



Solution Implementer's Series

Deploying Virtual Fabric and 8Gb/s Fibre Channel with IBM BladeCenter and VMware



**Solution
Implementer's
Series™**

Create robust, highly-available VMware ESXi 4.1 environments with best-of-breed Virtual Fabric and Fibre Channel storage networking and IBM BladeCenter HX5 blade servers

vmware®
READY

EMULEX®

Table of contents

Emulex Solution Implementer's Series.....	3
Solution deployment overview.....	3
Pre-installation	5
Updating the firmware.....	5
Configuring network connectivity.....	7
ESXi network requirements.....	7
Network support in a BladeCenter HX5 blade server	9
Implementing network connectivity in the POC environment	9
Configuring the Virtual Fabric switch.....	11
Configuring Fibre Channel connectivity.....	16
Configuring fabric zoning	16
Configuring the storage array.....	18
Configuring BfS.....	21
Configuring RAID on local storage.....	26
Deploying ESXi 4.1	27
Prerequisites.....	27
Using a DVD image to deploy ESXi 4.1	28
Post-installation considerations.....	29
Managing the solution.....	30
Using the OneCommand Manager plug-in for vCenter	31
Using the IBM System Storage Plug-in for vCenter	36
Advanced topics.....	41
Configuring multipathing	41
Setting up the storage stack for multipathing	42
Avoiding LUN thrashing	45
Validating the multipathing implementation.....	46
Summary.....	50
Appendix A – Bill of materials.....	51
Appendix B – Installing the IBM System Storage plug-in	52
Appendix C – Installing drivers and CIM providers.....	55
Installing the VFA NIC driver	55
Updating the 8Gb/s Fibre Channel expansion card driver.....	55
Installing the CIM providers.....	56
More information.....	57



Emulex Solution Implementer's Series

This document is part of the Emulex Solution Implementer's Series, providing Implementers (IT administrators and system architects) with solution and deployment information on popular server and software platforms. As a leader in I/O adapters – Fibre Channel (FC), Ethernet, iSCSI and Fibre Channel over Ethernet (FCoE) – the Emulex technology team is taking a lead in providing guidelines for implementing I/O for these solutions.

This solution described in this white paper was implemented and validated in May/June 2011.

Solution deployment overview

This white paper provides guidelines for deploying the IBM BladeCenter 8Gb/s Fibre Channel proof of concept (POC) environment shown in Figure 1.

The following topics are included:

- **Pre-installation:** Updating IBM system firmware and updating the server's Unified Extensible Firmware Interface (UEFI) for VMware® ESXi 4.1
- **Configuring network connectivity:** ESXi network requirements; network support in the IBM BladeCenter HX5 blade server; implementing network connectivity; configuring the 10GbE switch
- **Configuring Fibre Channel connectivity:** Setting up fabric zoning, storage and boot from SAN; configuring RAID on the server to support local storage
- **Installing ESXi:** Installing ESXi 4.1 from IBM's customized ESXi 4.1 DVD image, along with pre- and post-installation considerations
- **Managing the solution:** Using plug-ins from Emulex and IBM to enhance the manageability of the POC environment
 - **Changing the LUN queue depth:** Example of solution management: adjusting the HBA queue depth to enhance storage performance
- **Advanced topics**
 - **Implementing multipathing:** Creating multiple paths between the host and shared storage to enhance reliability



Note

As a matter of best practice for all deployments, Emulex recommends implementing a POC using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a POC, contact an IBM services representative or your IBM partner.

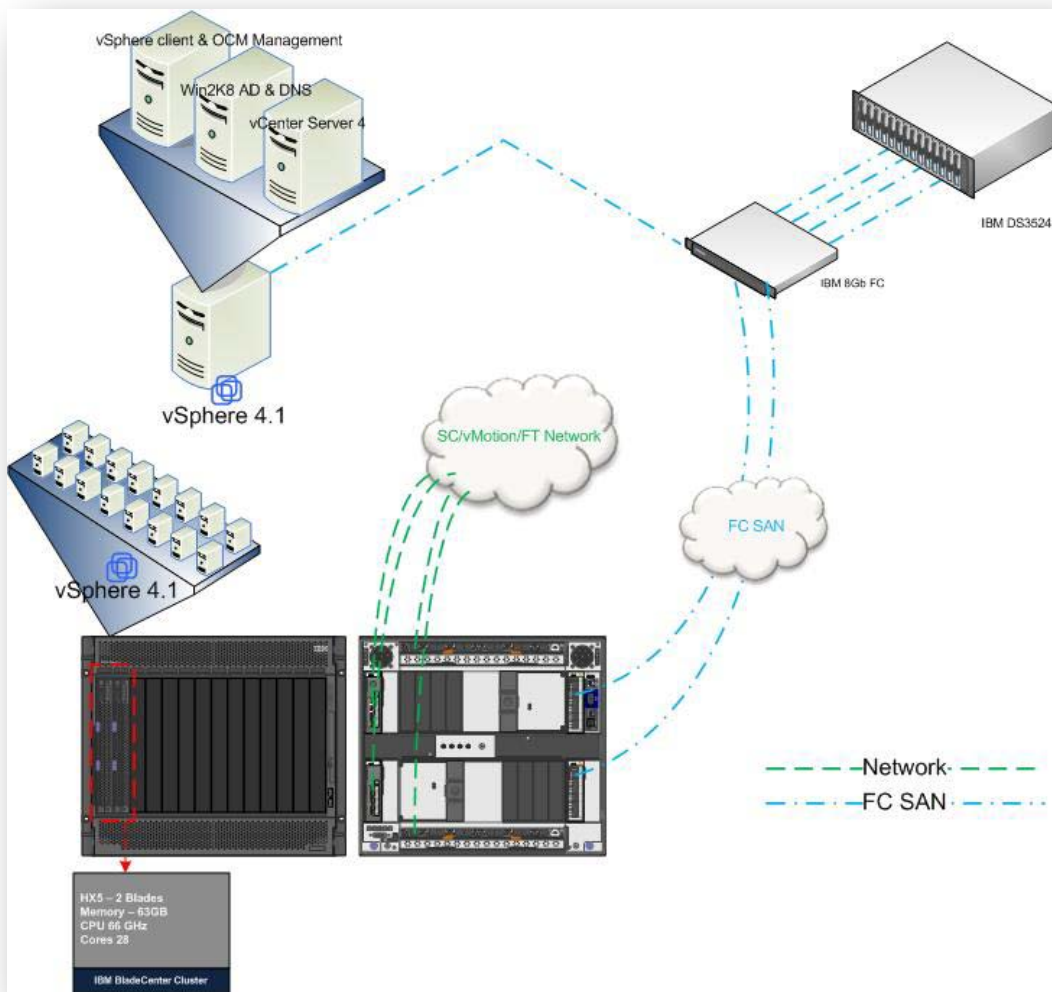


Figure 1 – High-level view of the IBM BladeCenter 8Gb/s Fibre Channel POC environment



The POC environment included the following key components:

- Two IBM BladeCenter HX5 Blade Servers¹ running ESXi 4.1
- IBM System Storage DS3524 Disk System (storage array)
- For network connectivity:
 - Emulex Virtual Fabric Adapter² (VFA)
 - BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter (Virtual Fabric switch)
- For Fibre Channel connectivity:
 - Two Emulex 8Gb Fibre Channel Expansion Cards³ (8Gb Fibre Channel expansion card)
 - Two Brocade 8Gb SAN Switch Modules for IBM BladeCenter (8Gb SAN switch)

Pre-installation

Before deploying ESXi 4.1 on the BladeCenter HX5 blade server, you should perform the following actions:

- Configure the BladeCenter H chassis
- Create and deploy Open Fabric manager configuration files
- Update all system firmware using the IBM UpdateXpress CD (see below)

Note

UpdateExpress 3.10 was used for this BladeCenter HX5 blade server.

Updating the firmware

You can use the IBM BladeCenter H Chassis Advanced Management Module (AMM) to update the firmware on the BladeCenter HX5 blade server. The procedure is as follows:

1. From a management workstation, download the **ibm_utl_uxsp_phsp03a-3.10_virtual_32-64.zip** UpdateXpress System Pack (for ESX 3/4, VMware 3x64). Place this file in the **/tmp** directory of the ESXi host using a tool such as WinSCP.
2. Extract the zip file to the **/temp** directory.
3. Log into the AMM's web interface on BladeCenter.
4. Select **MM control**→**Network protocols** and ensure Trivial File Transfer Protocol (TFTP) is enabled on the AMM.

¹ Only one blade server was configured.

² Referred to by IBM as a CFFh

³ Referred to by IBM as a CIOv



5. Select **Blade task**→**Firmware update** and select the target blade.
6. Use the browse button to point to the update file (**ibm_fw_uefi_hie130a-130_linux_32-64.bin**).
7. Click the update button to update the UEFI firmware.
8. Reboot the system to activate the new firmware.

Next, you should use IBM Open Fabric Manager (OFM) to create configuration files for the blade server.

Using Open Fabric Manager

By assigning Ethernet MAC and Fibre Channel WWN addresses to the BladeCenter slots that may be occupied by blade servers, OFM enables you to quickly replace and recover blade environments.

With OFM, system administrators need not worry about WWN or MAC addresses changing due to the replacement of a faulty device. OFM has its own MAC and WWN addresses that it manages for each slot, allowing you to replace a faulty device without affecting the consistency of addressing.

OFM lets you pre-assign WWN and MAC addresses, including boot targets, for up to 100 enclosures and 1400 blade servers. This can be even accomplished without any blade servers installed on the chassis.

IBM provides the following OFM offerings:

- Basic OFM functionality is delivered by the AMM
- Additional capabilities are delivered by the OFM-Advanced upgrade to the AMM

Creating a configuration file

Using OFM to create a configuration for a particular blade server bay is a simple process and is not described in this white paper. Figure 2 shows a sample configuration.

Creating a configuration file is highly recommended, since this allows OFM to assign all WWN and MAC addresses for the bay. Thus, if boot from SAN (BfS) has been implemented, you can easily replace one blade server with another similar device, with no re-configuration required.



IBM BladeCenter® H Advanced Management Module Welcome USERID [About](#) [Help](#) [Logout](#)

Bay 1: SN#YK11836BW1N8

- Monitors
- System Status
 - Event Log
 - LEDs
 - Power Management
 - Hardware VPD
 - Firmware VPD
 - Remote Chassis
- Blade Tasks
 - Power/Restart
 - Remote Control
 - Firmware Update
 - Configuration
 - Serial Over LAN
 - Open Fabric Manager
- I/O Module Tasks
- MM Control
- Service Tools
 - AMM Service Data
 - Blade Service Data
 - AMM Status
 - Service Advisor
- Scalable Complex

Open Fabric Manager Parameters for Bay: 1

Open Fabric Manager Mode: Enabled

Open Fabric Manager Profile: TempProfile BC-1Slot-1

Open Fabric Manager Parameters Version: Version 2

System Mgmt Processor Open Fabric Manager Capable: Yes

BIOS Open Fabric Manager Capable: Yes

Open Fabric Manager Status: n/a

Ethernet

OFM MAC Configuration

Slot Offset	Port	MAC Address A	VLAN tag A	MAC Address B	VLAN tag B	MAC range	Device supports OFM	Address Status
0	1	00:1a:64:76:00:00				No	Yes	Not used
0	2	00:1a:64:76:00:01				No	Yes	Not used
0	3	00:1a:64:76:00:02				No	Yes	Not used
0	4	00:1a:64:76:00:03				No	Yes	Not used
0	5	00:1a:64:76:00:04				No	Yes	Not used
0	6	00:1a:64:76:00:05				No	n/a	Not used
0	7	00:1a:64:76:00:06				No	Yes	Not used
0	8	00:1a:64:76:00:07				No	n/a	Not used

Fibre Channel

Slot Offset	Port	WWNN	WWPN	Target To Boot From	Device supports OFM	Address Status
0	3	2f:fe:00:00:c9:00:00:00	2f:fc:00:00:c9:00:00:00	none	Yes	Not used
0	4	2f:ff:00:00:c9:00:00:00	2f:fd:00:00:c9:00:00:00	none	Yes	Not used
0	5	2f:fe:00:00:c9:00:00:01	2f:fc:00:00:c9:00:00:01	none	Yes	Not used
0	6	2f:ff:00:00:c9:00:00:01	2f:fd:00:00:c9:00:00:01	none	n/a	Not used
0	7	2f:fe:00:00:c9:00:00:02	2f:fc:00:00:c9:00:00:02	none	Yes	Not used
0	8	2f:ff:00:00:c9:00:00:02	2f:fd:00:00:c9:00:00:02	none	n/a	Not used

Fibre Channel Target

Index	FC Target WWPN	FC Lun	Address Status
1	50:01:43:80:01:3b:f3:88	1	Not used

[Close](#)

Figure 2. Sample blade server configuration for Bay 1

Configuring network connectivity

This section outlines the network requirements generated by ESXi and the availability of NICs in a BladeCenter HX5 blade server. It also describes how network connectivity was implemented in the POC environment.

ESXi network requirements

ESXi can generate a range of network requirements, including the following:

- **Service console:** Allows vSphere client and vCenter Server to communicate with the host; a minimum of 100Mb/s is required for ESXi to communicate with the VMkernel



- **VM traffic:** Supports virtual machine (VM) network traffic; ideally, you would dedicate two 1GB NICs for VM traffic to enable load balancing and NIC teaming
- **vMotion:** The VMware vMotion (vMotion) feature enables the live migration of VMs from one ESXi host to another, supporting disaster recovery, planned downtime and fault tolerance
- **VMware Fault Tolerance:** The VMware Fault Tolerance (Fault Tolerance) feature enables two identical VMs on different ESXi hosts to remain in lockstep with each other; in order to minimize latency, you should use a dedicated NIC

Note

Like vMotion traffic, network traffic for the Fault Tolerance feature should be separated on to a different subnet.

- **iSCSI:** Connects iSCSI storage devices on the network
-

Note

If you are using software iSCSI, you may need to provide a path to shared LUNs that is separate from other VM traffic.

Table 1 provides recommendations for the minimum number of NIC ports and network throughput required by ESXi.

Table 1. Minimum network requirements for an ESXi 4.1 host with network redundancy

Network type	NIC ports	Throughput
Service console	2	100Mb/s
VM	2	1Gb/s-10Gb/s
vMotion	2	1Gb/s-2Gb/s
Fault Tolerance	2	1Gb/s-2Gb/s
iSCSI	2	1Gb/s-10Gb/s

While you may not need all these networks for a particular ESXi host, the service console, VM and vMotion networks are typically implemented.

If desired, a single dual-port VFA could be used to support the service console, VM traffic and vMotion.



Network support in a BladeCenter HX5 blade server

A BladeCenter HX5 blade server can be configured with an Emulex 10GbE VFA, which is supported by ESXi 4.1. This VFA can also be configured the following:

- Network adapters
- iSCSI and FCoE initiators

In addition, the server also features LAN-On-Motherboard (LOM), with two 1GbE network adapter ports.

