



## **SERVER CONNECTIVITY**

### **Technical Brief: How to Configure NPIV on VMware ESX Server 3.5**

Provides step-by-step instructions on how to configure NPIV on VMware ESX Server 3.5 in a Brocade fabric. Leveraging NPIV gives the administrator the ability to extend all of Brocade's advanced features and apply storage best practices in a virtualized data center.

**BROCADE**

---

## CONTENTS

<b>Introduction to NPIV</b> .....	<b>3</b>
Benefits of NPIV .....	3
<b>Server Virtualization Advantages</b> .....	<b>4</b>
Addressing and Access Control.....	4
NPIV in a Virtualized Environment .....	4
<b>Requirements to Implement NPIV</b> .....	<b>4</b>
Server Administrator Tasks .....	5
Storage Administrator Tasks .....	9
Verifying that the VM Is Functioning Correctly.....	21
<b>Troubleshooting</b> .....	<b>22</b>

## INTRODUCTION TO NPIV

N\_Port ID Virtualization (NPIV) is an extension to the Fibre Channel industry standard, which is available across the Brocade® Host Bus Adapter (HBA) product line and Brocade Storage Area Network (SAN) platforms. NPIV delivers significant advantages for running multiple Virtual Machines (VMs) and managing the workload across multiple physical servers.

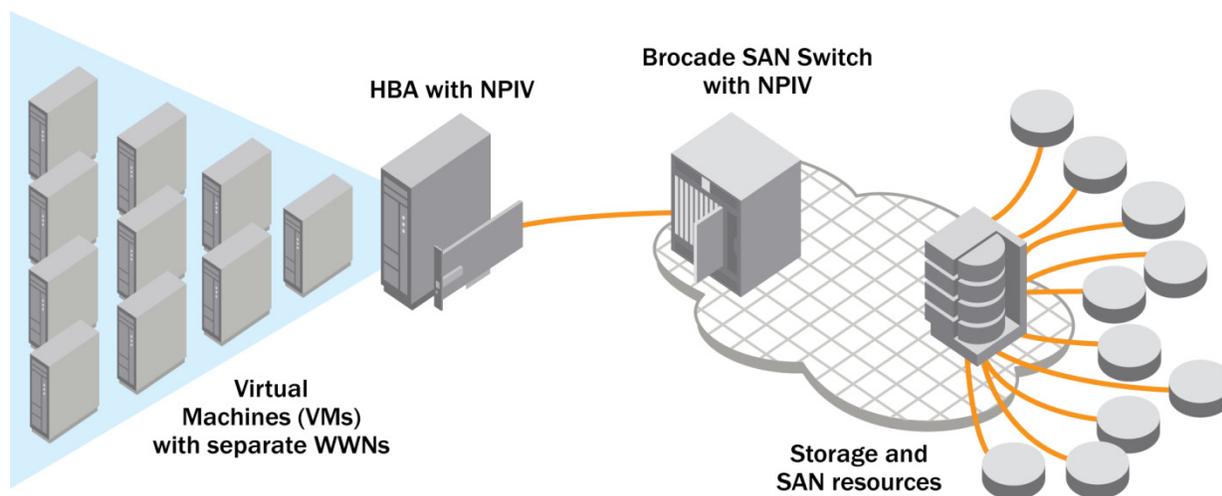
**NOTE:** The term “switches” is used in this document to reference Brocade backbone, director, switch, and embedded switch platforms unless otherwise noted.

## Benefits of NPIV

In a server virtualization environment, NPIV allows each VM to have a unique Fibre Channel (FC) World Wide Name (WWN), enabling multiple virtual machines to share a single physical HBA and switch port, a virtual HBA port, as shown in Figure 1. By providing a unique virtual HBA port, storage administrators can implement SAN best practices such as zoning for individual VMs. Administrators can also take advantage of SAN management tools, simplifying migration of VMs and their storage resources.

The benefits of deploying NPIV in your storage environment are real and available today:

- Maintaining fewer physical components reduces the number of points of failure, resulting in improved availability and network uptime.
- Less hardware, portable connections, and VM-level zoning all contribute to simplified SAN and server management.
- NPIV allows the SAN best practices that are available with physical servers to be used in virtual server environments.



**Figure 1.** NPIV in the SAN

---

## SERVER VIRTUALIZATION ADVANTAGES

A number of enhancements are being introduced into server virtualization products on the market, such as VMware® ESX Server, to augment existing support for Fibre Channel SANs, including NPIV and load balancing across FC ports.

### Addressing and Access Control

Each FC port in a fabric has a World Wide Name assigned to it by the equipment manufacturer, which uniquely identifies each node. WWNs play a critical role in determining the visibility and accessibility of storage LUNs (partitions in storage arrays) by servers connected to the fabric. Zoning is the mechanism by which FC ports are grouped together to restrict interference, add security, and simplify management. Zoning utilizes WWNs to allow access to storage. A server can see and access only storage LUNs that share a common zone with that server.

### NPIV in a Virtualized Environment

The hypervisor leverages NPIV to assign individual WWNs to each Virtual Machine, so that each Virtual Machine (VM) can be recognized as a specific end point in the fabric. The benefits of this approach are as follows:

- **Granular security.** Access to specific storage LUNs can be restricted to specific VMs using the VM WWN for zoning, in the same way that they can be restricted to specific physical servers.
- **Easier monitoring and troubleshooting.** The same monitoring and troubleshooting tools used with physical servers can now be used with VMs, since the WWN and the fabric address that these tools rely on to track frames are now uniquely associated to a VM.
- **Flexible provisioning and upgrade.** Since zoning and other services are no longer tied to the physical WWN “hard-wired” to the HBA, it is easier to replace an HBA. You do not have to reconfigure the SAN storage, because the new server can be pre-provisioned independently of the physical HBA WWN.
- **Workload mobility.** The virtual WWN associated with each VM follows the VM when it is migrated across physical servers. No SAN reconfiguration is necessary when the workload is relocated to a new server.
- **Applications identified in the SAN.** Since virtualized applications tend to be run on a dedicated VM, the WWN of the VM now identifies the application to the SAN.

## REQUIREMENTS TO IMPLEMENT NPIV

There are a few requirements in both the software and hardware to enable NPIV:

- **Switches.** NPIV needs to be supported on the switch connected to the HBA. All Brocade FC switches currently support NPIV—specifically starting in FOS 5.3.x or later, M-EOSc 8.1, and M-EOSn 9.6.0. (See the “[Troubleshooting](#)” section to find out how to see if NPIV is enabled on the switch.)
- **HBAs.** HBAs must support NPIV as well. The following vendors and models of HBAs are supported:
  - **Brocade.** Any 4 Gbit/sec or 8 Gbit/sec HBA
  - **Emulex.** 4Gbit/sec HBA running firmware level 2.70a5 or later. All Emulex 8 Gbit/sec HBAs running firmware 1.00a9 or later
  - **QLogic.** Any 4 Gbit/sec or 8 Gbit/sec HBA
- **Storage.** NPIV is completely transparent to storage arrays, so no specific support is required.

