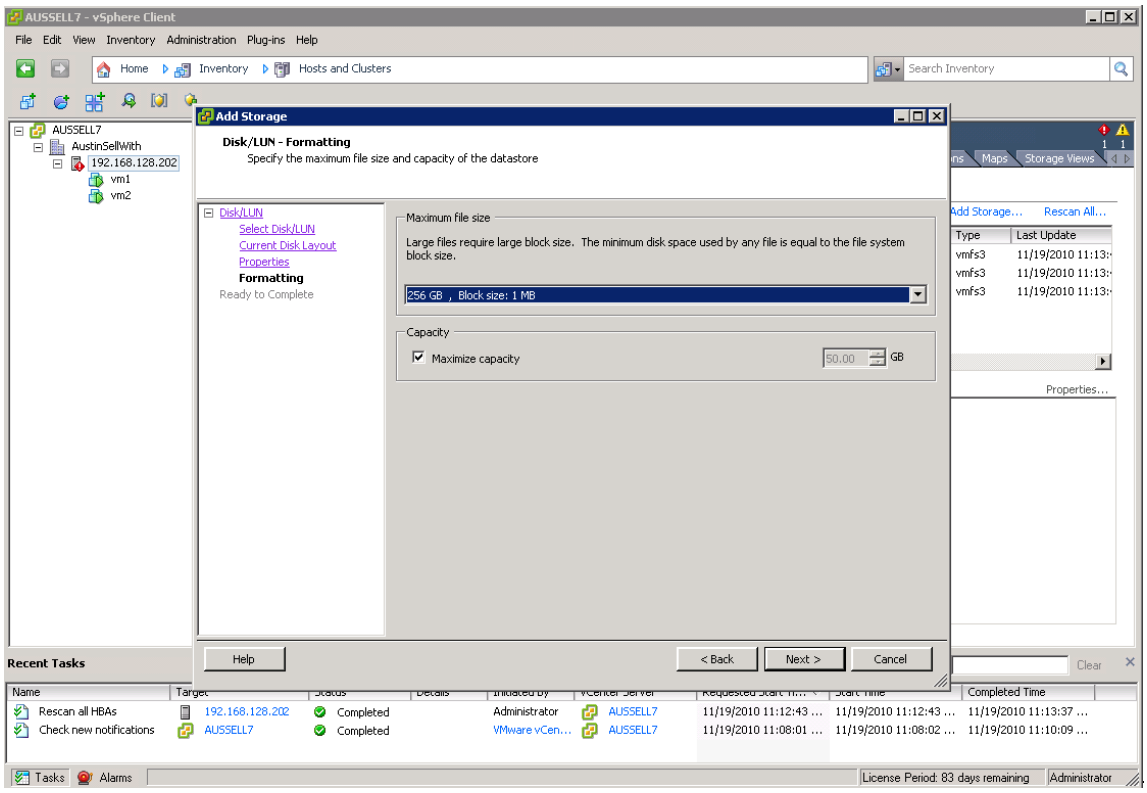


**Figure 13 Maximum File Size**

Review the disk layout and click ***Finish*** to add storage.