Implementing network connectivity in the POC environment

For simplicity, Emulex implemented network connectivity in the POC environment using the following 10GbE components:

- **NIC ports in the BladeCenter HX5 blade server**
 - Two 1GbE LOM ports – Used for the Service Console
 - Eight virtual NIC (vNIC) ports – Used for VM traffic, vMotion traffic and to support Fault Tolerance
- **IBM Virtual Fabric 10Gb Switch Module for IBM BladeCenter:** Ports three and four on a Virtual Fabric switch are enabled to support connections from the host.

Table 2 outlines network connectivity in the POC environment.

Table 2. Networks in the POC environment

Network type	NIC ports	Names	Bandwidth
Service console (Management Network)	2 (LOM)	vmnic0, 1	1Gb/s ⁴
VM	4	vmnic2, 3, 4, 5	2.5Gb/s per vmnic
vMotion	2	vmnic6, 7	2.5Gb/s
Fault Tolerance	2	vmnic8, 9	2.5Gb/s

⁴ NIC teaming was configured



Figure 3 presents a view of these networks from vCenter Server.

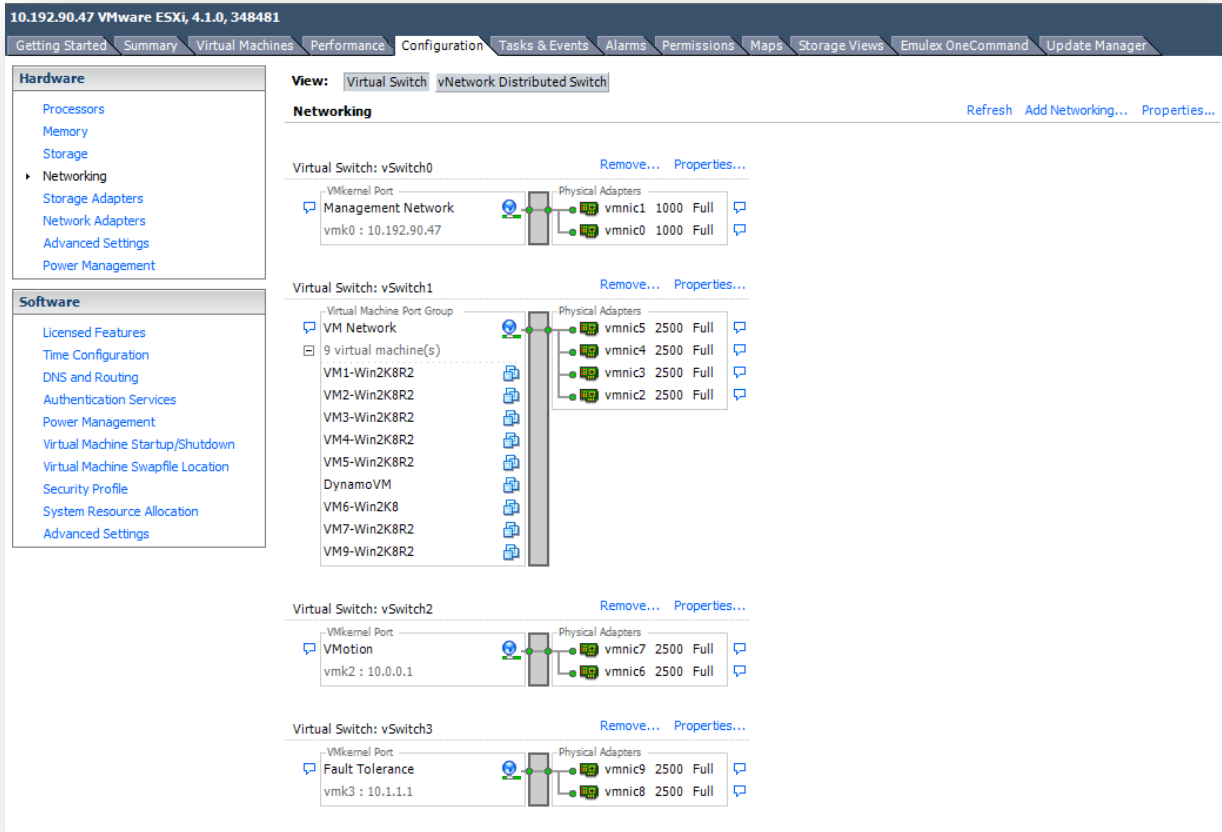


Figure 3. Showing the vNICs and LOM ports used to provides network connectivity on the HX5 Blade Server

The following section describes how to set up a Virtual Fabric switch to support the network connectivity required by the POC environment.



Configuring the Virtual Fabric switch

In general, the Virtual Fabric is configured⁵ via the Virtual Fabric switch rather than the VFA or the blade server.

Switch ports (10Gb/s) can be carved up to create four virtual NICs (vNICs) that connect to and are recognized by a physical NIC (pNIC) on the VFA. Each vNIC is configured via the switch to provide a bandwidth of 10Mb/s – 10Gb/s.

The switch is also used to create virtual groups from the vNICs, allowing you to control which vNICs appear on which external network and how much bandwidth is available on each network. In the POC environment, for example, Emulex created the following external connections via the Virtual Fabric:

- **VM traffic:** four vNICs (4 x 2.5Gb/s)
- **vMotion traffic:** two vNICs (2 x 2.5Gb/s)
- **Fault Tolerance traffic:** two vNICs (2 x 2.5Gb/s)

Before you can configure the Virtual Fabric switch, you must understand the following:

- How VFA PCIe Function IDs map to VFA ports
- How I/O bay modules on the Virtual Fabric switch map VFA ports

Table 3 outlines these mappings.

Table 3. VFA to Virtual Fabric switch mapping

VFA		Virtual Fabric switch		
PCIe Function ID	Port	I/O bay	Port alias ⁶	vNIC
0	0	7	INTx.1	1
2	0	7	INTx.2	2
4	0	7	INTx.3	3
6	0	7	INTx.4	4
1	1	9	INTx.1	1
3	1	9	INTx.2	2
5	1	9	INTx.3	3
7	1	9	INTx.4	4

⁵ For more information, refer to the IBM Redpaper, [“IBM BladeCenter Virtual Fabric Solutions.”](#)

⁶ Where “x” denotes the particular blade slot



Figure 4 shows how the Virtual Fabric in the POC environment was configured.

Note

The LOM was used for service console traffic.

If you are using software iSCSI, you may need to provide a path to shared LUNs that is separate from other VM traffic.

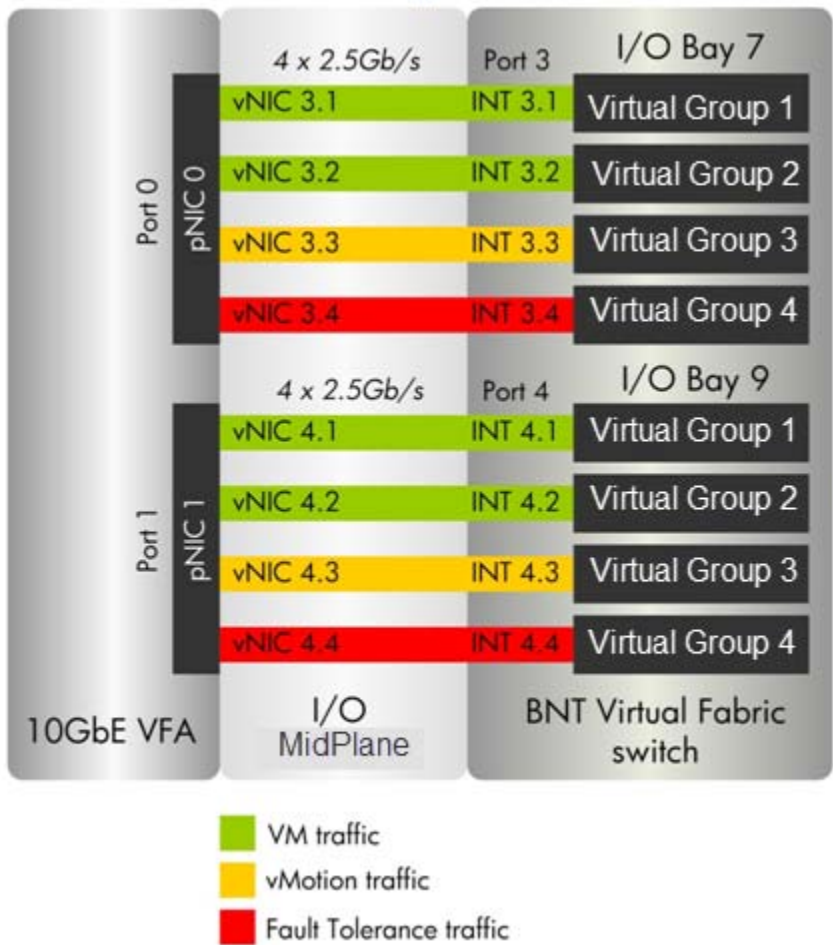


Figure 4. Virtual Fabric configuration for the POC environment



Perform the following steps to set up the Virtual Fabric switch:

1. Update the firmware for the Virtual Fabric switch to the latest level (6.5.2.3 for the POC environment).
Firmware downloads are available from the [BLADE Network Technologies](#) site⁷.
2. Log in to the Virtual Fabric switch on I/O Bay 7 using the IP address for the AMM.
3. Expand the selections shown on the screen by selecting **+** from **I/O Module Tasks**.
4. Select **Configuration** to show slots that are occupied.
5. Select the **Slot 7** tab.
6. After the screen has refreshed to show the appropriate IP address for the module in I/O Bay 7 (Slot 7), select **Advanced Options**.
7. Start a web session.
8. Log in to the Virtual Fabric.
9. Use BLADE OS to enable the ports needed to provide host connectivity (that is, Ports 3 and 4 on the Virtual Fabric switch), as shown in Figure 5.

Select the **CONFIGURE** tab to allow you to edit the settings.

Note that in the POC environment, pNICs 0 and 1 are physically connected to Ports 3 and 4 on the Virtual Fabric switch.

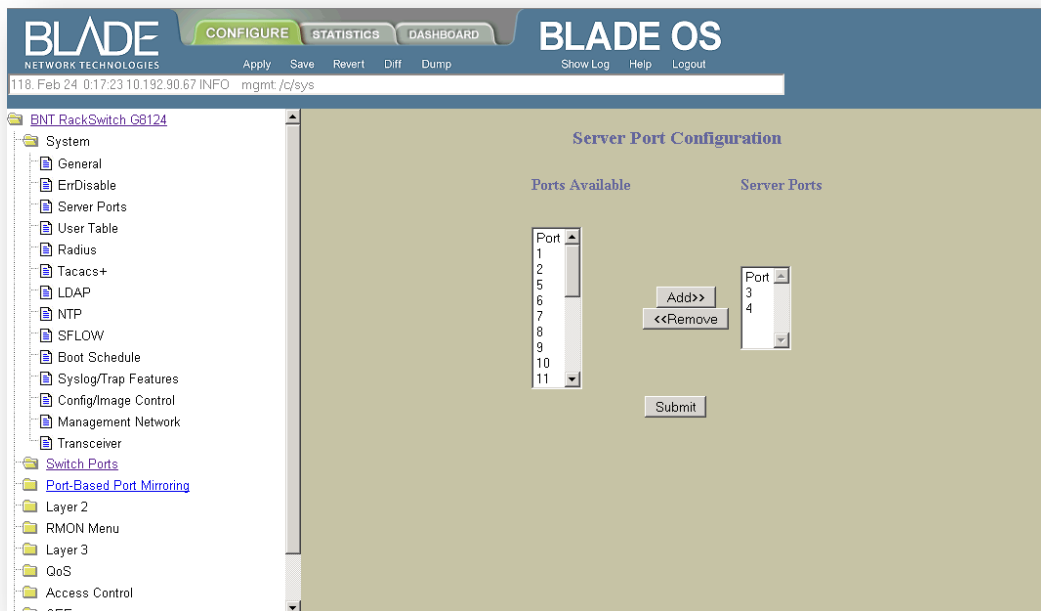


Figure 5. Using the BLADE OS user interface to enable server Ports 3 and 4 on the Virtual Fabric switch

⁷ Registration required



10. After configuring the switch ports, select **Submit** and then **Apply** to save the changes.

Note

Depending on the configuration, additional switch ports may need to be enabled to allow the VFA to log on to the network.

11. Select the **Virtualization** folder from the BLADE OS table of contents, followed by **General**.

12. Select **On** as the setting for **Global vNIC**; select **Submit** and then **Apply** to enable vNIC mode.

13. Select **vNICs** and enable the vNICs needed to accommodate network traffic in the POC environment (vNICs 3.1 – 3.4), as shown in Figure 6.

Select the maximum bandwidth for each vNIC. For example, a **Max Bandwidth in increments of 100Mbps** setting of **25** limits the particular vNIC to 2.5Gb/s, which is the requirement for the POC environment.

Select **Submit** and then **Apply**.



Figure 6. Using the BLADE OS user interface to configure vNIC 3.1 (INT 3.1)

The first pass through the above procedure is used to configure I/O Bay 7. You should now repeat Steps 2 – 13 for I/O Bay 9, logging in to the Virtual Fabric switch on I/O Bay 9 and selecting Slot 9 rather than Slot 7. Enable vNICs 4.1 – 4.4.



Setting up virtual groups

After you have enabled and configured the vNICs, you should allocate them to vNIC groups that provide the bandwidth required to support your network traffic.

vNIC groups can also include internal (INT) and external (EXT) switch ports and trunks, in effect creating separate virtual networks within the switch.

In the POC environment, Emulex created three vNIC groups (1 – 3) to separate VM, vMotion and Fault Tolerance traffic. The configuration procedure for setting up vNIC Group 1 (VM traffic) was as follows:

1. In the **CONFIGURE** tab of BLADE OS, select **vNIC Groups** from the **Virtualization** folder.
2. Click **1** to select vNIC Group 1.
3. As shown in Figure 7, change the **vNIC Group State** to **enabled**.
4. Add a **vNIC Group Vlan** name (**1701** in the POC environment).
5. Accept the default value for **Uplink Failover State** (that is, **disabled**).
6. Accept the default value for **Uplink Port Trunk** (that is, **None**).
7. Select the vNICs you wish to add to vNIC Group 1.
8. Select **Submit** and then **Apply**.