- NPIV can be used *only with Raw Device Mappings (RDM) disks*. VMFS disks do not support NPIV. For more information on RDMS, refer to the current [Fibre Channel SAN Configuration Guide for ESX Server 3.5](#).
- To implement NPIV, the physical HBAs port WWN on an ESX Server host must have access to all LUNs that are to be accessed by VMs.
- If you deploy a VM from a template or a clone, the new VM does not retain the WWN.

## Server Administrator Tasks

**NOTE:** Do not power on your VM for this procedure; your virtual machine needs to be powered off.

1. Assign or modify WWN or your Virtual machine by following the steps below. These steps are also found in the [ESX 3.5 SAN Configuration Guide](#).

### To assign a WWN to an existing VM using Virtual Center:

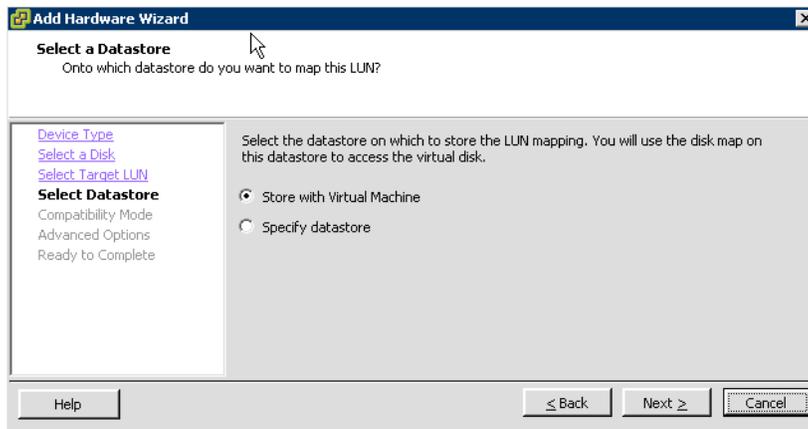
- a. From Virtual Center, select the VM to which you want to assign a WW, right click and select **Edit the Virtual Machine Settings**.
- b. Click the **Options** tab and click **Fibre Channel NPIV**.
- c. In the dialog box, select **Generate**. Here are the available options:
  - Leave Unchanged: Existing WWNs are retained.
  - Generate New WWN: New set of WWNs are generated.
  - Remove WWN assignment. Removes the WWNs from the VM.
- d. Click **Close** and exit configuration.

### To assign a WWN to a new VM using Virtual Center:

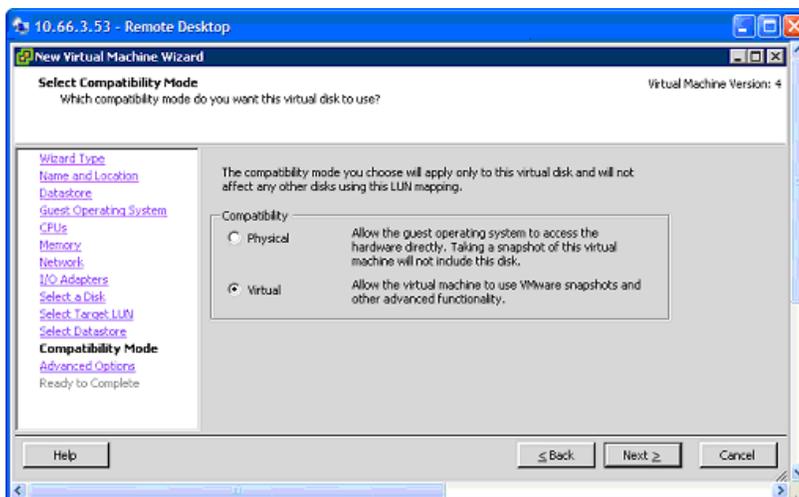
- a. Right-click on an ESX server or cluster and click **New Virtual Machine**.
- b. Click the **Custom** radio button.
- c. Give the VM a name and click **Next**.
- d. Enter the appropriate information and stop at the Select Disk Type screen.
- e. Select the **Raw Disk Mapping** and click **Next**.

- f. Select the datastore.

**NOTE:** To use VMotion for VMs that have NPIV enabled, make sure that the RDM file is located on the same datastore where the virtual machine configuration file resides.



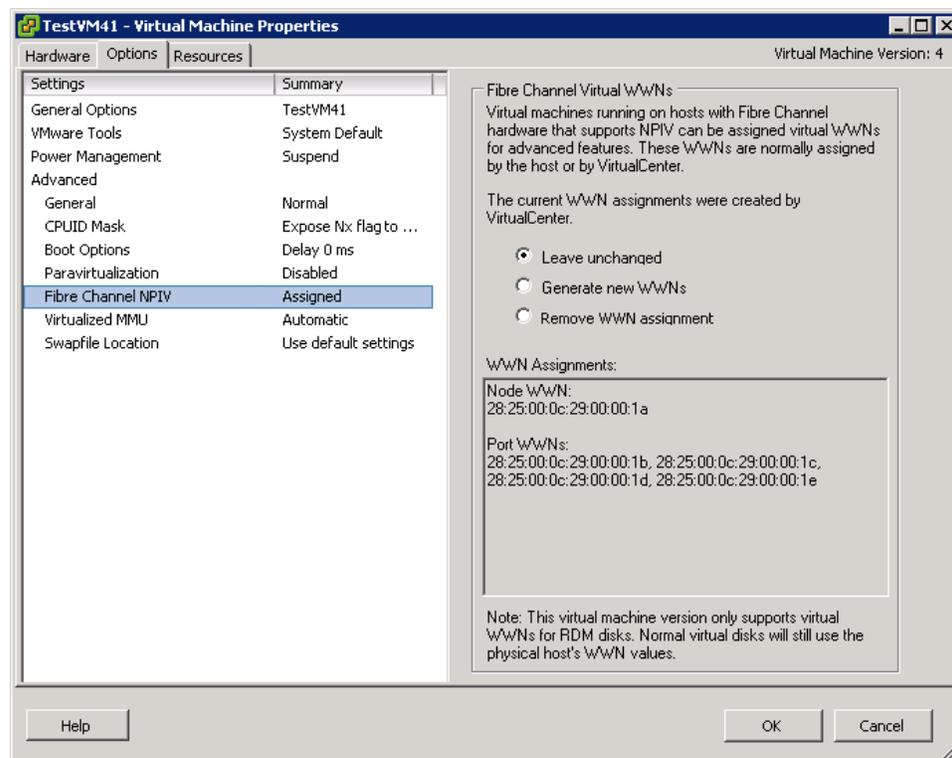
- g. Select the **Compatibility Mode**. For more information on the different compatibility modes, refer to the “[Fibre Channel SAN Configuration Guide](#).”



- h. Leave the defaults for the Advanced Options and click **Next**.
- i. In the last screen, click the **Edit the virtual machine settings before submitting** check box.
- j. Click the **Options** tab and click **Fibre Channel NPIV**.
- k. In the dialog box, select **Generate New WWN**.
- l. Click **Close** and exit configuration.

**NOTE:** Do not power on the VM guest after assigning the RDMs a virtual WWN. If you do so, it will use the physical HBA WWN.