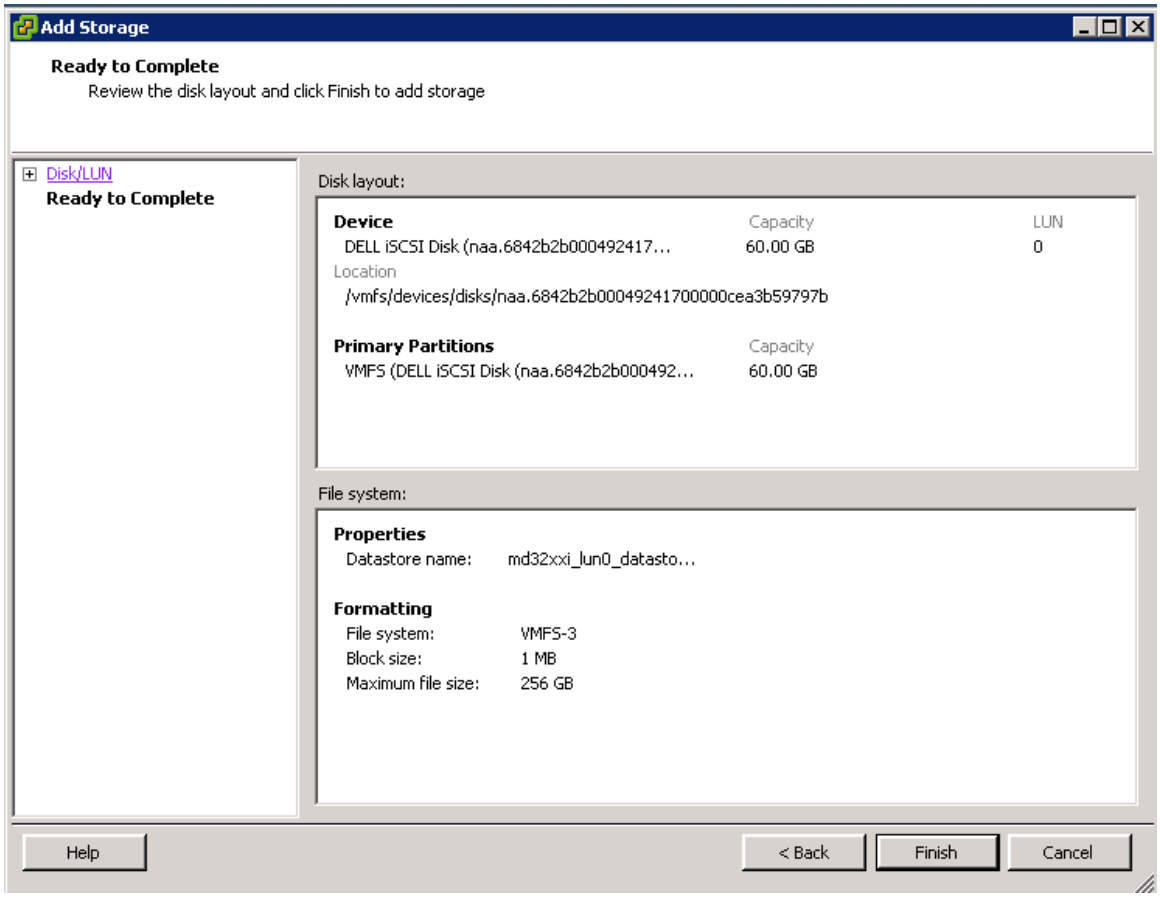


**Figure 14 Review disk layout**

The new storage is completed and ready to use with VMs.