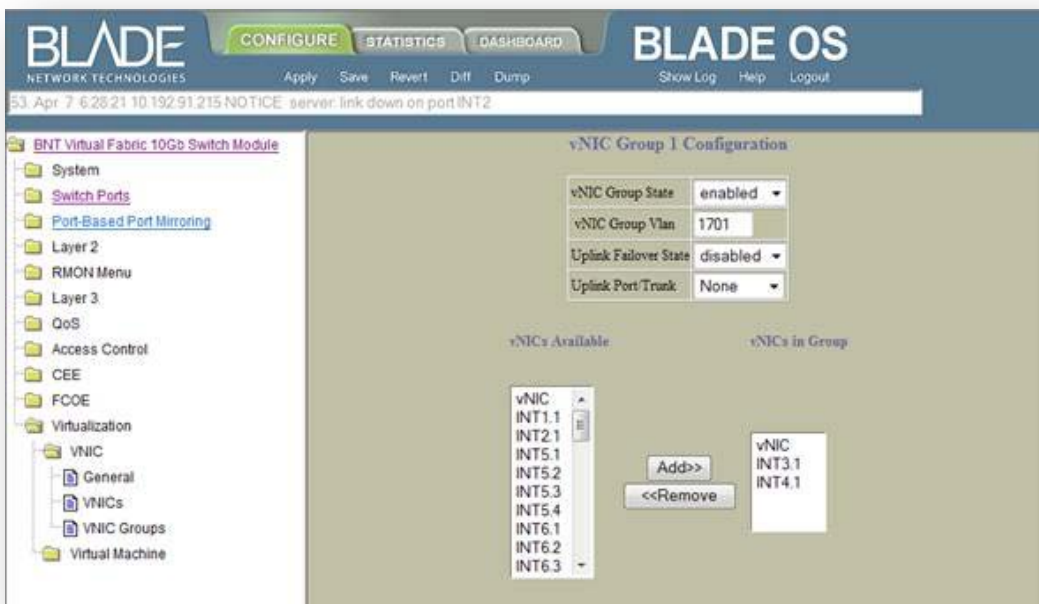


Figure 7. Using the BLADE OS user interface to add vNICs 3.1 and 3.2 to vNIC Group 1 (Note that, in the POC environment, vNICs 4.1 and 4.2 are also added to this group.)



Configuring Fibre Channel connectivity

Fibre Channel connectivity in the POC environment is achieved using the following 8Gb/s components:

- **8Gb/s SAN switch:** To provide redundancy, the POC environment includes two 8GB/s SAN switches that are supported by ESXi.

It is a good practice to use multiple 8Gb/s SAN switches; not only do you enhance redundancy but you can also restrict access to certain LUNs, if desired.

- **8Gb/s PCIe HBA:** An 8Gb/s Fibre Channel expansion card is installed in the BladeCenter HX5 blade server. Its two ports are connected to fabric switches; one port is configured for BfS

This section provides instructions for setting up the following:

- Fabric zoning
- Storage
- BfS

Configuring fabric zoning

Fabric zones are often implemented in an ESXi environment to enhance manageability while providing support for advanced features such as vMotion and Fault Tolerance that require multiple hosts to access the same LUN.

Zones can also enhance security. For example, consider what might happen if you were to connect a new server to the same SAN switch as an existing ESXi host. Without zoning or some other security measure, the new server would be able to access the same storage as the ESXi host and could potentially overwrite the existing file system, obliterating VM data and files. Since the POC environment features two 8Gb/s HBA ports, you should ideally provide two – or more – fabric switches, configuring each switch with a zone that includes a particular HBA port.



Figure 8 shows the zoning used in the POC environment.

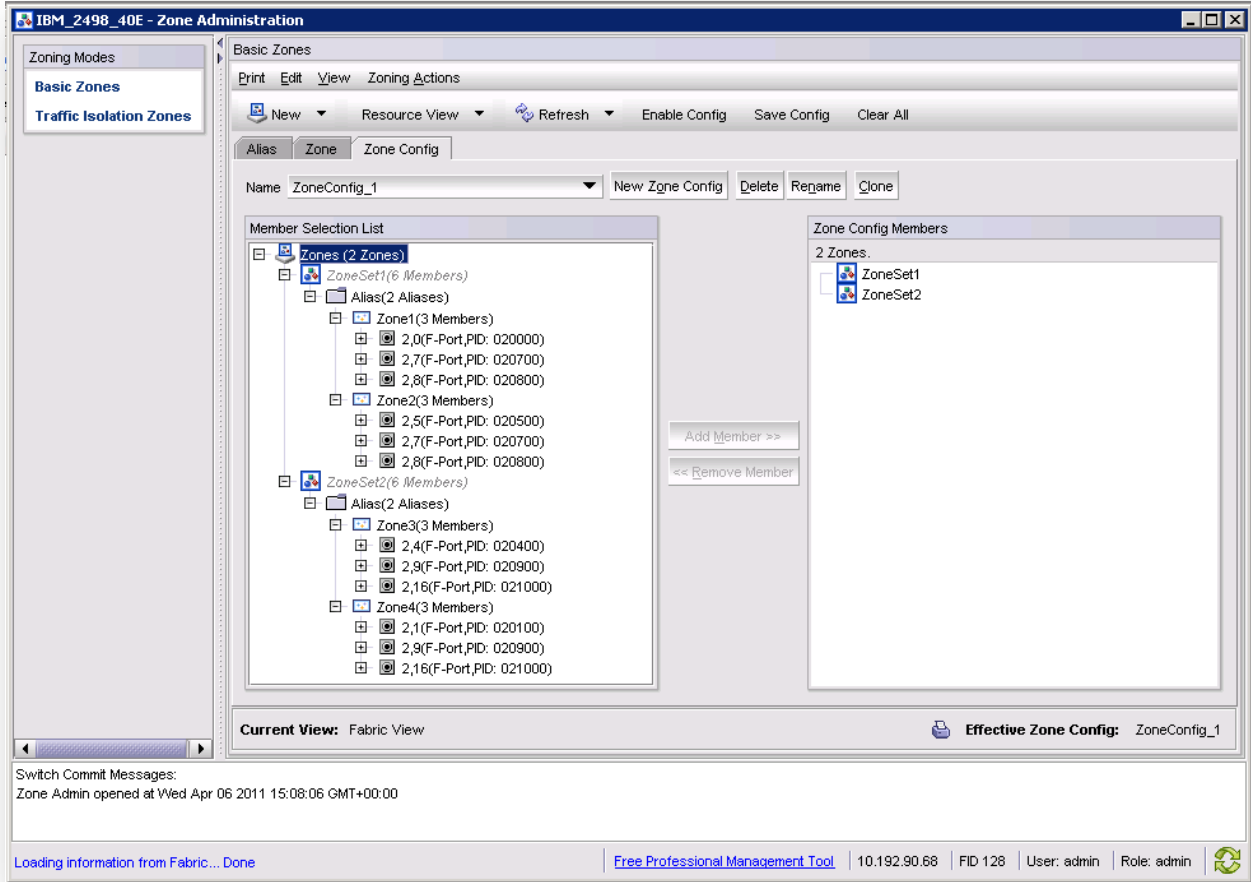


Figure 8. Port zoning in the POC environment, providing four paths to the LUN



The POC environment utilizes two 8Gb/s SAN switches and a total of four zones, as shown in Table 4.

Table 4. Zone configuration

HBA	Storage controller	Alias	Zone	Zone configuration	Switch
Port 0	Port A1	Zone 1	ZoneSet 1	Zone_Config_1	1
	Port B2	Zone 2	ZoneSet 1	Zone_Config_1	2
Port 1	Port A2	Zone 3	ZoneSet 2	Zone_Config_1	2
	Port B1	Zone 4	ZoneSet 2	Zone_Config_1	1

This zone configuration gives each ESX host four paths to a single LUN. The zones were added to an alias zone configuration, which was then activated.

At this stage of the deployment, no LUNs have been created; thus, LUNs cannot yet be bound to WWN ports on the 8GB/s Fibre Channel expansion card.

The next section describes how to configure the IBM System Storage DS3524 Disk System that is used in the POC environment.

Configuring the storage array

Emulex recommends using LSI SANtricity ES Storage Manager to help you configure the storage array so that you can present volumes to ESXi host(s) for use as data stores or Raw Device Mappings (RDM) volumes.

The process for configuring the array includes the following stages:

- Creating a volume group
- Creating a volume
- Mapping the volume to ESXi hosts
- Recording identifiers for the switch, storage array and storage controllers



Note

Mapping several volumes to several ESXi hosts can be performed in a single pass with careful planning.

The process is outlined below.

Creating a volume group

1. Log on to SANtricity ES Storage Manager.
2. Select the **Devices** tab; double-click the array to be managed.
3. Log in with the appropriate credentials.
4. Select the **Logical** tab.
5. Right-click **Total Unconfigured Capacity**, then select **Create Volume Group**.
6. Select **Create** to start the wizard that guides you through the process of creating a volume group.
7. Select the **Group Name** and use the **Automatic** option for drive selection choices.

Note

Advanced users can use the **Manual** option.

8. Select the **RAID** level then select the capacity required for your volume group; select **Finish**.
The wizard now creates the volume group.

Note

Each volume in the volume group uses the same RAID properties.

Consider using simple but unique names for volume groups.

Creating a volume

1. Select **Next** to invoke the volume wizard.
2. Enter the volume capacity and volume name; select **Use recommended settings**.



Note

If the volume is being shared between multiple ESXi hosts, use a naming convention that relates to the particular VMs or to the name of the file system (such as **Marketing VMs**).

Advanced users may wish to configure volume parameters such as I/O characteristics and controller ownership. Less advanced users can use the default settings since the storage array would then assign controller ownership so as to achieve load balancing.

3. Configure the host mode setting.
 4. Use default mapping and then select **VMware** from the drop-down menu.
-

IMPORTANT

Ensure that you specify the appropriate host type.

IMPORTANT

ESXi has a native multipath driver inbox. Thus, when you are using DS3524 storage, you should disable its auto volume transfer/auto disk transfer (AVT/ADT) feature, which is not supported by ESXi.

5. If you wish to create more volumes, select **Yes**. Otherwise select **No** so that you can map the volume to the ESXi host(s).
6. If the hosts have not already been added, click the **Mappings** tab and then select **Undefined Mappings**. Right-click the volume you have just created and then select **Define Additional Mapping**.
7. Enter a host or host group name (such as **IBMX3850X5ZEUS**); select the LUN ID.
8. The LUN ID must be the same for all hosts accessing the shared LUN
9. Select **Add** and then close the window.
10. Within ES Storage Manager, select the interface required for the host (**Fibre Channel**); select **Add by selecting a known unassociated host port identifier**.
11. Select the host port identifier via the drop-down arrow, which associates the host with the host port identifier.
12. Give the host an alias (such as **3850X5ZEUSPort1**); select **Add** and then **Next**.
13. Select the host type (**VMWARE**).



14. Answer **Yes**, indicating that the host will share access to the same volumes with other hosts.
15. After indicating that the volume will be shared, you should now specify the remaining hosts that will share the volume. Enter a name for each host group or select an existing group.
16. Select **Next** and then **Finish**.

Mapping the volume

1. Select the volume name from the undefined mappings, then right-click the volume and select **Define additional mappings**.
2. Select the host or host group you wish to map to the volume.
3. Verify the volume by selecting the host from the **Host Group** or **Default Group** in the topology table of contents.
4. Enter a LUN; select the volume name and add LUN to Host or Host Group.
5. Select **Close**.

This completes the storage configuration process.

Note that, in many implementations, the DS3524 storage array may be accessed by multiple ESXi hosts; indeed, it may be shared heterogeneously with systems running different operating systems. Storage zoning can be used to protect individual volumes, helping to avoid data corruption or the accidental deletion of file systems.

Next, you should configure the 8GB/s Fibre Channel expansion card to support BfS.

Configuring BfS

Enterprise server manufacturers such as IBM continue to offer local disk storage; however, with the growth of virtualization and the increased use of BfS⁸, server configurations are evolving. For example, IBM now offers a diskless server, which would allow you to deploy ESXi via a USB flash drive.

BfS capability is often configured in an ESXi environment, where its benefits include:

- Enhanced manageability
- Faster deployment
- Easier backup⁹
- Enhanced disaster recovery capabilities

⁸ Whereby the storage array owns the operating system LUN

⁹ Since the array owns the LUN, array-based copies can be made without server intervention.



The BfS configuration process includes the following stages:

- Loading the latest boot code to the 8GB/s Fibre Channel expansion card
- Configuring the ESX host
- Specifying the desired boot volume
- Placing the 8GB/s Fibre Channel expansion card first in the boot order

This vendor-specific procedure describes how to configure an Emulex 8Gb/s Fibre Channel expansion card for BfS.

Configuring the 8Gb/s Fibre Channel expansion card for BfS

IBM and Emulex have jointly created a simple, unique method for configuring Fibre Channel expansion cards through the UEFI.

Note

For more information on configuring BfS in an ESXi environment, refer to VMware's [Fibre Channel SAN configuration guide](#).

After configuring the storage array and provisioning volumes, you can perform the following procedure to enable the 8Gb/s Fibre Channel expansion card for BfS:

Note

This configuration is performed within the blade server's UEFI. In the event of a legacy (non-BfS) boot, all settings are shared with legacy mode; thus, there is no longer a need to configure legacy BIOS settings.

1. Ensure that you have presented the boot (target) LUN to the hosts. Also ensure no other LUNs are visible at the time of the ESXi 4.1 install to make the process easier.
2. Ensure that you know the WWPNs for the 8GB/s Fibre Channel expansion card you plan to use.
3. Ensure that the 8GB/s Fibre Channel expansion card is running the latest bootBIOS code¹⁰, which can be downloaded from the [Emulex website](#).
The bootBIOS code used in the POC environment is version 2.12a6.
4. After updating the boot code, reboot the host.
5. When prompted, select **<F1>Setup during boot**.
6. Select **System Settings** and click **Enter**.

¹⁰ The boot code works with the server's BIOS to allow the system to boot from a connected Fibre Channel drive.