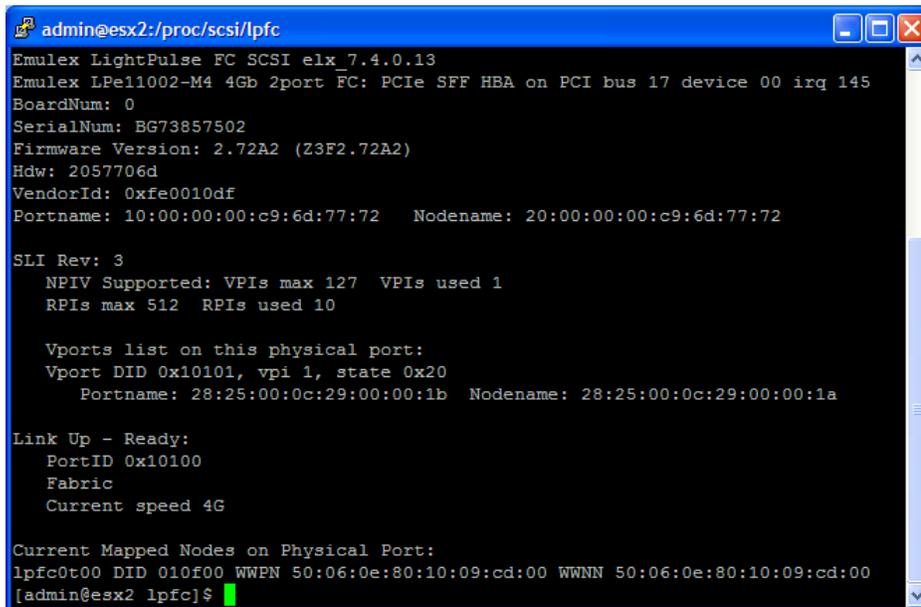
2. Once you have assigned a WWN to the VM, go back into the VM settings by right-clicking the VM and choosing **Edit Settings** from the menu.
3. Click the **Options** tab and click **Fibre Channel NPIV** in the left pane.



4. Record the Node WWN and the Port WWNs. Give this information to the storage administrator so that (s)he can zone and LUN-mask the back-end storage. (Please follow the Storage Administrator steps below to zone and LUN-mask)
5. After the storage administrator completes zoning and configuration of the storage, you can safely power on the VM.
6. To verify that I/O traffic is actually going through the virtual port (depending on the host HBA). SSH into the physical ESX Server (in all cases, replace the number "1" at the end with the number of your HBA:
  - a. Brocade HBA –cat /proc/bfa/1

```
[root@esx1 bfa]# cat 2
Chip Revision: Rev-C
Manufacturer: Brocade
Model Description: Brocade-825
Instance Num: 1
Serial Num: ALX0417D0A9
Firmware Version: FCHBA1.0.0
Hardware Version:
Bios Version:
Optrom Version:
Port Count: 2
WWNN: 20:00:00:05:1e:0d:61:4a
WWPN: 10:00:00:05:1e:0d:61:4a
Instance num: 1
Target ID: 0 WWPN: 50:06:01:60:3a:60:0f:7b
VPort list:
WWPN: 28:03:00:0c:29:00:00:07
[root@esx1 bfa]#
```

- b. QLogic HBA—`cat /proc/scsi/qla2300/1`
- c. Emulex HBA—`cat /proc/scsi/lpfc/1`



```
admin@esx2:/proc/scsi/lpfc
Emulex LightPulse FC SCSI elx_7.4.0.13
Emulex LPe11002-M4 4Gb 2port FC: PCIe SFF HBA on PCI bus 17 device 00 irq 145
BoardNum: 0
SerialNum: BG73857502
Firmware Version: 2.72A2 (Z3F2.72A2)
Hdw: 2057706d
VendorId: 0xfe0010df
Portname: 10:00:00:00:c9:6d:77:72  Nodename: 20:00:00:00:c9:6d:77:72

SLI Rev: 3
  NPIV Supported: VPIs max 127  VPIs used 1
  RPIs max 512  RPIs used 10

Vports list on this physical port:
Vport DID 0x10101, vpi 1, state 0x20
  Portname: 28:25:00:0c:29:00:00:1b  Nodename: 28:25:00:0c:29:00:00:1a

Link Up - Ready:
  PortID 0x10100
  Fabric
  Current speed 4G

Current Mapped Nodes on Physical Port:
lpfc0t00 DID 010f00 WWPN 50:06:0e:80:10:09:cd:00 WWNN 50:06:0e:80:10:09:cd:00
[admin@esx2 lpfc]$
```

Everything is working if you see the virtual WWNs (also know as vPorts) listed.

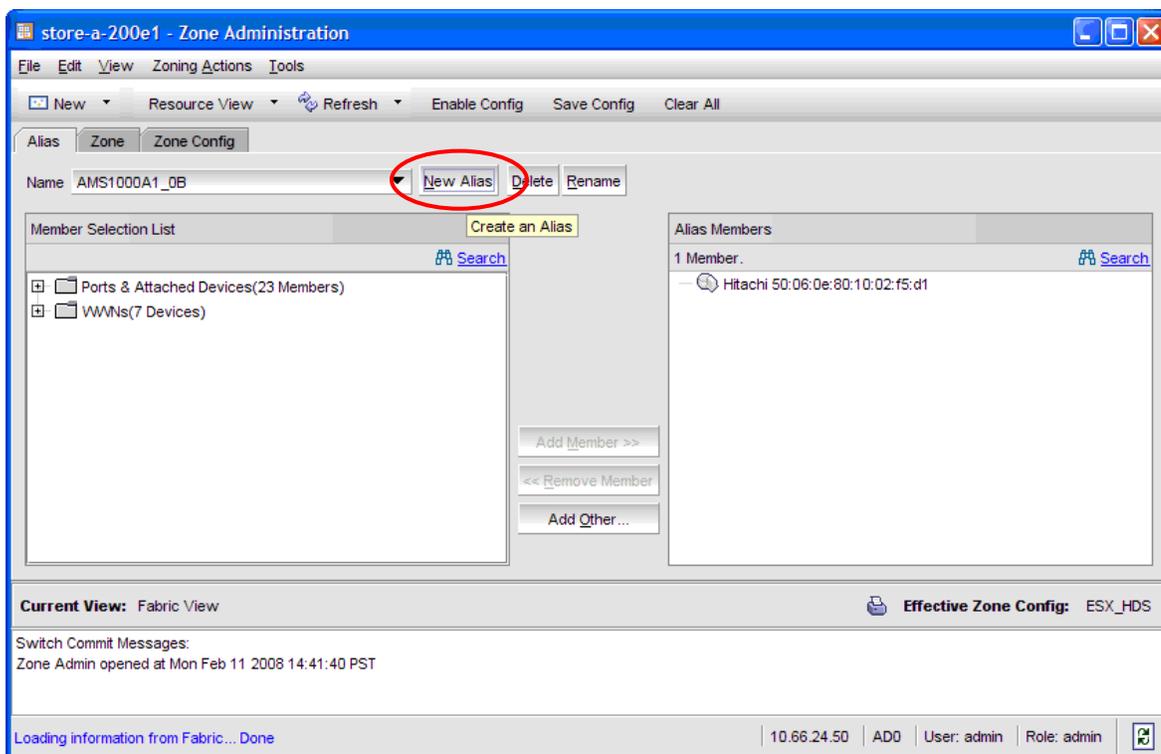
## Storage Administrator Tasks

Prior to zoning the VM, be sure that you have zoned all the physical HBA WWNs to the correct storage array port WWN/s. Storage best practices are to zone a single initiator to single target to maintain security and reduce interference. Once you receive the virtual WWNs from the Server administrator, you then need to create unique zones for each VM to the storage array port WWN/s, However, if you have a very small environment, then you may find it easier to place all the physical HBA port WWNs into one zone and then individually zone each VM to an array port.