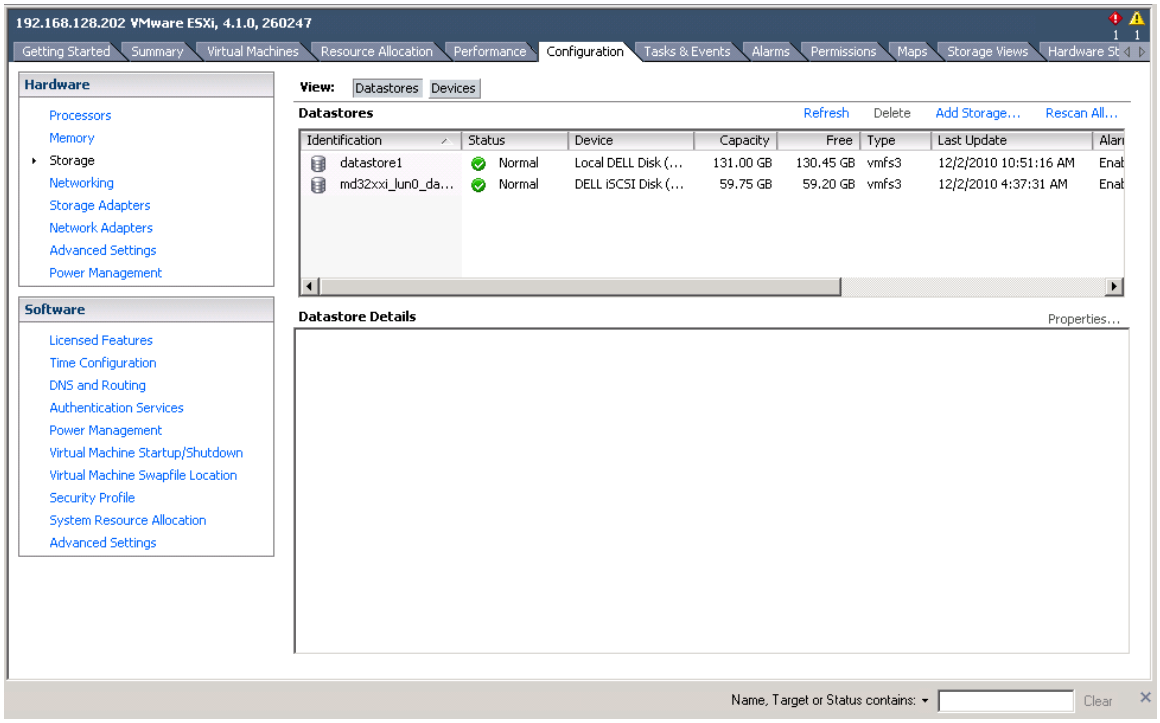


**Figure 15: Configured Paths**

## STEP12: CONFIGURING VMWARE PATH SELECTION POLICY

Features enabled by configuring the iSCSI Software Initiator the way we have is that now we can take advantage of several multipathing policies, such as Most Recently Used(MRU) and Round Robin(RR). This allows for greater and better bandwidth utilization than in previous versions of ESX.

To verify the Path Selection Policy on a datastore, select the volume from the vCenter GUI. ***Configure -> Storage***. Highlight the datastore and select properties.

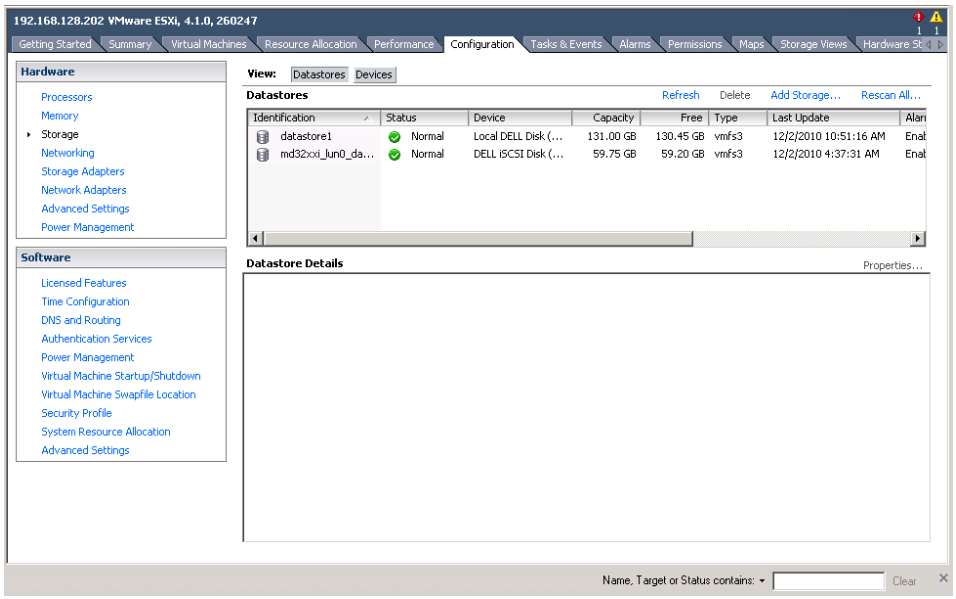


From the properties window select ***Manage Paths***.