7. Select **Emulex Configuration Utility**, as shown in Figure 9.

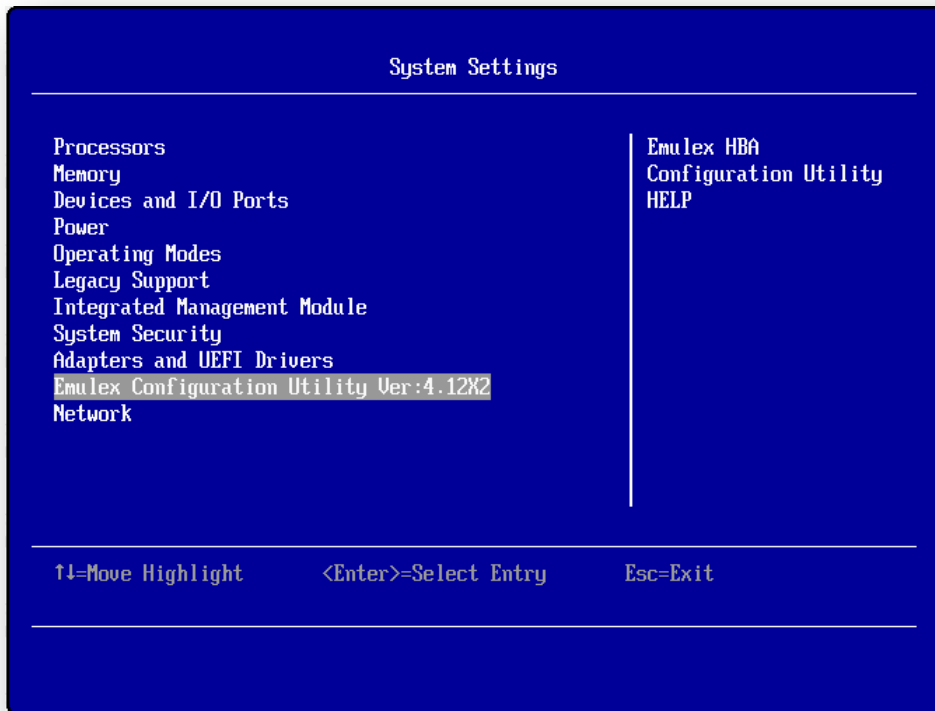


Figure 9. Selecting the Emulex Configuration Utility from UEFI

8. Select the Emulex adapter (**Emulex HBA**) from the list provided and click **Enter**.



9. Select **Set Boot from SAN**, as shown in Figure 10. To select **Enable**, click **Enter**.

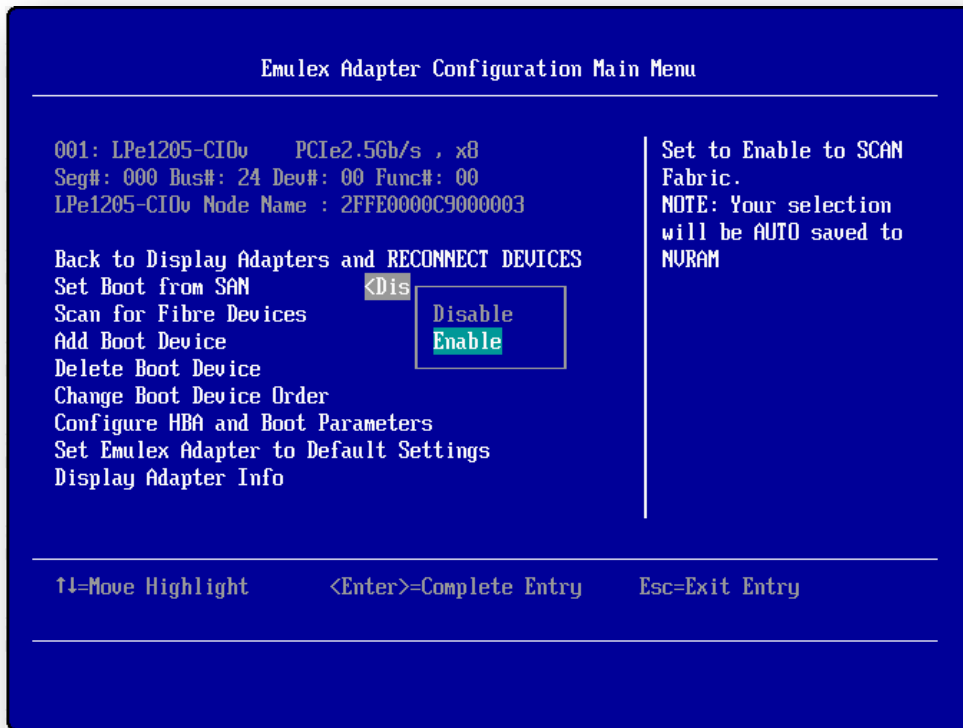


Figure 10. Selecting and enabling BfS

10. Select **Scan for Fibre Devices** to ensure you can see the LUNs that were previously provisioned by your SAN administrator.
11. Select **Add Boot Device**.



12. Select the desired target, as shown in Figure 11.

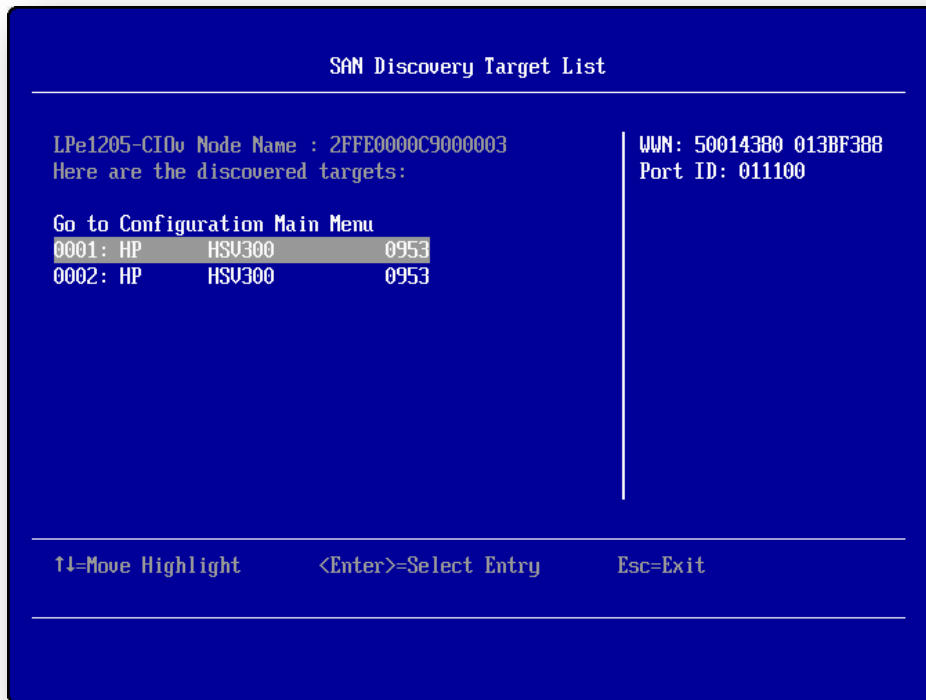


Figure 11. Selecting the discovered target storage processors



13. Select the LUN you wish to use as the boot LUN, as shown in Figure 12.

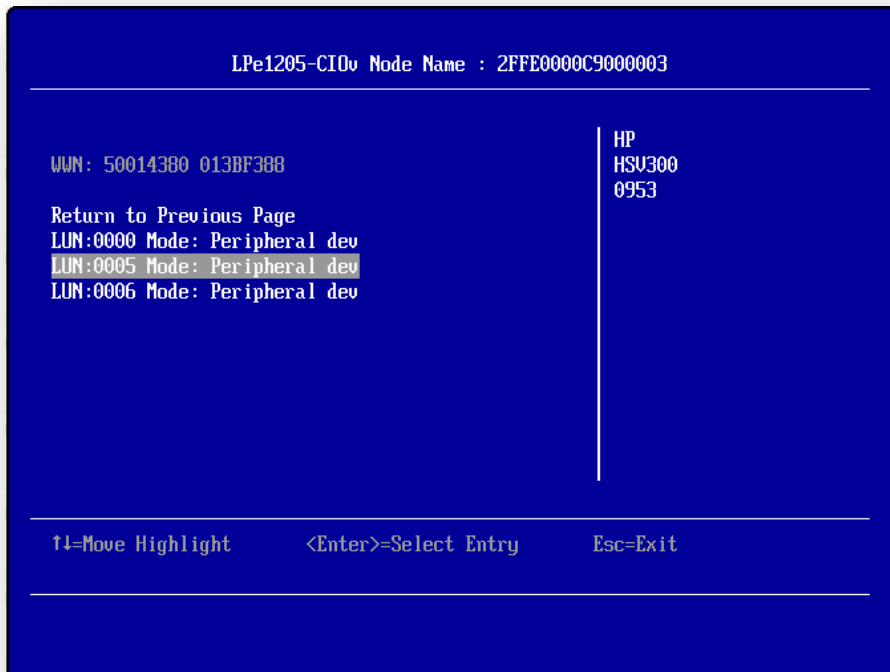


Figure 12. Selecting the boot LUN

14. Select **Commit Changes**.

15. Click the **Esc** key until you arrive at the **System Configuration and Boot Management** screen.

16. Select **Save Settings** and **Exit Setup** to reboot the host.

You can now install ESXi 4.1 software on the target LUN.

Configuring RAID on local storage

If you are installing ESXi 4.1 locally on a solid-state drive (SSD) or hard disk drive rather than using BfS, you should set a RAID1 configuration on the ESXi host, allowing you to maintain a backup of the operating system on a second drive in case of failure. In most cases, two identical drives are sufficient for an ESXi 4.1 installation.

When the ESXi host starts or reboots, press **CTRL A** to enter the configuration menu for the server's RAID controller. It is not important to configure for performance workloads.



Note

When testing the capability to test to boot from local storage, Emulex disconnected the fabric to prevent the host from booting from SAN or attempting to scan the entire array.

You are now ready to install ESX 4.1 on the BladeCenter HX5 blade server.

Deploying ESXi 4.1

This section describes how to install ESXi 4.1 on the BladeCenter HX5 blade server, providing information on the following topics:

- **Prerequisites:** Downloading the latest drivers
- **Installation:** Installing ESXi 4.1 from a image
- **Pre- or post-installation considerations:** Installing or updating ESXi drivers for the Emulex adapters; updating the VMkernel to support DS3500 arrays

Prerequisites

Downloading the latest drivers

Emulex recommends downloading any necessary drivers before installing ESXi 4.1.

The POC environment uses the following Emulex drivers:

- **8Gb/s Fibre Channel expansion card driver:** The Emulex lpfc820 CNA Driver is an in-box driver that is installed during the installation of ESXi 4.1. If you do not wish to rely on the in-box driver, you should download the latest driver¹¹ from the [Emulex website](#), for installation following the deployment of ESXi 4.1.
- **VFA NIC driver:** The NIC driver for the VFA is not provided in-box and should be downloaded from the [VMware website](#) to be installed following the deployment of ESXi 4.1.

In addition, you should update the Emulex Common Information Model (CIM) providers. Download the update package from the [Emulex website](#) (**Management** tab), for installation following the deployment of ESXi 4.1.

¹¹ VMware ESX/ESXi 4.1 Driver CD for Emulex FC/FCoE CNA Driver



Note

Depending on the number of drivers you need to install, it may be beneficial to script the install process.

You are now ready to install ESXi 4.1.

Using a DVD image to deploy ESXi 4.1

ESXi 4.1 can be deployed in several ways. For example, if your servers are managed through IBM Director or IBM Tivoli, you can use these applications for the rapid deployment of ESXi 4.1 on to a particular server. In the POC environment, a DVD ISO file was used.

Note

A fully-featured ESXi 4.1 image can be downloaded free-of-charge from VMware. This image, which is valid for 60 days, may be useful in a POC deployment.

The procedure for installing ESXi 4.1 from an image includes the following steps:

1. Either download the image from the [VMware website](#) or use the [VMware vSphere Hypervisor 4.1 U1 with IBM Customization](#).
2. Store the image on the network where it can be accessed by the BladeCenter HX5 blade server's AMM.

If necessary, configure the AMM.

3. Using a web browser, enter the IP address of the AMM port.
 4. Log in to the AMM.
-

Note

If the AMM NIC port was not configured prior to this procedure, you may need to contact your server administrator to obtain log-in information.

5. Select the remote management keyboard, video and mouse (KVM) feature.
6. Open the DVD image window and select the ISO image.
7. Power on the server, either manually or from within the KVM.

More information

If you are unfamiliar with the installation process, you should first review the instructions provided by [IBM](#). Additional instructions are provided by [VMware](#).



Post-installation considerations

Installing the VMware drivers and Emulex OneCommand Manager

Before deploying any VMs, you should ensure that the latest versions of the ESXi 4.1 drivers for Fibre Channel and the VFA NIC have been installed on the host. You may also wish to install Emulex OneCommand Manager to simplify the management of Emulex adapters in an ESXi environment.

ESXi is bundled with a range of in-box drivers that are automatically installed when the appropriate hardware is detected. Thus, for example, a BladeCenter HX5 blade server equipped with an 8Gb/s Fibre Channel expansion card automatically loads the in-box Fibre Channel driver. However, the NIC driver for the VFA is not provided in-box and must be installed manually.

Note

To verify that you are running the latest, approved Fibre Channel driver, check the [VMware Compatibility Guide](#), which lists the versions of drivers that have been tested with ESXi 4.1 (or later).

Installing or updating the 8Gb/s Fibre Channel expansion card driver

For detailed information on updating or installing the Fibre Channel driver, refer to [Appendix C – Installing drivers and CIM providers](#).

Installing the VFA NIC driver

After the ESXi 4.1 deployment is complete, locate the NIC driver that you downloaded prior to the ESXi install. Detailed information on installing the NIC driver is available on the [Emulex website](#).

Installing OneCommand Manager

To centralize and enhance the management of Emulex devices in an ESXi environment, you can also install OneCommand Manager, which is part of the VMware Software Kit that is available free-of-charge on the [Emulex website](#) (under the **VFA Software Kits** tab).

For detailed information on installing OneCommand Manager, refer to [Using the OneCommand Manager plug-in for vCenter](#).



Updating the VMkernel to support DS3500 storage arrays

Before deploying VMs, you should consider updating the VMkernel, using vCenter Server Advanced Settings to better support your storage array. For example, you may wish to disable the **Use.DeviceReset** option.

Table 5 shows advanced VMkernel options that are specific to DS3500 arrays. For more information, refer to the appropriate draft [IBM Redbooks](#) publication.

Table 5. Advanced VMkernel options associated with DS3500 storage

Action	Advanced option
Disable the Use.DeviceReset option for IBM midrange storage subsystem logical drives.	Disk.UseDeviceReset=0
Enable the Use.LunReset option.	Disk.UseLunReset=1
If the logical drives are not being used for either RDM or Microsoft Cluster Server (MSCS), disable the ResetOnFailover option.	Disk.ResetOnFailover=0
If you are using RDM or MSCS across multiple ESX hosts, enable the ResetOnFailover option.	Disk.ResetOnFailover=1 Enable Disk.RetryUnitAttention Disk.RetryUnitAttention=1
Enable logging on the ESX host.	Scsi.LogMultiPath=1 Scsi.PrintCmdErrors=1
If you are working with FlashCopy or Remote Mirror logical drives, enable LVM.EnableResignature .	LVM.EnableResignature = 1

Managing the solution

With the growing acceptance of virtualization, the IT management burden is increasing, causing vendors to create specific plug-ins to simplify the management of an ESXi environment. For example, the Emulex OneCommand Manager plug-in for vCenter enables the centralized management of Emulex adapters via the vCenter console, providing a user experience similar to that delivered by a standalone OneCommand Manager implementation. Meanwhile, the IBM System Storage plug-in for vCenter allows you to monitor and manage DS5000, DS4000 and DS3000 storage arrays that have been deployed in conjunction with ESXi.

In addition, since developing one of the first tools for the management of x86 servers, IBM has continued to expand their management capabilities. For example, IBM Systems Director can be used to manage any IBM system, whether x86, UNIX or mainframe.



Note

For more information on the configuration and management of BladeCenter HX5 blade servers, refer to [IBM Redbooks](#).

This section provides more information on management plug-ins.

Using the OneCommand Manager plug-in for vCenter

The OneCommand Manager (OCM) plug-in enables single-pane-of-glass management for Emulex adapters via vCenter Server. This plug-in builds on existing Emulex drivers and providers for ESX and ESXi hosts; no additional agents are required.

Capabilities

Delivered via the vCenter Server console, the capabilities of the OCM plug-in include:

- Monitoring Emulex adapters deployed within a single ESXi host or an ESXi cluster
- Managing parameter settings
- Simultaneously updating the firmware of multiple Emulex adapters

A command-line interface (CLI) is also provided.