1. Log in to the Brocade switch using Brocade Web Tools, Brocade Data Center Fabric Manager (DCFM™), or the Fabric OS® (FOS) Command-Line Interface (CLI) to configure Zoning. This paper will explain how to Zone using Brocade Web Tools and the CLI.

**To zone using Brocade Web Tools (assuming that all the physical HBAs ports and array ports are zoned):**

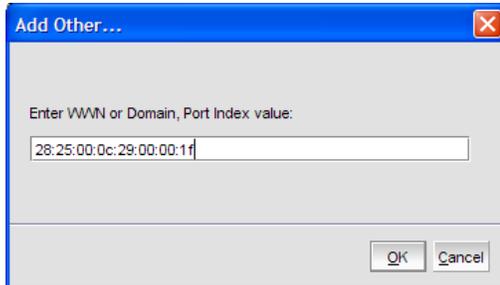
- a. Log in to Web Tools by opening up Internet Explorer and typing in the FQDN or IP address (e.g., <https://5100edge1.brocade.com>)
- b. Log in to the switch, and click **Zone Admin**.
- c. Click **New Alias**, and type a name for the VM guest Node Name (for example, vmguest1).



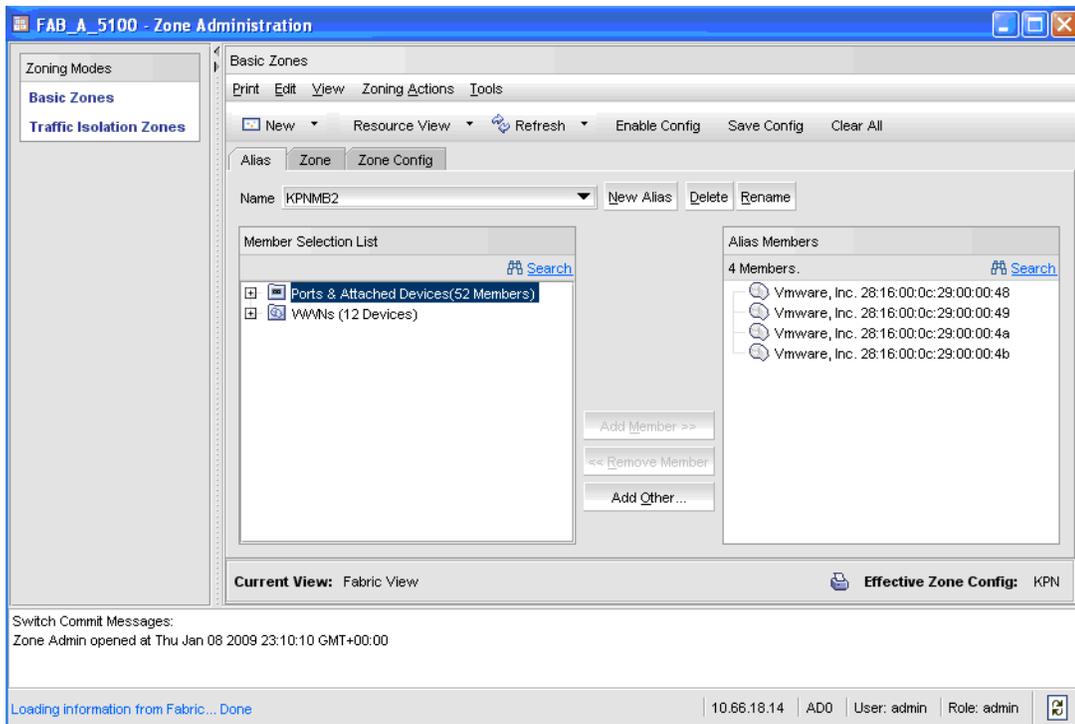
- d. With the Alias selected from the drop-down men, click **Add Other**.
- e. In the dialog box that displays, enter the **Port WWNs** generated by vCenter and click **OK**.

**NOTE:** When NPIV is enabled, four WWN pairs (WWPN & WWNN) are created for each VM. When the VM is powered on with NPIV enabled, it uses each of these WWN pairs in sequence to try to discover an access path to the storage. The number of virtual WWNs (Also know as a VPORT) that are instantiated equals the number of physical HBA ports present on the host, up to the

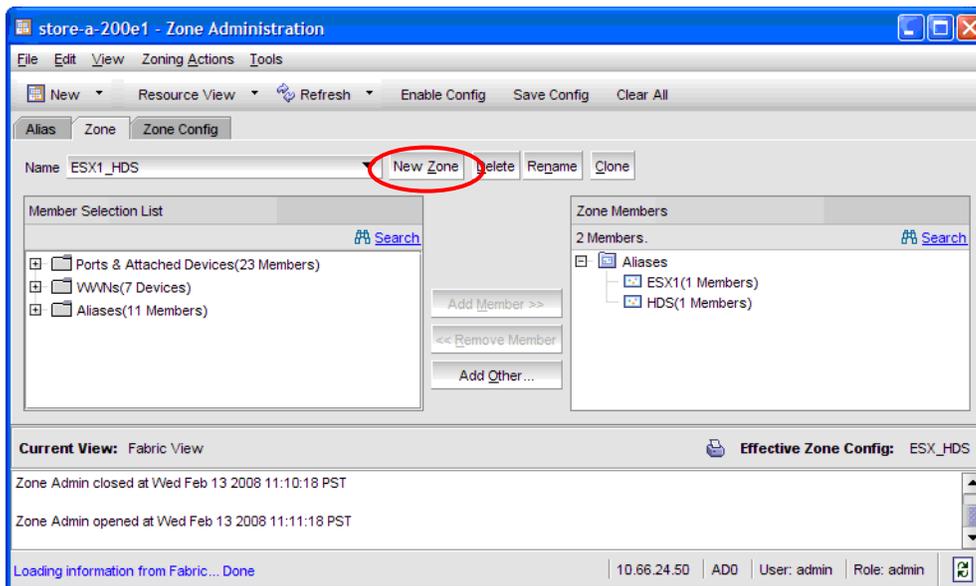
maximum of four per VM. A VPORT is created for each physical HBA port that has access to the storage. For example if you have only two physical ports connected to the ESX Server, then use the first two ports. If you have four ports on the ESX Server, then use all four port WWNs. If the HBA doesn't support NPIV or if zoning and LUN masking is not correctly setup, then the VM will fallback to the physical HBA port WWN.