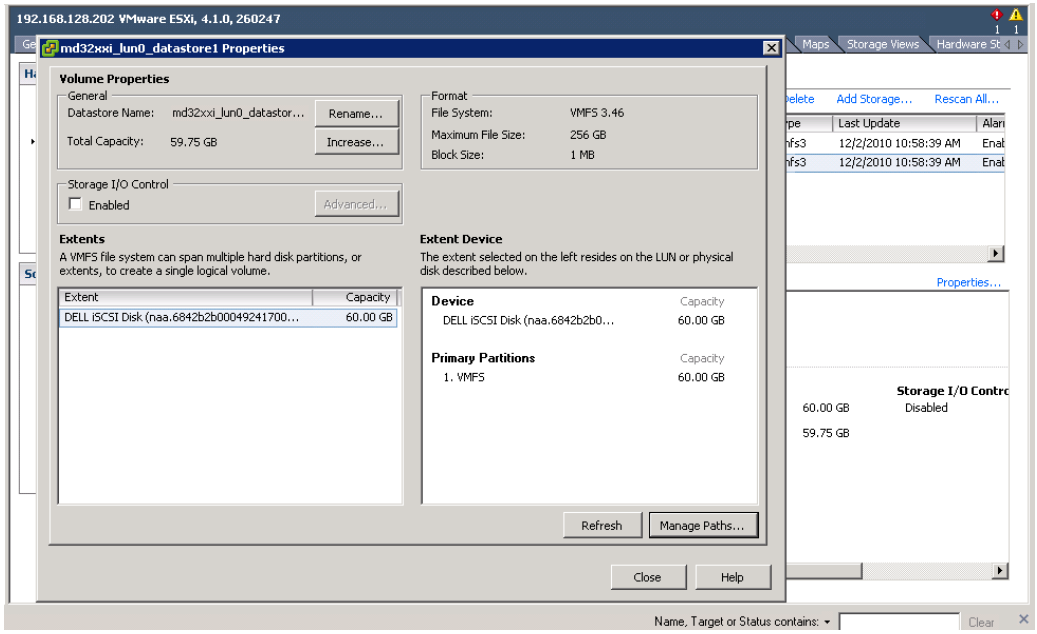

The path selection can be either MRU or RR and the Storage Array Type should be VMW_SATP_LSI.

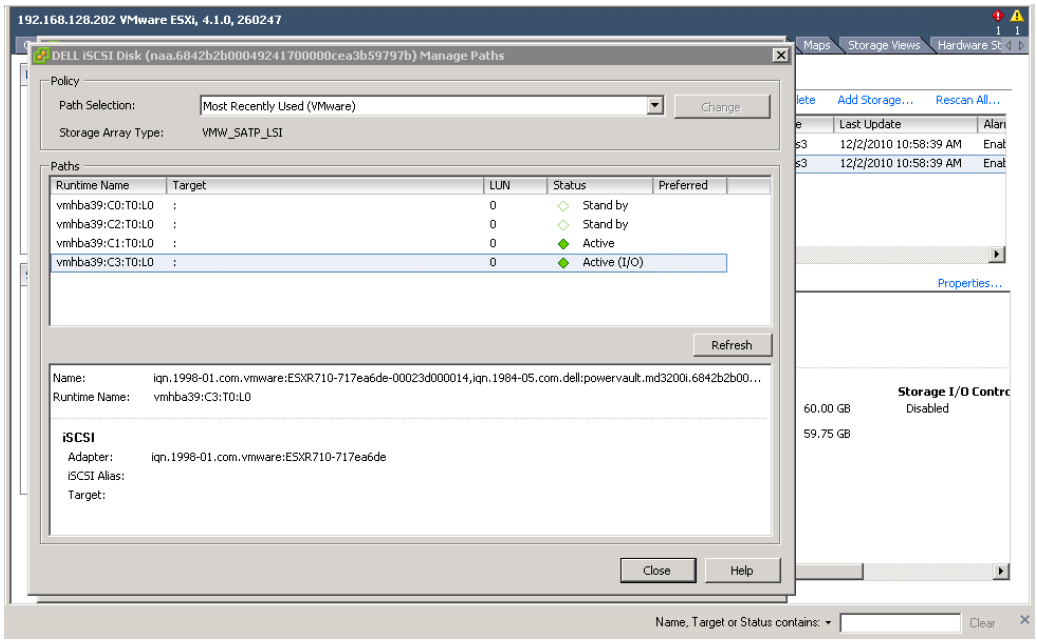


# CONTACT INFORMATION

HTTP://SUPPORT.DELL.COM/SUPPORT/TOPICS/GLOBAL.ASPX/SUPPORT/PRODUCT_SUPPORT/PRODUCT_SUPPORT_CENTRAL?C=US&CS=555&L=EN&S=BIZ**References**

VMware vSphere 4.1 Documentation:

http://www.vmware.com/support/pubs/vs_pages/vsp_pubs_esxi41_e_vc41.html

Dell/VMware alliance home page:

http://www.dell.com/vmware

# APPENDIX A

# Steps for using the iSCSI Offload functions of a NIC

Lower CPU utilization and potentially lower power consumption justify offload adapters(a.k.a. h/w initiators).