Either the OCM plug-in or the CLI can provide a single-pane view (as shown in Figure 13) to help you better manage and monitor your SAN infrastructure.

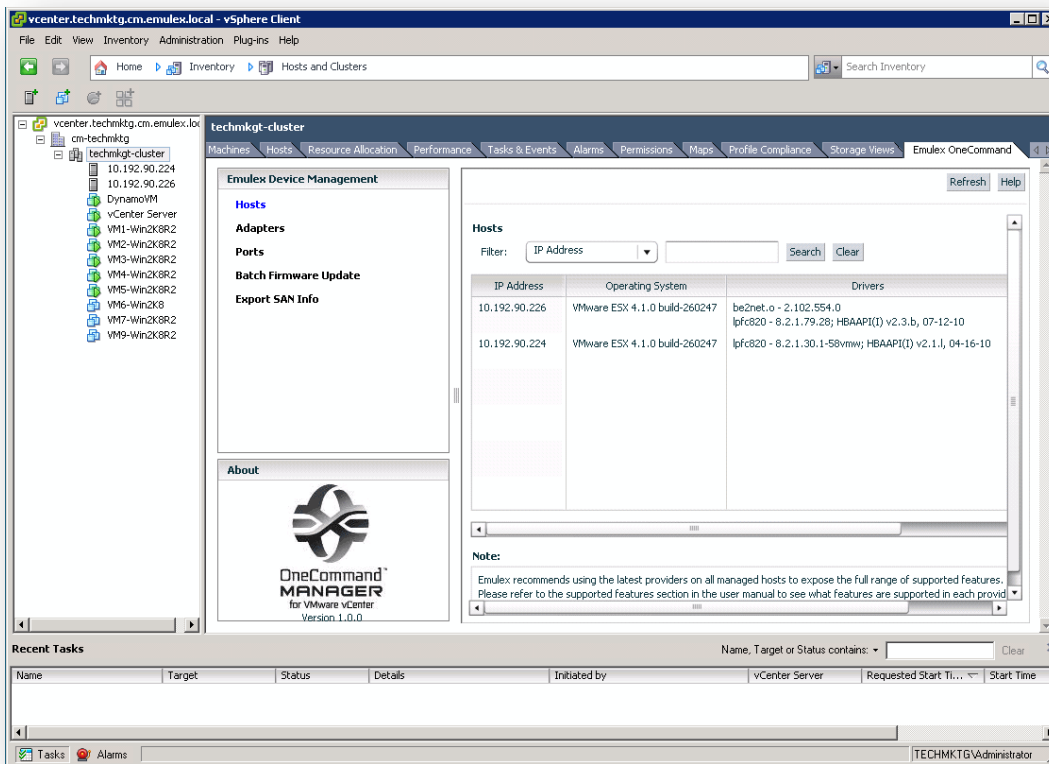


Figure 13. View of **technmkt-cluster** from the OCM Plug-in

OCM can be used to perform value-add activities such as changing the LUN queue depth to enhance storage performance, as described below.

Installing the OCM plug-in

The process for installing the OCM plug-in is as follows:

1. Download the plug-in from the [Emulex website](#) (**Management** tab), using a Windows management workstation.
2. Run the executable file.
3. Select the **OCM for VMware vCenter Registration** icon that has been placed on the desktop.



4. Register the OCM plug-in by providing the vCenter Server name or IP address, as shown in Figure 14.

The screenshot shows the registration window for OneCommand Manager for VMware vCenter. The window has a dark green header with the OneCommand logo and the text 'MANAGER for VMware vCenter'. Below the header are four input fields: 'vCenter Server Name', 'vCenter Server HTTPS Port' (with '443' entered and a checkbox), 'Username', and 'Password'. At the bottom are three buttons: 'Reset', 'Register', and 'Unregister'. The footer features the 'EMULEX' logo.

Figure 14. OCM for VMware vCenter register/ unregister window

5. Login to vCenter Server via vSphere client to ensure that the OCM plug-in is enabled.

If you see an **Emulex OneCommand** tab in vCenter Server, this indicates the plug-in has been installed successfully.

Note

For the plug-in to function correctly, you must have administrator's privileges. The **Apply** button (see Figure 13) does not appear unless you have such privileges.



Changing the LUN queue depth

In effect, the LUN queue depth determines how many commands the HBA can accept on a per-LUN basis. At times, an ESX host may detect a queue-full condition that causes SCSI commands to be aborted; thus, increasing the queue depth may help you enhance I/O performance.

Note

The LUN queue is a single component of an I/O path that has multiple layers, each of which can impact I/O performance.

Example

Provided for demonstration purposes, this example describes how to use the OCM plug-in for vCenter to change the depth of the LUN queue. You no longer need to run a separate application on a management server in order to carry out this and other administrative tasks.

Perform the following steps:

1. Use vSphere Client to log into vCenter Server.
2. Select the ESXi host on which you wish to change the LUN queue depth.
3. Select the **Emulex OneCommand** tab.
4. Select the particular adapter from the **Emulex Device Management** pane.
5. Select the port and then select the **Driver Parameters** tab.



6. Scroll down to the **lun-queue-depth** parameter; make the desired change, as shown in Figure 15.

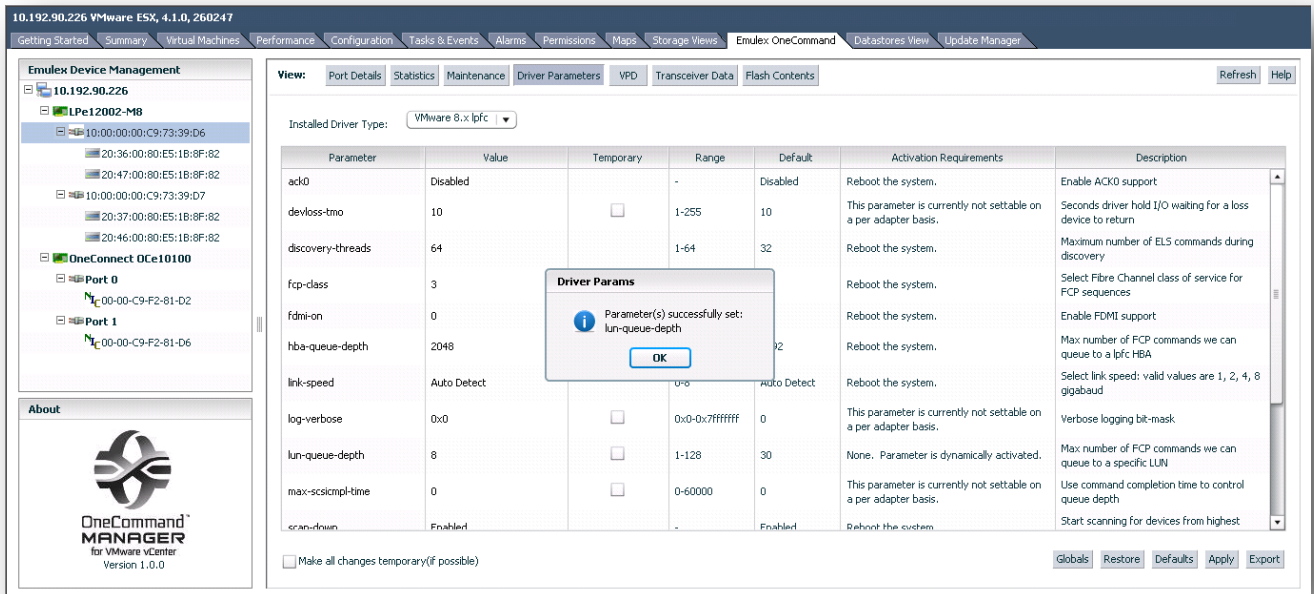


Figure 15. Successfully changing the LUN queue depth parameter from its default value of 30 to 8

7. Select **Apply**, which appears after the parameter has been entered.

Effect of multiple VMs on queue depth

According to [VMware](#), if only one VM is active on a LUN, you only need to set the maximum queue depth (**lun-queue-depth**). If, however, multiple VMs are active, the value of the **Disk.SchedNumReqOutstanding** VMkernel setting must also be considered.

With multiple queues, the maximum queue depth is equal to the lower of the following settings:

- Adapter queue depth
- **Disk.SchedNumReqOutstanding**

Note

VMware recommends setting **Disk.SchedNumReqOutstanding** and the adapter queue depth to the same value.

Refer to the [VMware Knowledgebase](#) for more information on changing the queue depth.



Using the IBM System Storage Plug-in for vCenter

IBM provides a free management plug-in – the [System Storage Plug-in for vCenter](#) – for customers running DS3000, DS4000 and DS5000 storage arrays in an ESXi 4.x environment. This tool gives you the ability to monitor and manage DS family arrays from within a vCenter Server window. You can provision the entire storage subsystem from a single pane without needing to switch between multiple interfaces.

The System Storage Plug-in not only enables the array to be viewed from within vCenter Server, it also provides management capabilities allowing you to add a particular host or host group to a storage subsystem.

The System Storage Plug-in provides support for the following:

- vCenter Server 4.0 and 4.1
- Fibre Channel connectivity
- Windows Server 2003, 2003 R2, 2008 and 2008 R2 operating systems



Figure 16 shows a view of a single storage array within vCenter Server.

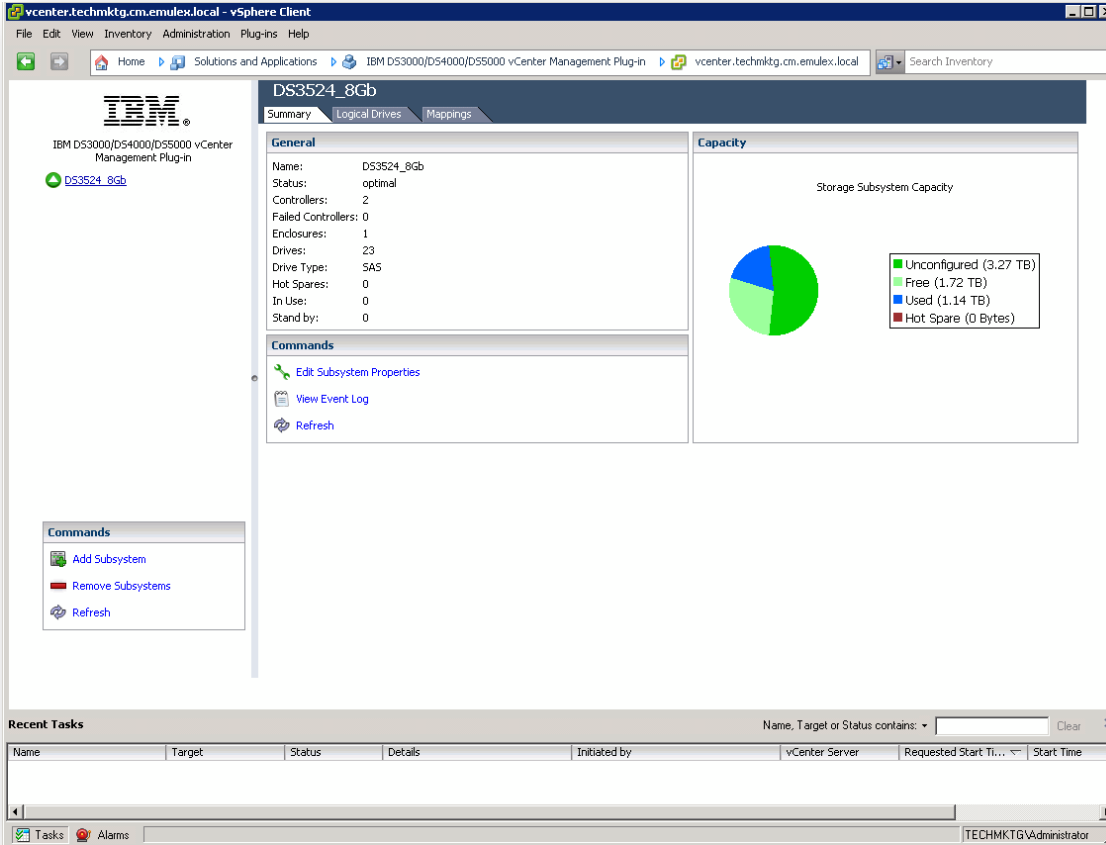


Figure 16. View of a DS3524 array from vCenter Server

For information on installing the plug-in, refer to [Appendix B – Installing the IBM System Storage plug-in.](#)

The next section outlines the procedure for using the System Storage Plug-in to configure an ESXi host to the storage subsystem.



Configuring an ESXi host to a storage subsystem

1. Right click on the host and select **Configure ESX Host to Storage Subsystem** as shown in Figure 17.

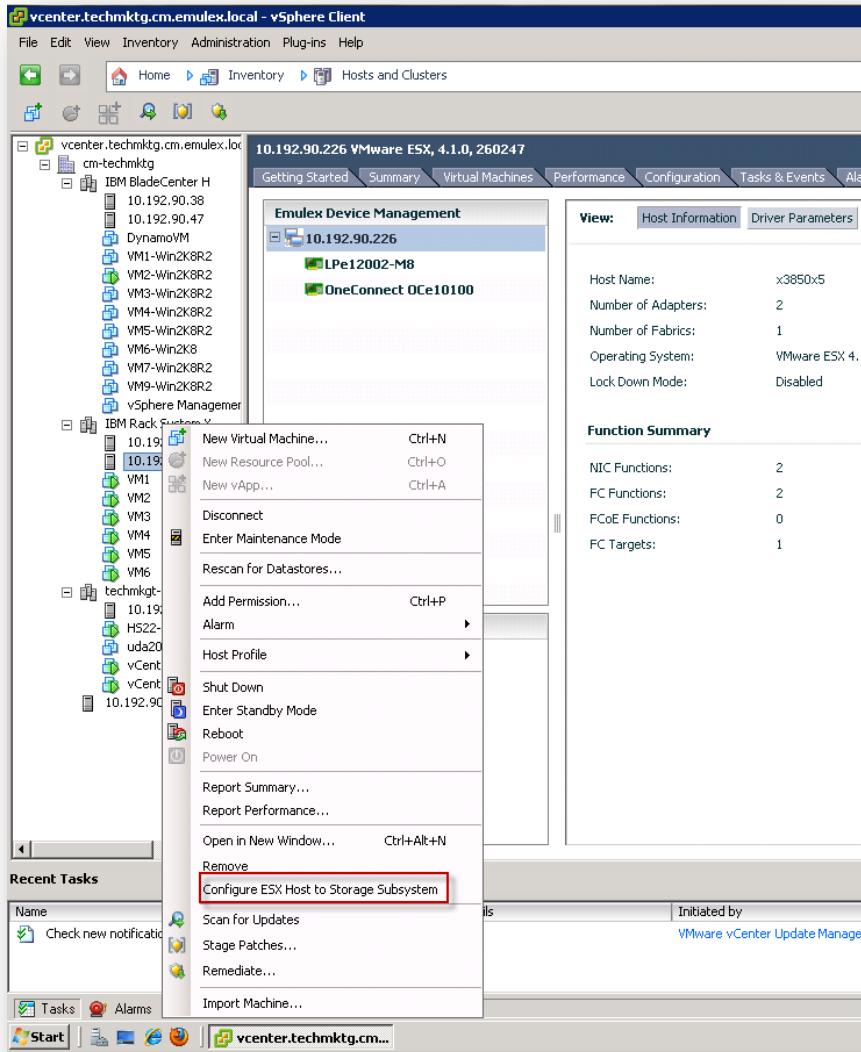


Figure 17. Selecting the **Configure ESX Host to Storage Subsystem** option



2. Select the desired storage array. You now have the option to add a host or host group or have IBM System Storage Plug-in make a suggestion for you. In the example shown in Figure 18, ESX host **10.192.90.226** was added to a host group called **VMwareCLSTR**.