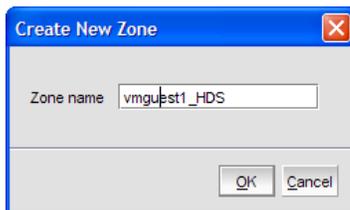
Once you have completed adding Port WWNs, the screen should look similar to this:



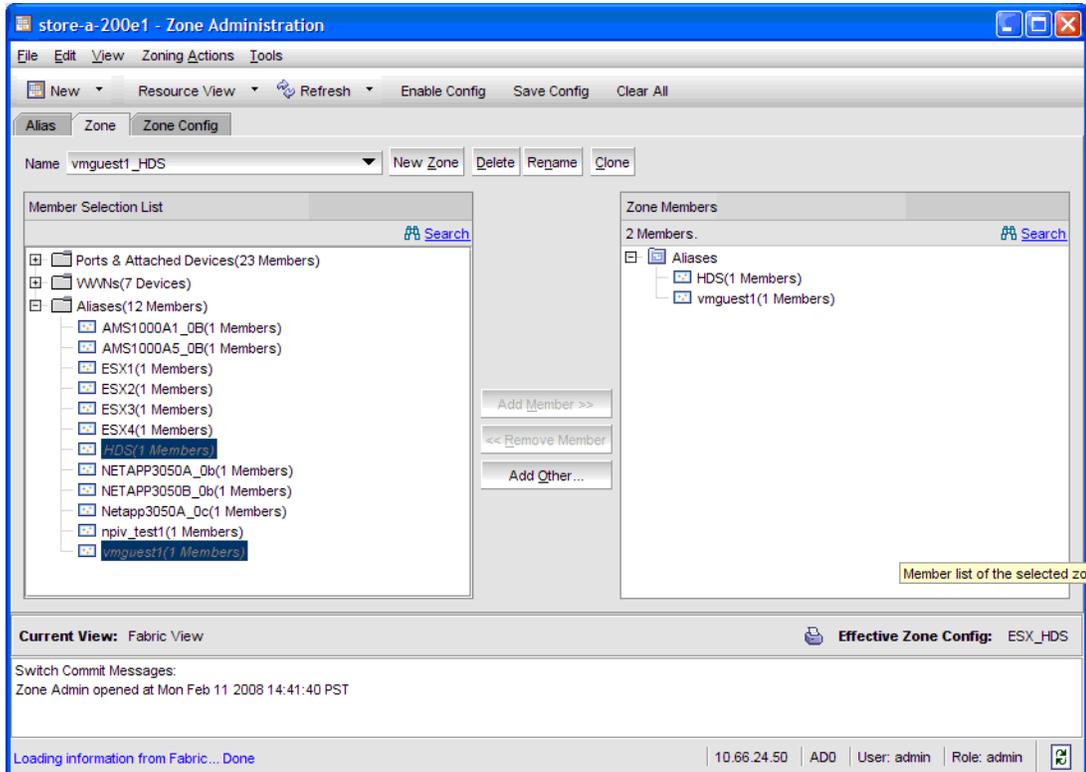
- f. Click the **Zone** tab and click **New Zone**.



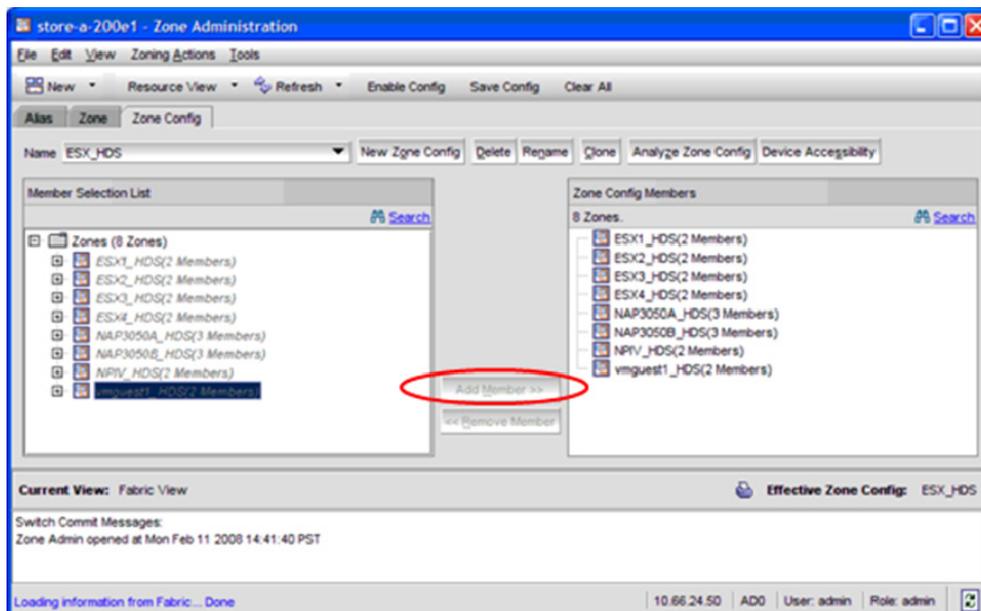
- g. Enter a name for the zone and click **OK**.  
**NOTE:** A SAN best practice is to zone one initiator to one target so that you increase security and to restrict interference from other hosts.



- h. Click the **plus sign (+)** to expand the Aliases folder in the Zone tab.
- i. From the Member Selection List, select the newly created alias (for example, vmguest1) and the alias of the storage array ports. Click the **Add Member** button, which adds the aliases to the zone.



- j. Click the **Zone Config** tab.



- i. From the Member Selection List, click the **plus sign (+)** next to Zones, select the newly created zone, and click the **Add Member** button.

- ii. Click **Save Config** at the top.  
Be patient as it could take about 15 – 30 seconds to commit. A status pane window at the bottom will let you know when the configuration has been committed.
- iii. Once the configuration has been committed, click **Enable Config**.  
This will also take 15 – 30 seconds to commit.

**To zone using the FOS CLI, perform all the steps for each switch in the fabric:**

- a. Telnet to the Brocade switch and create a new alias:

```
alicreate "vmguest1","28:25:00:0C:29:00:00:1b, 28:25:00:0C:29:00:00:1C;  
28:25:00:0C:29:00:00:1D; 28:25:00:0C:29:00:00:1E"
```

**NOTE:** When NPIV is enabled, four WWN pairs (WWPN & WWNN) are created for each VM. When the VM is powered on with NPIV enabled, it uses each of these WWN pairs in sequence to try to discover an access path to the storage. The number of virtual WWNs (Also know as a VPORT) that are instantiated equals the number of physical HBA ports present on the host, up to the maximum of four per VM. A VPORT is created for each physical HBA port that has access to the storage. For example if you have only two physical ports connected to the ESX Server, then use the first two ports. If you have four ports on the ESX Server, then use all four port WWNs. If the HBA doesn't support NPIV or if zoning and LUN masking is not correctly setup, then the VM will fallback to the physical HBA port WWN.