With hardware-based iSCSI storage, you use a specialized third-party adapter capable of accessing iSCSI storage over TCP/IP(a.k.a. iSCSI Offload, ISOE adapters). This iSCSI initiator handles all iSCSI and network processing and management for your ESX/ESXi system.

The list of hardware initiators supported on the MD3200i can be found here.

You must install and configure the hardware iSCSI adapter for your host to be able to access the iSCSI storage device. For installation information, see either Broadcom or Emulex documentation.

With respect to Broadcom NICs, this site from VMWare identifies 'Functionality Caveats', which include:

- Broadcom Hardware iSCSI does not support Jumbo Frames or IPv6.
- Dependent hardware iSCSI does not support iSCSI access to the same LUN when a host uses dependent and independent hardware iSCSI adapters simultaneously. This site gives further information regarding dependent & independent adapters. Dell support in this regard extends only to dependent adapters.

# Considerations involving subnet configuration

It is recommended that you have different ports on different subnets due to throughput and pathing considerations. The decision to use multiple subnets may depend upon the specific architecture and networking topology and may not be essential in all cases.

Multiple subnets should be allocated by the number of array ports per controller. With the MD3200i you only get an active path to multiple ports on the same controller if they are on different subnets. Likewise, with VLANs, for multiple target ports on the same controller you only get an active path if the target ports are on different VLANs. Since the MD3220i has 4 ports per controller you get your best throughput with 4 subnets.

On an MD3200i/MD3600i, the LUNs are assigned out of disk groups. Only one controller at a time has access to the LUNs. The alternate controller has access to those LUNs but only as a

failover alternate path. This LUN/controller relationship is reported by the MD3200i/MD3600i to VMware which then dictates the paths (and the "optimal" condition of the array).

On the MD3200i/MD3600i, setup each port on a different subnet with one disk group per controller, each disk group with 4 LUNS (one LUN per port).

Regardless of the pathing method in VMware, without separate subnets a bottle-neck at the controller *<will arise>* since pathing is at the disk group level rather than the LUN level for the MD3200i/MD3600i.

Configuring RR to the same LUN across both controllers is not recommended.

Here are the ESXi CLI commands to switch the pathing policy. You run this for each device. Note, this can also be accomplished via the GUI.

esxcli nmp device setpolicy --device < device name > --psp VMW_PSP_MRU
esxcli nmp device setpolicy --device < device name > --psp VMW_PSP_RR
esxcli nmp device list

The number of servers, and physical NICs as well as the I/O access characteristics will affect how you configure the MD3200i*/MD3600i*. Performance depends on many variables. For example, if you have lots of random data going to one LUN and lots of sequential data on another LUN, you might want them to be on separate disk groups as all that random movement could inhibit performance of the sequential data. Now, if all the data is sequential, you might get better performance keeping both the LUNs on a large disk group and take advantage of spindles. So, a lot depends on the nature of your data.

A performance tuning guide that talks about various variables that you can play with and how each might impact performance is available here.

# Using the MD Configuration Utility(MDCU)

This utility provides a wizard for configuring ports on an MD300xi. A user would employ this on a guest OS under an established VM.

# Considerations involving Volume Not On Preferred Path

Virtual Disk Not On Preferred Path(VNOP) condition is the result of a Virtual Disk (a.k.a. LUN) being assigned to the alternate controller from its original assignment when first created.

This situation can arise as a combination of default settings for Virtual Disk creation in the MD3200i and the creation of active iSCSI sessions between the host and the MD3200i for iSCSI

access. Insufficient active paths have been defined to allow multipath access preferred by the MD3200i storage array for performance and redundancy features.

The difference between current path and MD3200i preferred path can be a configuration issue, and in most cases is not a functional path or hardware failure. There are corrective actions possible to resolve the reported Virtual Disk Not On Preferred Path fault condition without replacement of any hardware.