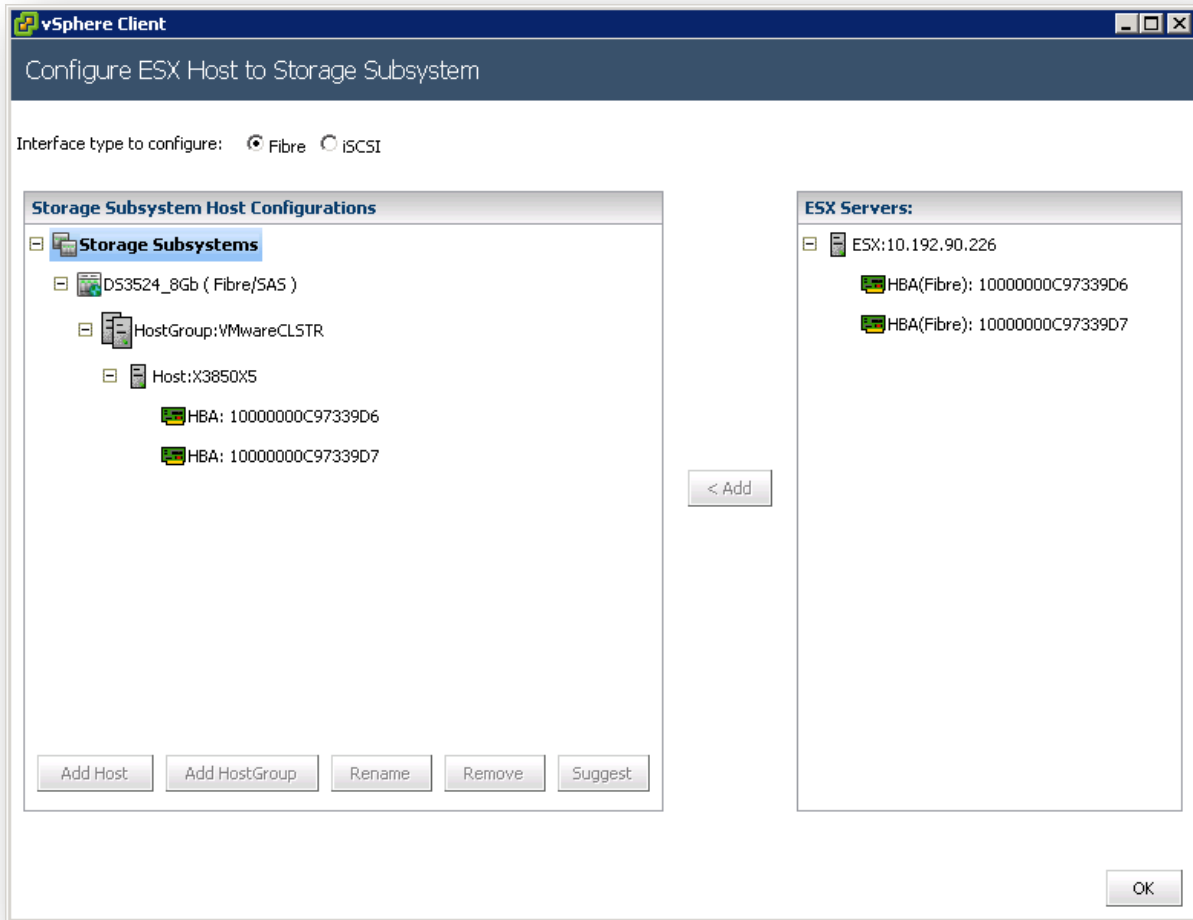


Figure 18. Ports from ESX host **10.192.90.226** added to the host group called **VMwareCLSTR**.



Using the System Storage Plug-in to view datastores

Once the IBM System Storage Plug-in has been enabled, the **Datastores View** tab can be used to obtain information (such as size, type of connectivity, status and path count) about datastores provisioned by the array. The following views are available:

- **Datastore View** (as shown in Figure 19)
- **Logical Drive View**

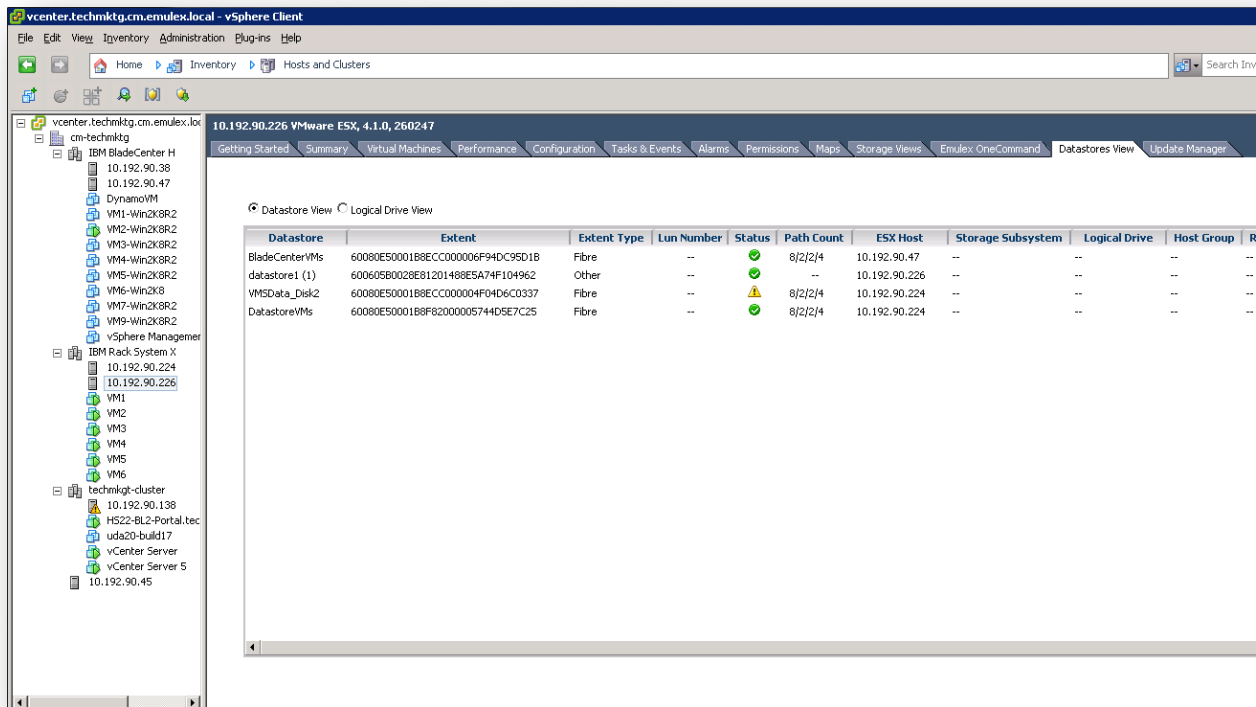


Figure 19. Datastores View tab in the System Storage Plug-in



Advanced topics

The remaining sections of this white paper outline options for enhancing the availability or performance of the POC environment.

Information is provided on the following topics:

- **Configuring multipathing:** Configuring multiple paths between an ESXi host and storage to eliminate a single point of failure
- **Enabling the VMware Storage I/O Control feature:** Minimizing the latency associated with block-based storage, thus enhancing storage performance

Configuring multipathing

To enhance fault tolerance, the POC environment featured multipathing between the BladeCenter HX5 blade server and the DS3524 storage array.

ESXi 4.1 is Asymmetric Logical Unit Access (ALUA)-compliant, which simplifies the process for setting a multipath configuration. When an ALUA array is detected by ESXi, it automatically selects the optimal path to the target LUN. Furthermore, VMware provides a multipathing driver in-box. Since this driver is supported by the DS3524 storage array, there is no need to install additional IBM software.

This section provides information on the following topics:

- Setting up the ESXi storage stack for multipathing
- Avoiding LUN thrashing
- Validating the multipathing implementation in the POC environment



Setting up the storage stack for multipathing

In ESXi, the VMkernel is responsible for managing interactions with storage via the Pluggable Storage Architecture (PSA), a modular system (shown in Figure 20) that also provides multipathing support for arrays specified in the [VMware Hardware Compatibility Guide](#).

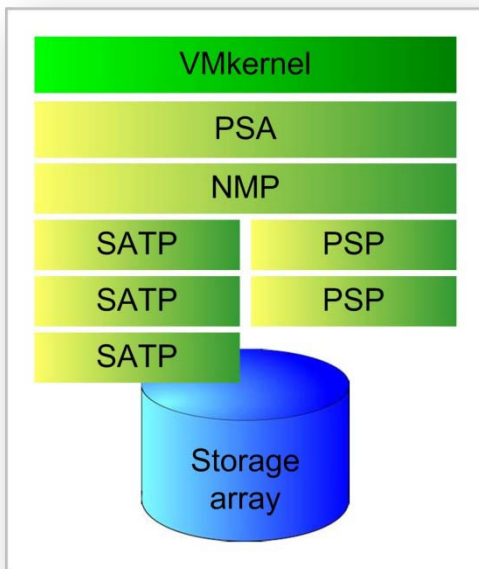


Figure 20. Simplified view of the ESXi storage stack

The stack includes the following components:

- **PSA:** The PSA is responsible for managing storage and coordinating the simultaneous operation of multiple plug-ins
- **NMP:** The generic Native Multipathing (NMP) plug-in manages lower-level plug-ins and is responsible for operations that are not array-specific, such as monitoring, path probing and multipath configuration. Alternatively, you can deploy a third-party multipathing plug-in (MMP) such as EMC Power Path.
- **SATP:** There is a device-specific Storage Array Type Plug-in (SATP) for each supported array. The SATP makes the array-type known to the VMkernel and is able to communicate requirements (such as device discovery and failover) to the array; it also understands error codes generated by the array.

SATPs are global; thus, a particular array can only be associated with a single SATP.

The default SATP is **VMW_SATP_LSI**, which has been designed to accommodate LSI and similar arrays.



- **PSP:** A Path Selection Plug-in (PSP) is responsible for selecting a physical path for each I/O request. PSPs work in conjunction with SATPs.

The tested environment uses the Most Recently Used (MRU) PSP – **VMW_PSP_MRU** – which always tries to select the path that the ESX host had used most recently. If this path is unavailable, the PSP switches to an alternate path. Other policy options are Fixed and Round Robin (RR).

PSPs are used on a per-LUN basis; thus, for example, one LUN can be set to use an MRU policy while another LUN on the same array can be set to use a RR policy.

Identifying the PSP setting for a particular LUN

If you wish to determine the current PSP setting(s), use vSphere Client to log into vCenter; alternatively, you could use vSphere Management Assistant (vMA).

Refer to the [VMware Hardware Compatibility Guide](#) for information on multipathing configurations for your array that are supported by ESXi.

Using vCenter Server

Perform the following steps:

1. Log in via vSphere client to vCenter Server.
2. Select the desired ESXi host.
3. Select the **Configuration** tab and then **Storage**→**Datastore**.
4. Select **Properties**→**Manage Paths**.



Figure 21 presents a typical view of the Policy Path Selection.

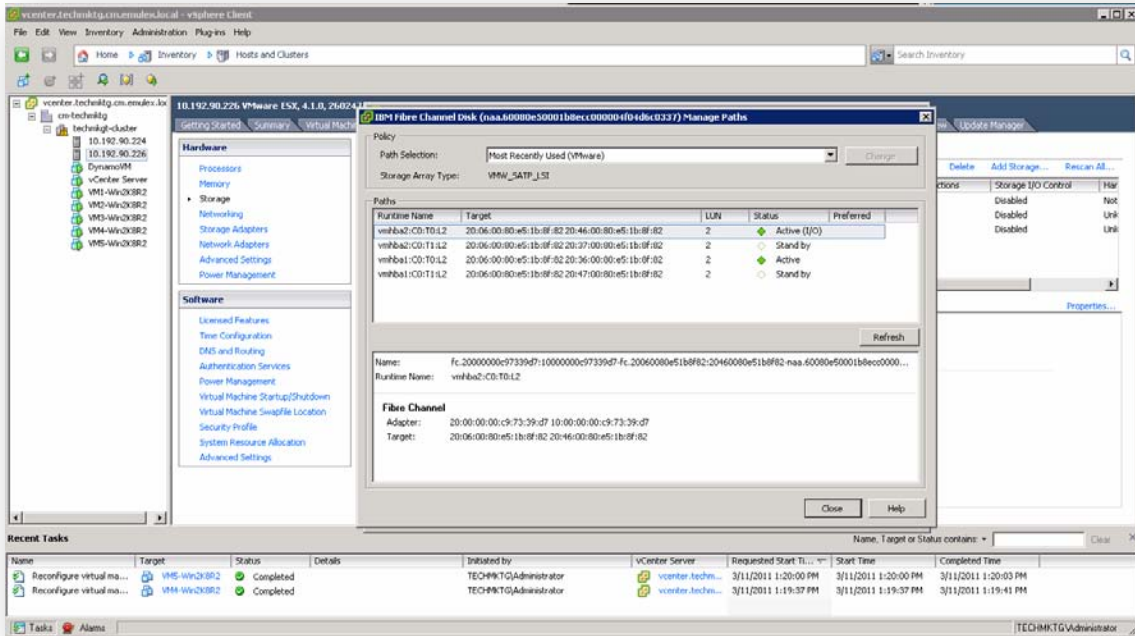


Figure 21. ESX multipathing settings for the DS3524 array set to Most Recently Used

In this example, the storage array type is an active/passive **VMW_SATP_LSI**. As shown in the **Path Selection** box, the current PSP setting is **MRU**.

Using vMA

Alternatively, you can use the vSphere CLI, part of vMA, to show current PSP settings. Use the following command:

```
esxcli nmp satp psp -l
```

To view all storage devices controlled by the NMP, use the following command:

```
#esxcli nmp device list
```



Identifying supported configurations

Figure 22 presents supported multipathing configurations for a DS3500 storage array, obtained using the [VMware Hardware Compatibility Guide](#).