- b. Create a Zone to map the guest to storage: Once again the best practices is one initiator to one target.

```
zonecreate "vmguest1_HDS","vmguest1:HDS"
```

- c. Add the zone to the configuration:

```
cfgadd "mycfg", "vmguest1_HDS"
```

- d. Enable the configuration:

```
cfgenable "mycfg"
```

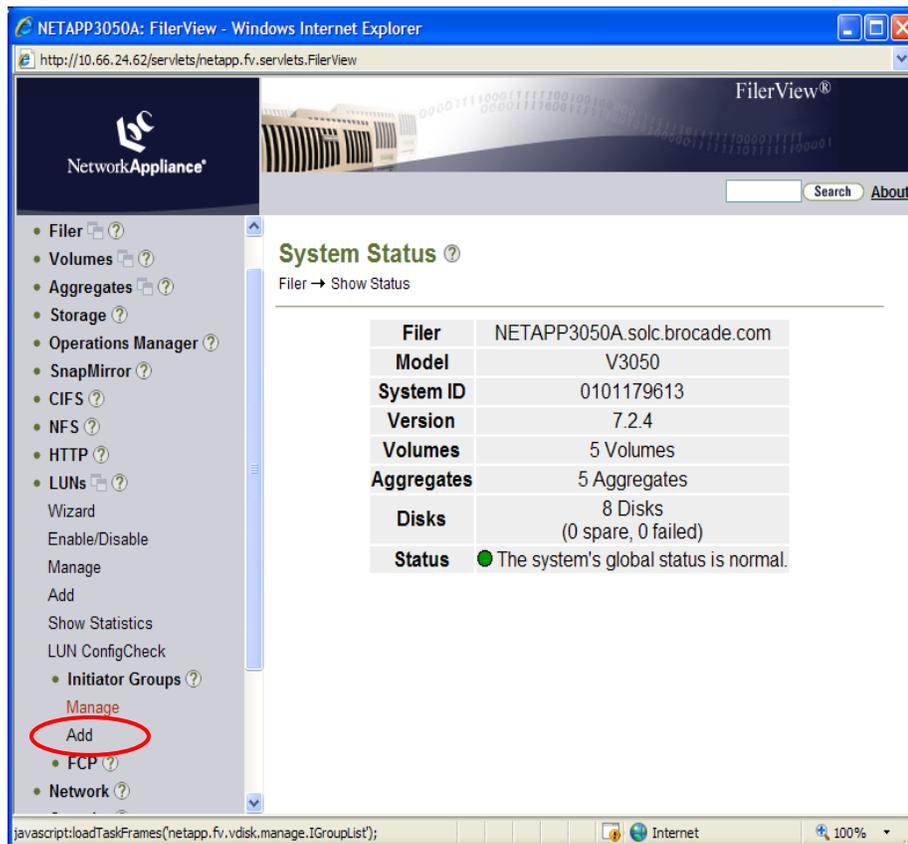
- e. Press Y to confirm you want to enable the configuration.

2. Once you have created the zone on the Brocade switch, log in to the storage array and add the Node WWN and Port WWNs of the VM to the LUN(s) that the VM will access. Depending on your storage array, the Node WWN may not be necessary.

You can configure back-end storage using NetApp FilerView or HDS AMS storage array.

#### Using NetApp FilerView:

- a. Log in to FilerView on the NetApp appliance (for example, [http://netapp3050a.brocade.com/na\\_admin](http://netapp3050a.brocade.com/na_admin))
- b. If you do not already have an Initiator Group created (An Initiator Group makes the LUNs visible to connected hosts), then create one by clicking **Add**, located below Initiator Groups in the left navigation pane.

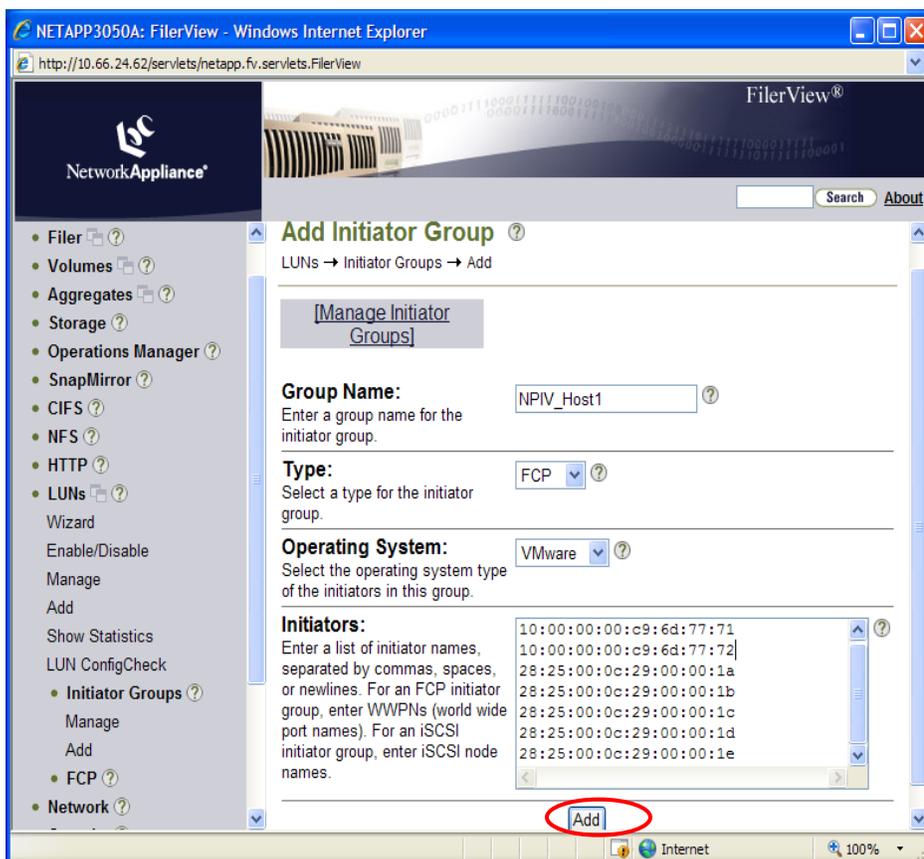


The screenshot shows the NetApp FilerView web interface in a Windows Internet Explorer browser window. The address bar shows the URL <http://10.66.24.62/servlets/netapp.fv.servlets.FilerView>. The page title is "NETAPP3050A: FilerView - Windows Internet Explorer". The main content area displays the "System Status" page for the Filer, with a link to "Show Status". A table provides the following system information:

<b>File</b>	NETAPP3050A.solc.brocade.com
<b>Model</b>	V3050
<b>System ID</b>	0101179613
<b>Version</b>	7.2.4
<b>Volumes</b>	5 Volumes
<b>Aggregates</b>	5 Aggregates
<b>Disks</b>	8 Disks (0 spare, 0 failed)
<b>Status</b>	● The system's global status is normal.

The left navigation pane contains a tree view with the following items: File, Volumes, Aggregates, Storage, Operations Manager, SnapMirror, CIFS, NFS, HTTP, LUNs, Wizard, Enable/Disable, Manage, Add, Show Statistics, LUN ConfigCheck, Initiator Groups (with a sub-item "Add" circled in red), FCP, and Network.