Release	Mode	Path Policy	Firmware	Device Driver(s)	Configuration
ESX 4.1	VMW_SATP_LSI	VMW_PSP_MRU	Refer to vendors Interoperability Matrix	N/A	8G FC Switched
ESX 4.1	VMW_SATP_LSI	VMW_PSP_MRU	Refer to vendors Interoperability Matrix	N/A	FC Switched
ESX 4.0 U1 ¹	VMW_SATP_LSI	VMW_PSP_MRU	Refer to vendors Interoperability Matrix	qla2xxx version 821.k1.43vmw, lpfc820 version 8.2.0.30.52vmw	FC Switched
ESX 4.0 U1	VMW_SATP_LSI	VMW_PSP_MRU	Refer to vendors Interoperability Matrix	N/A	8G FC Switched
ESX 4.0 ¹	VMW_SATP_LSI	VMW_PSP_MRU	Refer to vendors Interoperability Matrix	qla2xxx version 821.k1.38vmw, lpfc820 version 8.2.0.30.49vmw	FC Switched
ESX 3.5 U5 ^{2,3}	N/A	MRU	Refer to vendors Interoperability Matrix	lpfc_740 version elx_7.4.0.40.3, bfa version 1.0.0.02, qla2x00_707 version 7.07.04.01	8G FC Switched
ESX 3.5 U5 ²	N/A	MRU	Refer to vendors Interoperability Matrix Equivalent	lpfc_740 version elx_7.4.0.40.3, bfa version 1.0.0.02, qla2x00_707 version 7.07.04.01	8G FC Switched

Notes:

Storage partners using ESX 4.0 or later may recommend VMW_PSP_RR for path failover policy for certain storage array models. If desired, contact the storage array manufacturer for recommendation and instruction to set VMW_PSP_RR appropriately.

1. Please refer to vendors Interoperability Matrix for minimum ESX Update Release
2. End-to-end connectivity at 8Gbps FC speed is supported with 8G FC arrays only if the product details has prefix of "8G FC" or is footnoted. Otherwise support for 8G FC arrays is limited to up to 4Gbps speed only
3. Patches ESX350-201002403-BG, ESX350-201002408-BG, and ESX350-201002409-BG are required.

Figure 22. ESX multipathing settings for the DS3524 array

Note

Emulex suggests using the settings recommended by VMware.

Avoiding LUN thrashing

LUN thrashing can occur if paths from two different ESXi hosts contend for the same LUN via different storage controllers on the same storage array.

To avoid thrashing, ensure hosts always access LUNs via the owning storage controller.

Note

After a LUN has been created – but before it is mapped to a host – the array can assign an owning controller to the LUN automatically.



Validating the multipathing implementation

In order to validate the multipathing implementation in the POC environment, Emulex created a failure, as described in this section.

Performance degradation

Even with multipathing, you can expect some performance degradation if a failure were to occur on the SAN. This is especially true if the failure is on the primary path, which also happens to be the preferred path to the LUN, such as a controller, an HBA, a cable or a fabric switch.

If you set up an optimal multipathing configuration that includes fabric zones and the appropriate cross-connection of fiber cables between switches and hosts, ESXi hosts should still be able to access the SAN, though in a degraded fashion. At this point, you should use the storage management software supplied with the array (such as IBM DS Storage Manager or SANtricity ES Storage Manager) to troubleshoot and address the problem.

Testing failover

To test multipathing, Emulex simulated an error on the storage array by physically removing the primary controller. This controller owns the LUNs for the following virtual disks:

- Guest OS
- Data disk

Note

The virtual disks are maintained in the same disk group but on different LUNs..

After the controller went offline, there was no feedback from vCenter Server regarding the failure; all alarms remained at their default settings. However, as shown in Figure 23, the System Storage plug-in for vCenter provided a notification indicating that attention was required – but without delivering any additional information.

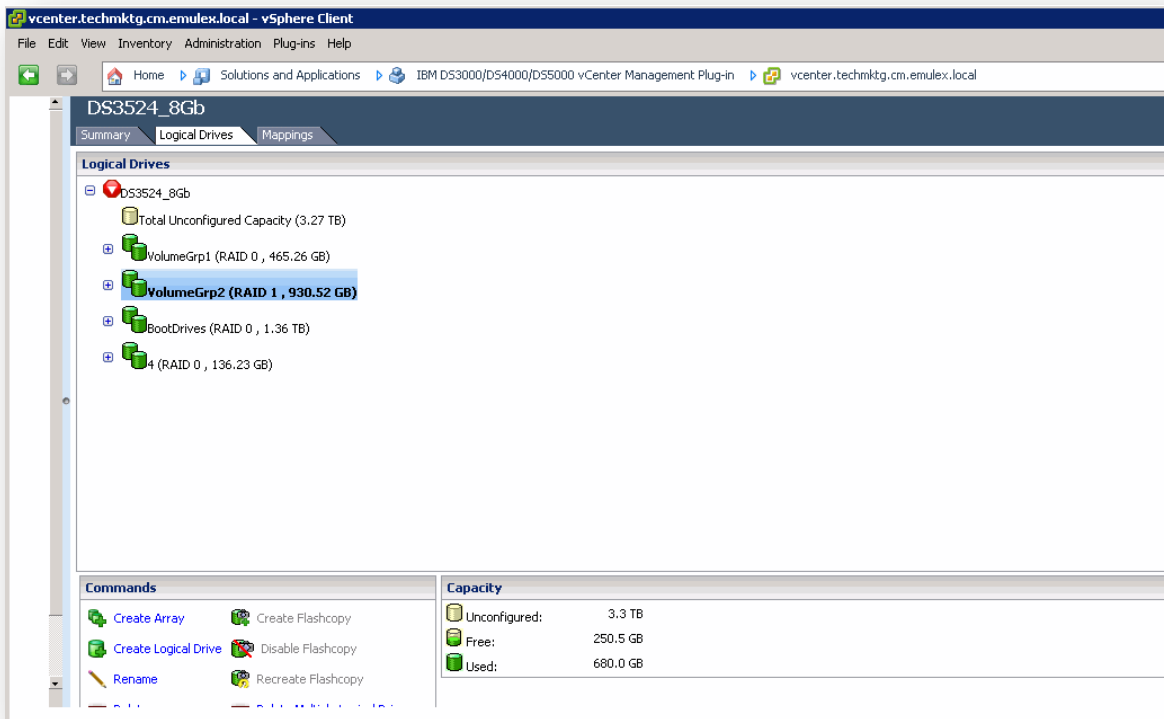


Figure 23. The System Storage plug-in for vCenter provided a notification of the storage failure



To obtain more information about the failure, use out-of-band management software that has been installed on a management server. This example uses SANtricity ES Storage Manager, as shown in Figure 24.

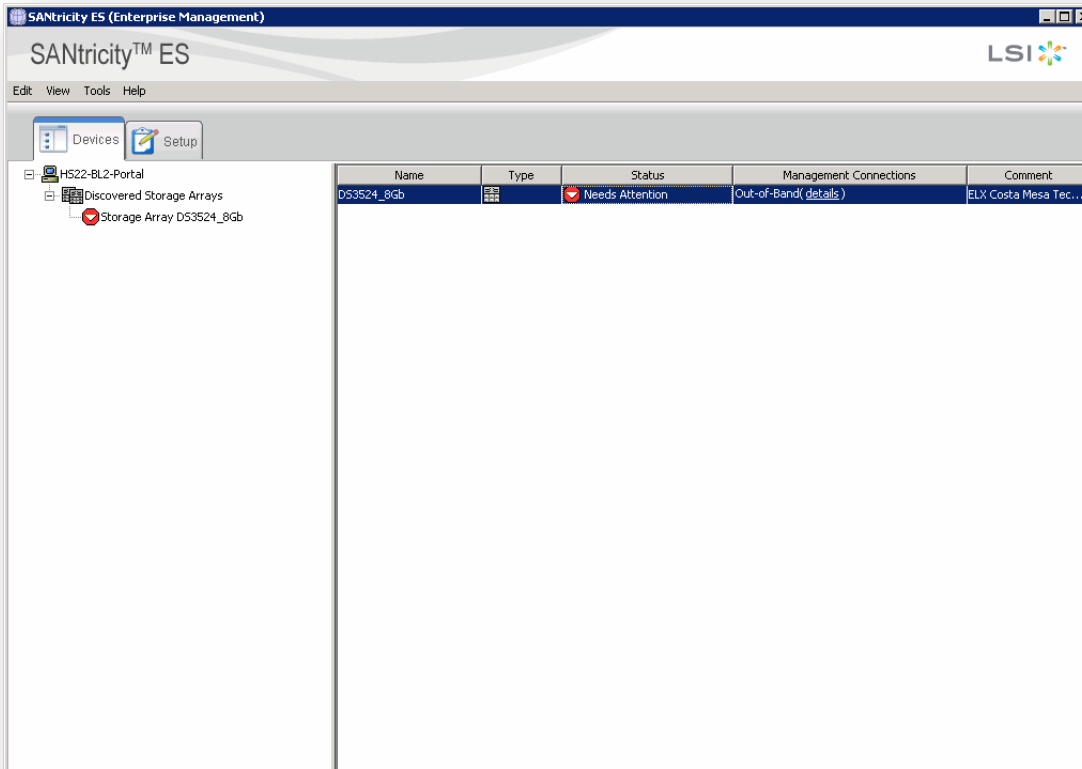


Figure 24. SANtricity ES Storage Manager providing a notification about the storage failure



Double-click on the storage array associated with the notification. On the following screen, select **Needs Attention** (above the menu tabs), as shown in Figure 25.

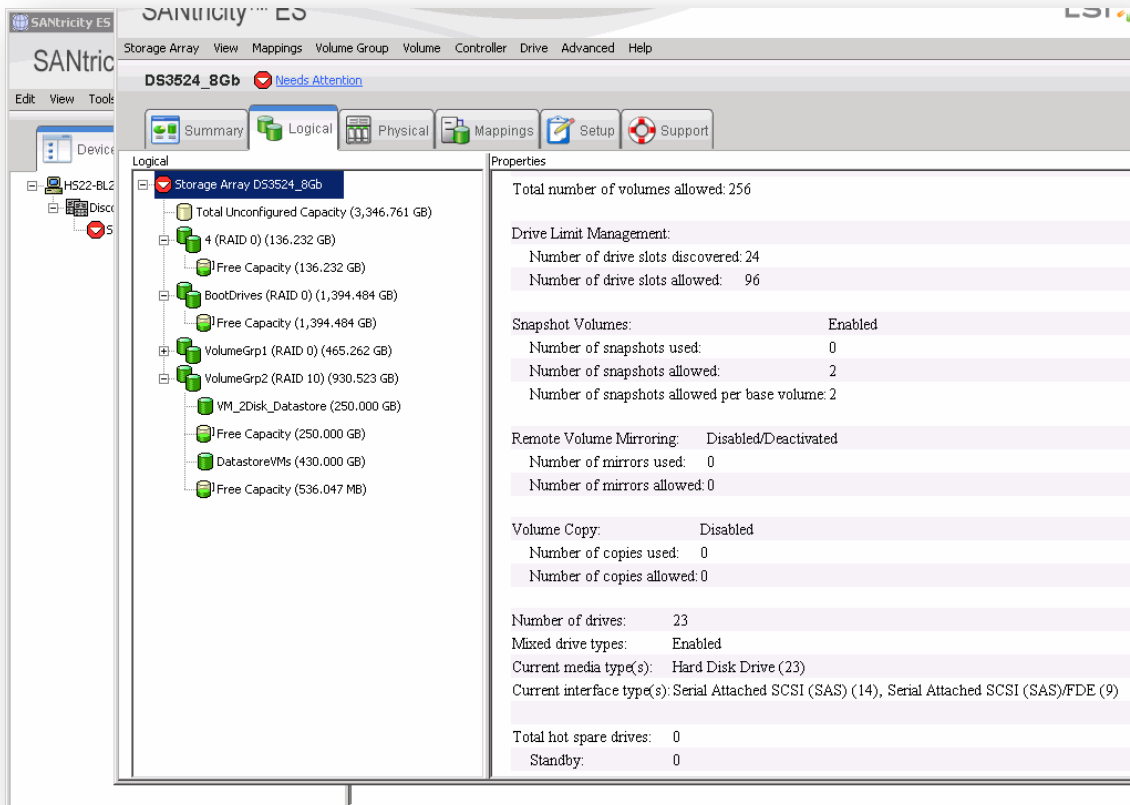


Figure 25. Selecting the option to obtain more information about the storage failure

The resulting Recovery Guru screen provided the following notification: **Volume Not on Preferred Path.**

Addressing the notification

The guru's assessment was accurate since Emulex had physically removed the storage controller. In response, Emulex placed the controller offline, physically replaced the device and placed it online. The procedure was as follows:

1. Using SANtricity ES Storage Manager, select the controller in the Array Management Window.
2. Select **Advanced**→**Recovery**→**Place Controller**→**Offline**.
3. Complete the instructions provided and then select **Yes**.



4. Re-run the Recovery Guru to check for any other issues.
5. Replace the storage controller.
6. Select the controller in the Array Management window.
7. Select **Advanced**→**Recovery**→**Place Controller**→**Online**.
8. Complete the instructions provided and then select **Yes**.
9. Re-run the Recovery Guru to check for any other issues.

Redistributing volumes

After a failure has occurred on the storage array and has been corrected, use the Recovery feature of SANtricity ES Storage Manager to redistribute volumes to their preferred controller. Perform the following steps:

1. Select **Advanced**→**Recovery**→**Redistribute Volumes** menu option.

Note

If the Redistribute Volumes menu is grayed out (unavailable), volumes are already associated with their preferred controllers.

2. Re-run the Recovery Guru to check for any other issues.
All volumes should now be back online and verifiable in the System Storage plug-in; SANtricity ES Storage Manager should show healthy status.
3. If problems persist, contact your IBM Technical Support Representative.

Summary

In general, it takes careful planning to achieve the optimal configuration for an ESXi infrastructure featuring 8Gb/s Fibre Channel storage. The POC environment described in this white paper provides a baseline for the configuration and deployment of ESX 4.1 on certified IBM hardware such as BladeCenter HX5 blade servers and DS3500 storage arrays to create a robust, highly-available, virtualized infrastructure.



Appendix A – Bill of materials

Qty	Part number	Description
Blade server		
1	41Y4842	IBM BladeCenter H Chassis
1	69Y3048	IBM HX5 Type-7872AC1 Blade Server
2	M392B5273CH0-CH9	4GB DDR III SDRAM
7	HMT125V7BFR8C-H9	2GB DDR III SDRAM
1	18JDF25672PDZ1G4F1	2GB DDR III SDRAM
2	59Y6048	High Speed Expansion Card SSD Storage
1	46M6142	Fibre Channel Expansion Card
Server connectivity		
1	49Y4235	Emulex 10Gb Ethernet Virtual Fabric Adapter for IBM System x
1	46M6140	Emulex 8Gb Fibre Channel Expansion Card (CIOv)
Storage array		
1	1746A4D	IBM System Storage DS3524 Express Dual Controller Storage
2	68Y8432	8GB FC 4 Port Daughter Card
12	49Y1841	146GB SAS 2.5 HDD
Network switch		
1	46C7191	BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
4	44W3308	IBM 10GBase-SR 10Gb Fibre SFP+ Transceiver
Fabric switch		
2	44X1927	Brocade 8Gb SAN Switch Module for IBM BladeCenter
8		6-foot Fibre Channel cable
Storage management plug-ins		
1		IBM System Storage plug-in for vCenter Server 4.x
1		OneCommand Manager plug-in for vCenter



Appendix B – Installing the IBM System Storage plug-in

Perform the following steps to install the IBM DS3000 DS4000 DS5000 vCenter Management Plug-in (System Storage plug-in):

1. After registering on the portal, download the plug-in from <http://ibmdsseries.com/>.
2. Run the executable file from a Microsoft Windows client; follow the instructions provided by the installer.
3. Specify the IP address of the management server, as shown in Figure B-1.

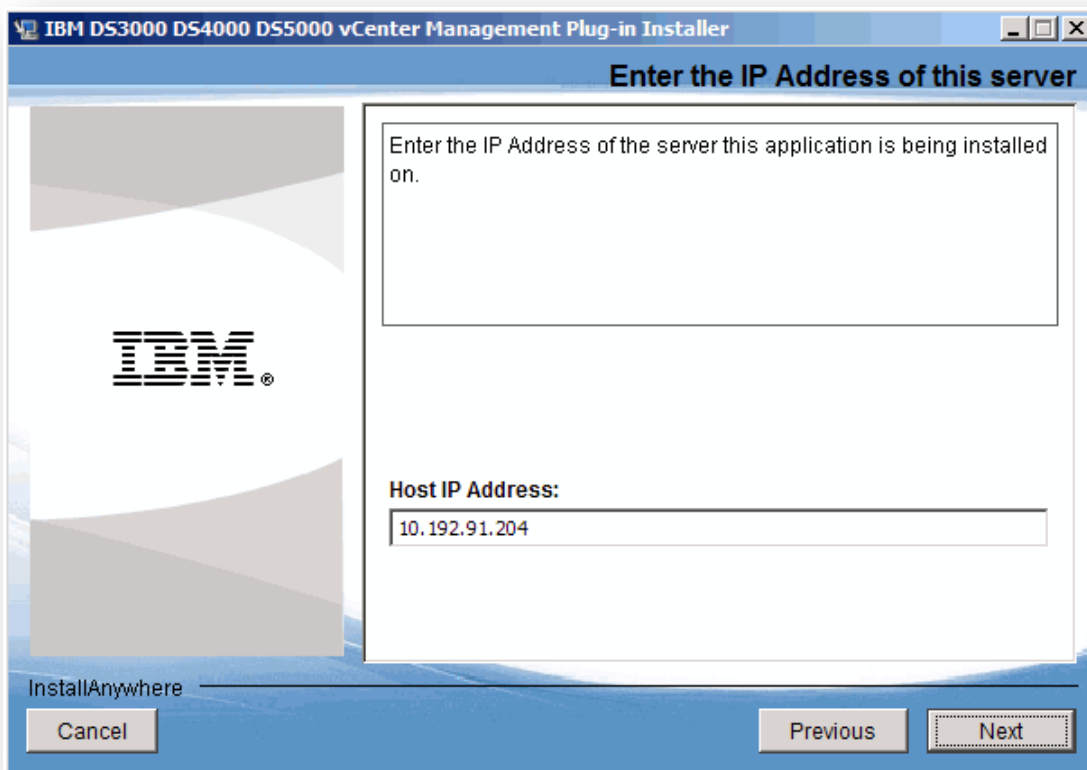


Figure B-1. Specifying the IP address of the management server

4. Specify the IP address of the vCenter Server host.
5. Enter an email address for administrator, followed by the user ID required to log the administrator into vCenter Server.
6. Specify the administrator's password.



7. Install the appropriate Secure Socket Layer (SSL) certificate and then select **Ignore** to continue.
8. When you reach the **Install Complete** screen, press **Done**, as shown in Figure B-2.

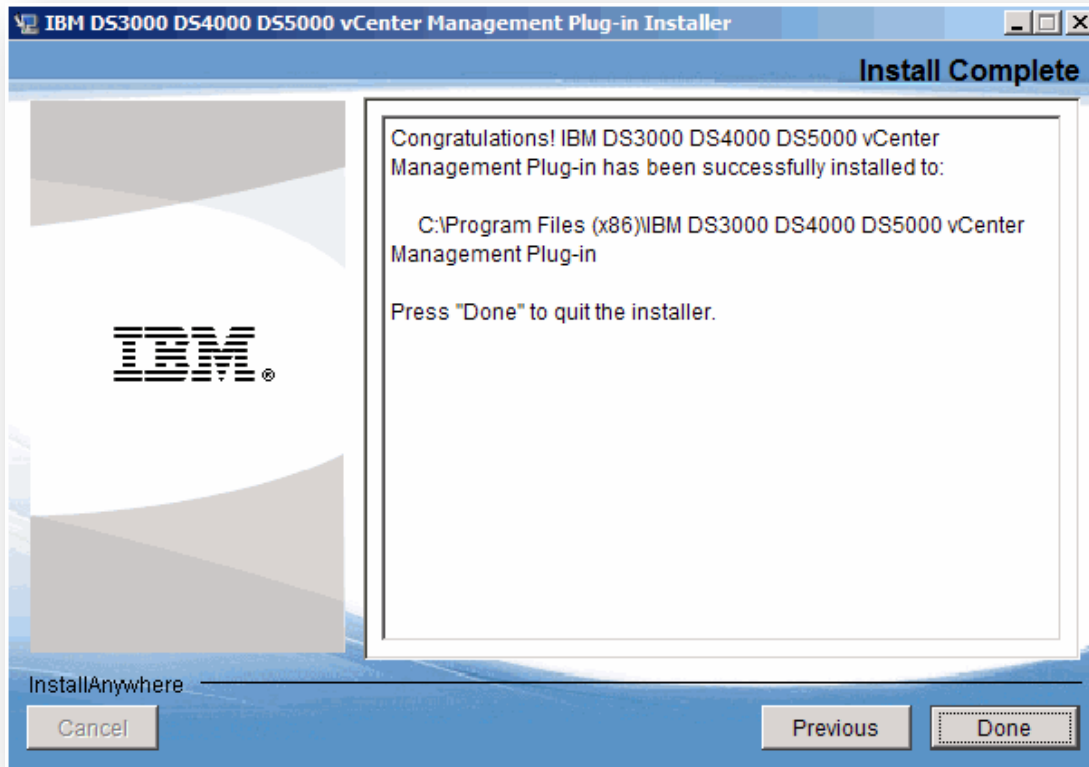


Figure B-2. Completing the installation

9. To use the features provided by the System Storage plug-in, log into vCenter Server via vSphere Client.



10. From the **Plug-ins** menu, select **Manage Plug-ins**, as shown in Figure B-3.

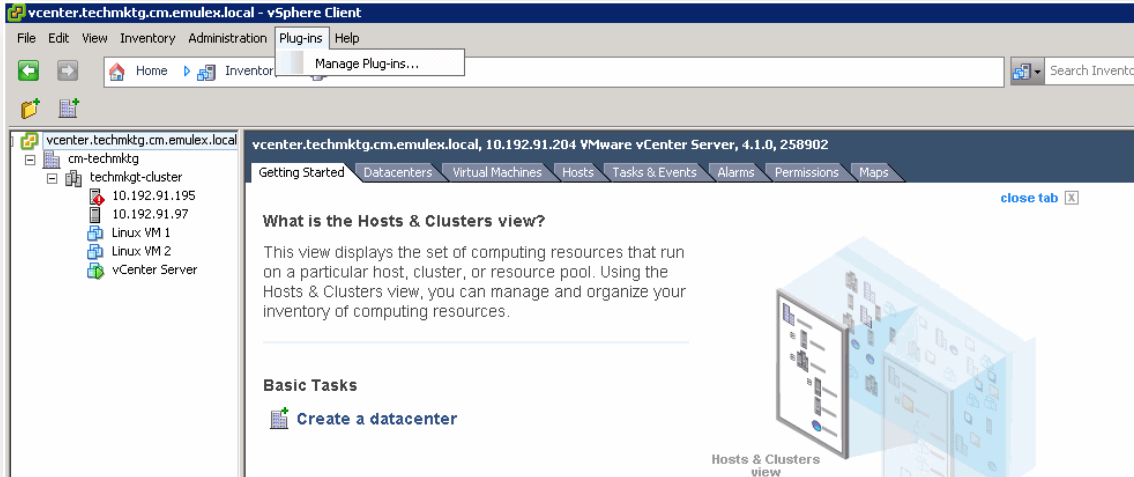


Figure B-3. Selecting management plug-ins

11. Verify that the System Storage plug-in (**IBM DS3000 DS4000 DS5000 vCenter Management Plug-in**) is present, as shown in Figure B-4.

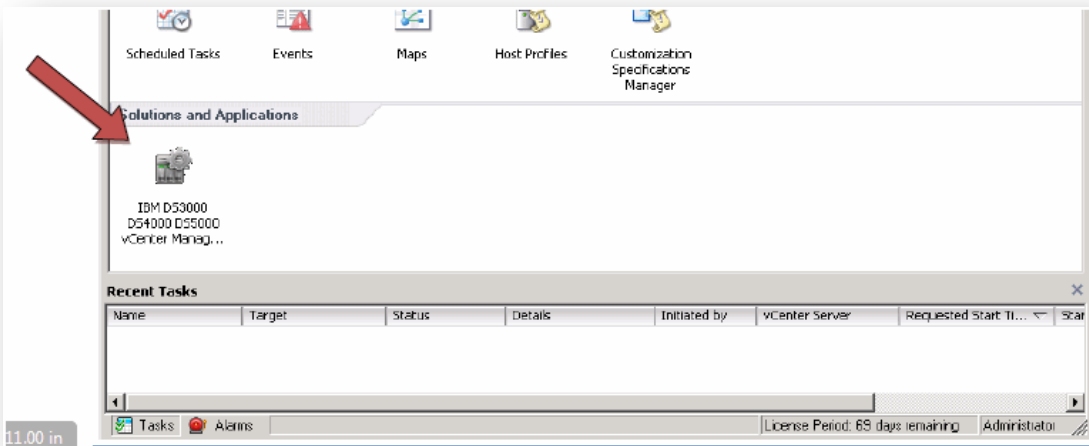


Figure B-4. Validating that the System Storage plug-in is present



Appendix C – Installing drivers and CIM providers

This appendix provides instructions for the following:

- Installing the Emulex VFA NIC driver for ESXi
- Updating the Emulex 8Gb/s Fibre Channel expansion card driver for ESXi
- Installing Emulex CIM providers for UCNA and Fibre Channel

Installing the VFA NIC driver

Download the VFA NIC driver from the [VMware website](#). Use the following command to install the driver:

```
vihostupdate.pl --server xxx.xxx.xxx.xxx --install --bundle  
c:\temp\SVE-be2net-2.102.404.0-offline_bundle-281453.zip
```

Updating the 8Gb/s Fibre Channel expansion card driver

IMPORTANT

Install or update the driver and CIM providers before installing OCM or the OCM plug-in for vCenter.

Carry out the following steps to update the driver¹²:

1. [Download](#) the driver, which is an ISO image.
2. Create a CD/DVD from the image.
3. Mount on CD/DVD on a Windows client system.
4. From the vCenter Server, place the host in maintenance mode.
5. Remotely login to the ESXi host.
6. Run the following command to install the driver:

```
vihostupdate.pl --server xxx.xxx.xxx.xxx --install --bundle  
D:\offline-bundle\
```

where:

- **xxx.xxx.xxx.xxx** is the ip address of the ESXi management network
 - **D:** is the CD/DVD drive
7. When prompted, enter the appropriate user name and password for logon.
 8. After the update is complete, reboot the ESXi host for the changes to take effect.
 9. Exit vCenter Server maintenance mode.

¹² The IBM product name for this driver is *IBM BladeCenter LPe1205-CIOv 8Gb/s HBA driver*.



Installing the CIM providers

IMPORTANT

Before installing the CIM providers, ensure that the drivers for your Emulex adapters have been installed.

Carry out the following steps to install the CIM providers:

1. [Download](#) and install the CIM provider package.
2. Copy the package to a temporary directory on the remote system.
3. Place the host in maintenance mode.
4. Run the following command to install the package (offline bundle):

```
vihostupdate.pl --server x.x.x.x --install --bundle C:\temp\elx-  
esx-4.1.0-emulex-cim-provider-3.2.30.1-offline_bundle-364582.zip
```

5. When prompted, enter the appropriate user name and password for logon.

You should then see that the update has completed successfully, as shown in Figure C-1.

```
C:\Program Files\VMware\VMware vSphere CLI\bin>vihostupdate.pl --server 10.192.90.38 --install --bundle "C:\Users\Administrator\Downloads\elx-esx-4.1.  
0-emulex-cim-provider-3.2.30.1-offline_bundle-364582(1).zip"  
Enter username: root  
Enter password:  
Please wait patch installation is in progress ...  
The update completed successfully, but the system needs to be rebooted for the changes to be effective.  
C:\Program Files\VMware\VMware vSphere CLI\bin>_
```

Figure C-1. Successful CIM provider installation



More information

IBM Sizing Guides – VMWare ESX 4.0 WLE-based Sizing Guide	http://www-304.ibm.com/partnerworld/wps/sizing/sizingguide/sqListView.jsp?guide_id=sgq75032541151219002 30&data_guide_private=
IBM System Storage DS3500: Introduction and Implementation Guide	http://www.redbooks.ibm.com/redpieces/abstracts/sg247914.html?Open
IBM Midrange System Storage Implementation and Best Practices Guide	http://www.redbooks.ibm.com/abstracts/sg246363.html
BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter	http://www.redbooks.ibm.com/abstracts/tips0708.html
VMware Implementation with IBM System Storage DS4000/DS5000	http://www.ibm.com/storage
Installing VMware ESXi Server 4.1 Installable – IBM System HX5 (7872)	http://www-947.ibm.com/support/entry/portal/docdisplay?brand=5000020&Indocid=MIGR-5085425
IBM BladeCenter Virtual Fabric Solutions	http://www.redbooks.ibm.com/redpapers/pdfs/redp4673.pdf
IBM Fix Central (download device driver, BIOS, firmware)	http://www.ibm.com/support/fixcentral/
VMware vSphere 4.1	https://www.vmware.com/support/pubs/vs_pages/vsp_pubs_esxi41_e_vc41.html
VMware Fibre Channel SAN Configuration Guide	http://www.vmware.com/pdf/vsphere4/r41/vsp_41_san_cfg.pdf
VMware device-driver update CD's	http://downloads.vmware.com/info/datacenter_download/vmware_vsphere_4/4#drivers_tools/

To help us improve our documents, please provide feedback at implementerslab@emulex.com.

© Copyright 2011 Emulex Corporation. The information contained herein is subject to change without notice. The only warranties for Emulex products and services are set forth in the express warranty statements accompanying such products and services. Emulex shall not be liable for technical or editorial errors or omissions contained herein.

LightPulse and OneCommand are registered trademarks of Emulex Corporation. IBM is a registered trademark in the U.S. and other countries.

VMware is a registered trademark of VMware Corporation.

12-0135 August, 2011



www.emulex.com

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