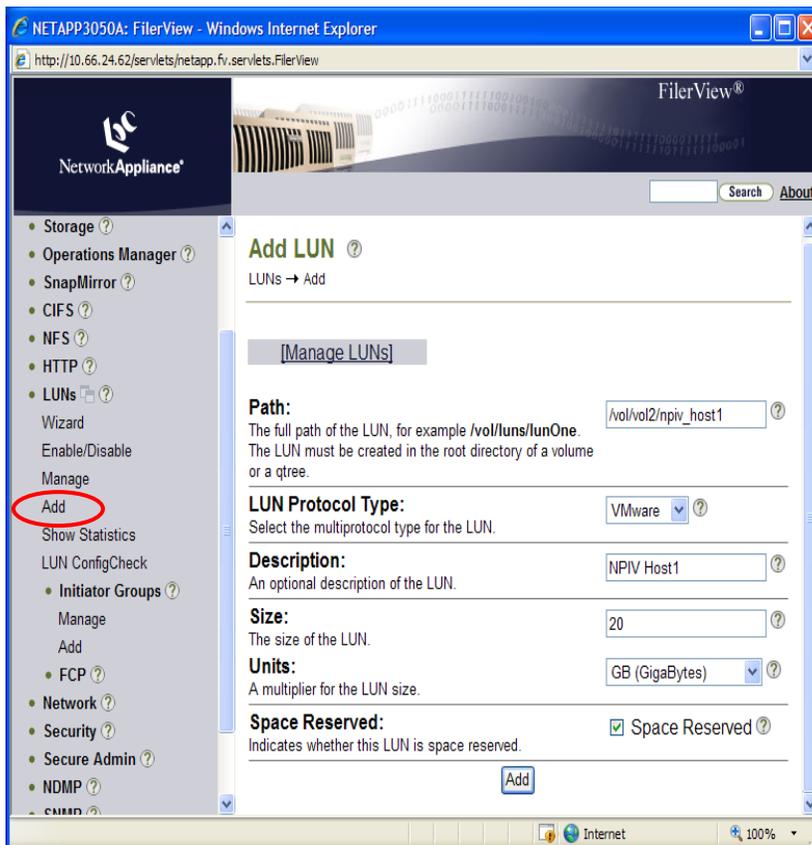
- c. In the Add Initiator Group screen, assign it a group name, a type (FCP) and an OS (VMware).
- d. Enter the VMware guest Node WWN and Port WWNs that you received from vCenter and the physical HBA port WWNs of the ESX host in the Initiators section. If you are using VMotion, be sure that you add the other ESX hosts physical HBAs WWNs that participate in the VMotion process. This will ensure that the NPIV connection is maintained when the VMotion takes place. Otherwise the VM will default back to the physical HBA WWN and will not be using NPIV. Also
- e. Click **Add**



**REMINDER:** The number of physical HBA ports in the ESX Server device determines how many Port WWNs you need to enter in the initiators group.

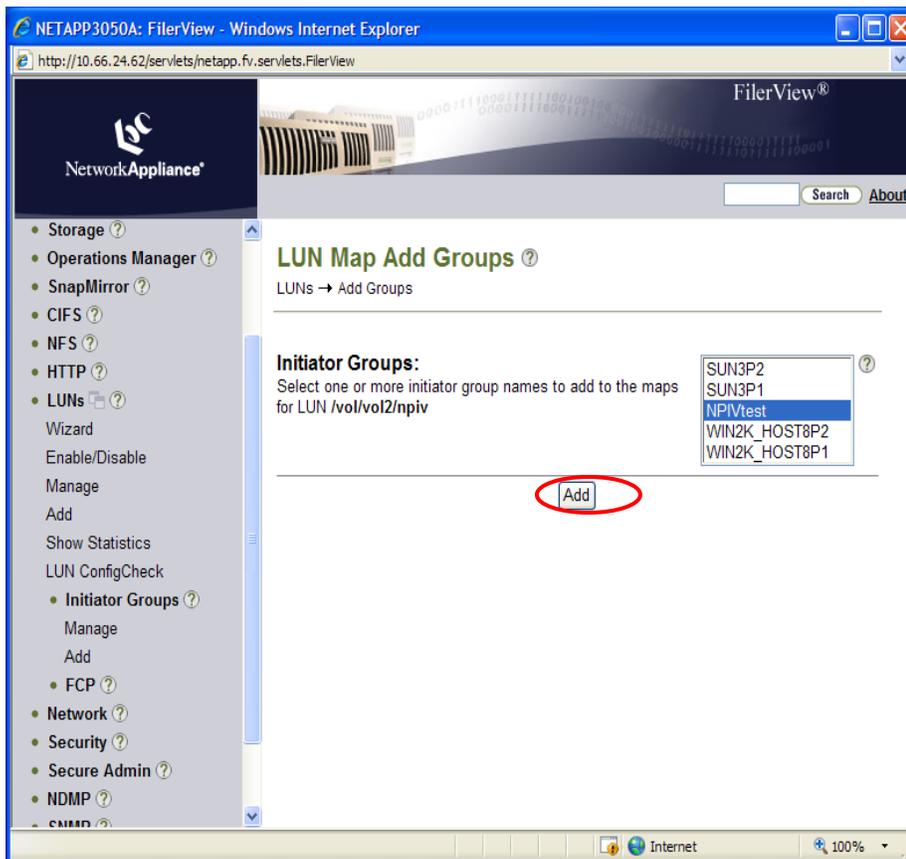
- f. Fill in the requested information and click **Add**.

- g. Select **Add** under LUNs in the left navigation pane.



- h. In the Add LUN screen, enter the required information and click **Add**.
- i. Once the LUN is created, click **Manage** under LUNs in the left navigation panel and select the LUN you created.
- j. Click **Map LUN**.
- k. Click **Add Groups to Map**.

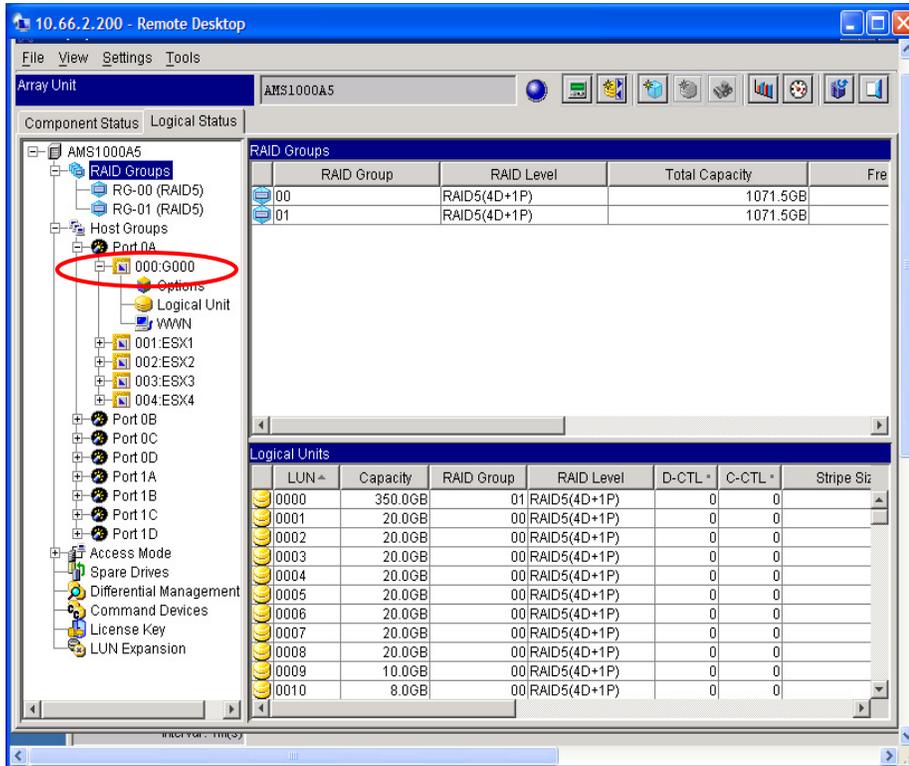
- I. Select the Initiator Group you just created and click **Add**.



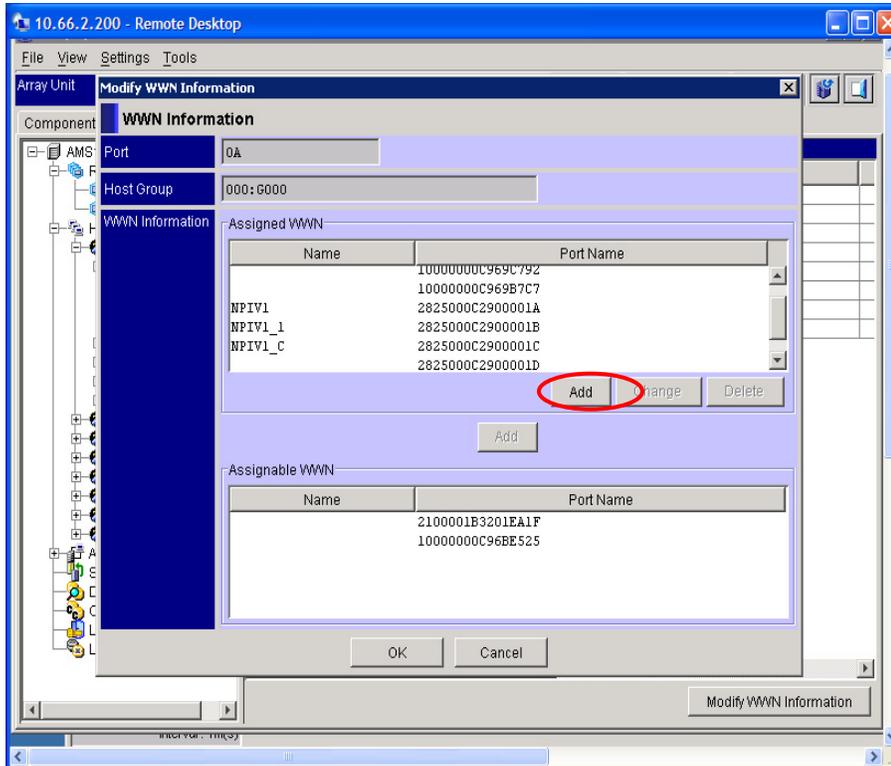
- m. Now inform your VMware Administrator that (s)he can power on the VM.

**Using HDS AMS Storage Navigator:**

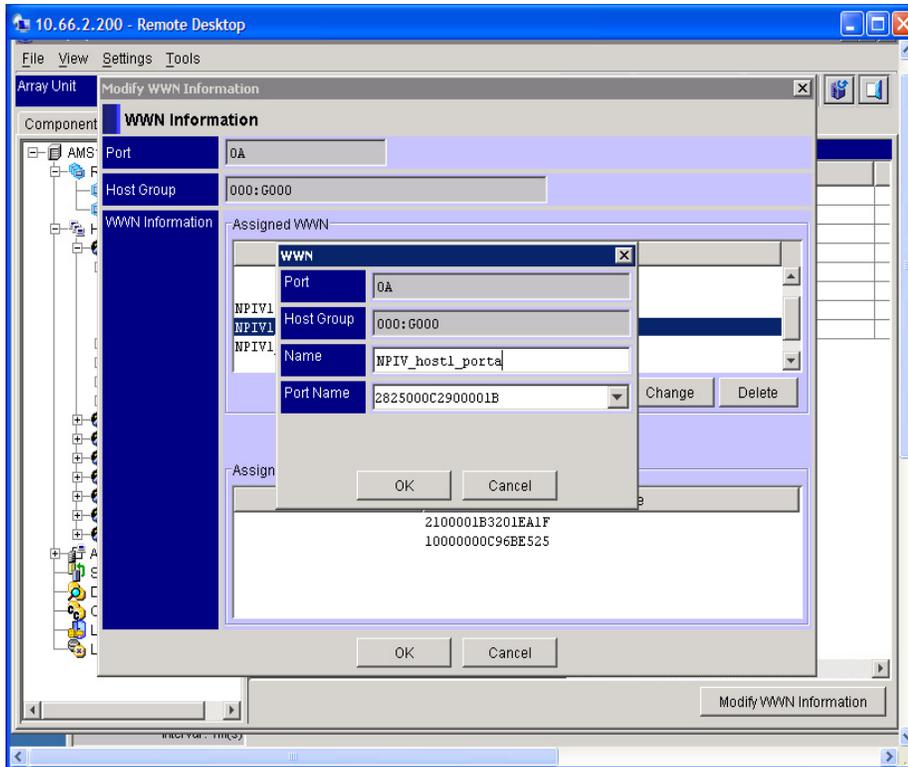
- a. In the Logical Status tab, select the port group in which the LUNs are located in the left navigation pane.



- b. Click **WWN** to select the WWNs that are available to map to this port group.
- c. Click **Modify WWN Information** at the bottom right.
- d. In the WWN Information screen, click **Add**.



- e. In the WWN dialog box, add the Virtual WWNs that the Server Administrator gave you. Also add the physical HBA WWNs of the ESX Server host and the physical HBA WWNs of the other ESX Server if you are using Vmotion.
- f. Give the WWN a user-friendly name, enter the port WWN without any colons, and click **OK**.



- g. Click **OK** again
- h. Now inform your VMware Administrator to power on the Virtual Machine.

## Verifying that the VM Is Functioning Correctly

1. Log in to the Brocade switch.

2. Issue the **switchShow** command.

You should see the NPIV ports with a number greater than 1 (“one”) as shown below. If you see only one NPIV port, then the setup has not been successful.

```
FAB_B_5100:admin> switchshow
switchName:      FAB_B_5100
switchType:      66.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    1
switchId:        fffc01
switchWwn:       10:00:00:05:1e:58:03:e8
zoning:          ON <KPN>
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 1

Area Port Media Speed State      Proto
=====
0 0 id N4 Online F-Port 4 NPIU public
1 1 id N4 Online F-Port 3 NPIU public
2 2 id N4 Online F-Port 3 NPIU public
3 3 id N4 Online F-Port 2 NPIU public
4 4 id N4 Online F-Port 50:01:43:80:01:34:a2:19
5 5 id N4 Online F-Port 50:01:43:80:01:34:a2:1d
```

You can also issue the **nsShow** command to show the ports logged in to the fabric. You should see the Virtual Machine NPIV ports.

```
N 010003; 3;28:16:00:0c:29:00:00:4e;28:16:00:0c:29:00:00:4c; na
FC4s: FCP
NodeSymb: [43] "Emulex LPe1105-HP FU2.72A2 DUelx_7.4.0.13-2"
Fabric Port Name: 20:00:00:05:1e:58:03:e8
Permanent Port Name: 20:11:00:05:1e:05:cc:60
Port Index: 0
Share Area: No
Device Shared in Other AD: No
Redirect: No
```

You could also use the service console on the ESX Server and check the /proc nodes of the HBA to get the details. The procedure to find this information is documented in the “Server Administrator Tasks” section in this document.

