



# Using EonStor® FC-host Storage Systems in VMware Infrastructure 3 and vSphere 4

## Application Note

### ***Abstract***

This application note explains the configure details of using Infortrend FC-host storage systems with VMware Infrastructure 3 and vSphere 4 to deliver a reliable, flexible and efficient data center.

## VMware Virtualization

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

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

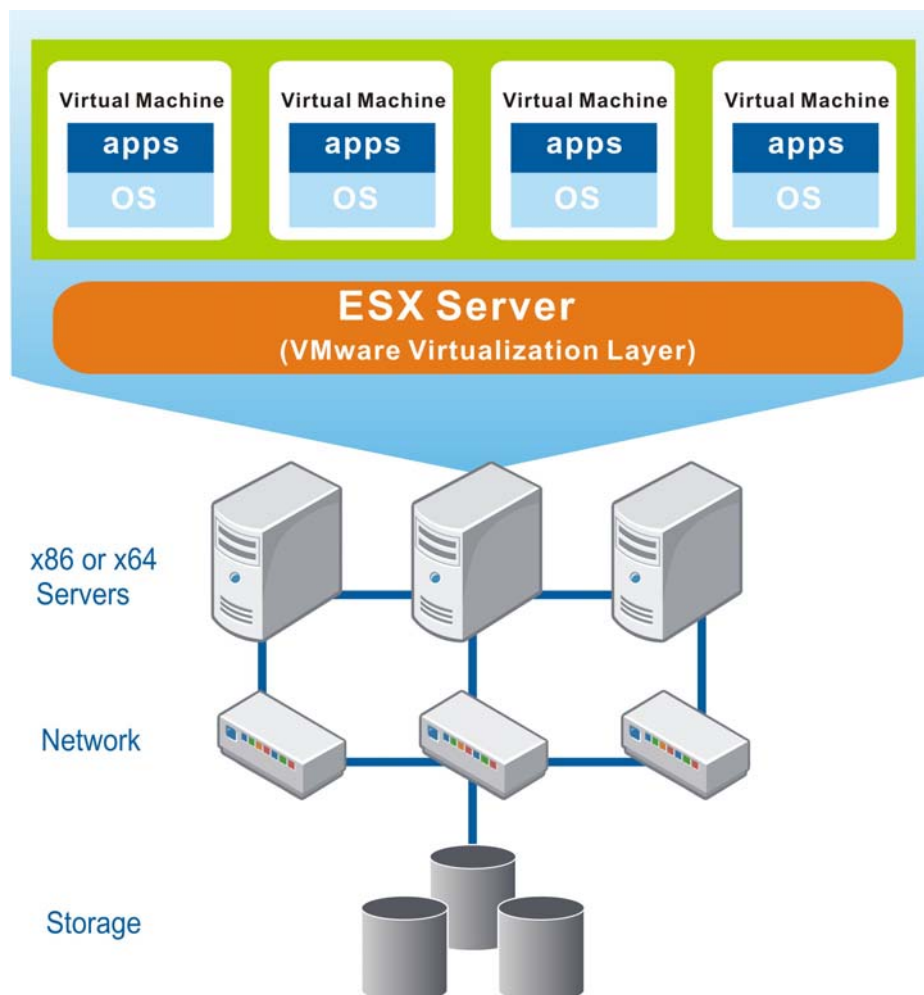


Figure 1. VMware ESX Server

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

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


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

## EonStor Storage Systems in VMware Virtualized

### Environments

Storage Area Network (SAN) is an architecture bringing all storage resources into a pool and making them shared among multiple hosts. It is the most suitable storage architecture for virtualized environments because when the operating systems (OS) and applications are moved around virtual machines, the data do not have to be copied and moved accordingly. Infortrend provides both FC-host RAID arrays and iSCSI RAID arrays to support SAN in VMware virtualized environments. EonStor arrays' features of centralized management and configuration platform, redundancy designs, and smart scalability can bring many benefits to the environments.

In a data center implemented with numerous storage systems, a powerful storage management tool is absolutely necessary. Infortrend's proprietary management suite, SANWatch, allows system managers to configure, administer and monitor multiple EonStor arrays locally or remotely through a user-friendly graphic interface. As shown in **Figure 2**, by integrating Infortrend's storage management tool (circled in red) into a VMware virtualized environment (circled in blue) can make the infrastructure more robust and easier to maintain.



**Figure 2.** Using SANWatch Management in a VMware Virtualized Environment

Besides the software complement, the fault-tolerant hardware modules and RAID functionalities of EonStor arrays further improve system availability. When controllers, power supplies, cooling modules or hard drives fail, they can all be hot-swapped without causing downtime or data loss.

Along with data growth, the capacity of a single EonStor subsystem can be scaled by connecting to expansion enclosures, and the scaling will not downgrade their industry-leading performance. By allowing users to start with the capacity they need now and expand for more when future needs arise, the flexible way of scaling minimizes

waste of investment and in turn contributes to enhancing virtualization's core value of optimized resource utilization.

To ensure seamless integration of EonStor Storage Area Network (SAN) storage into VMware virtualized environments, Infortrend has performed comprehensive testing to verify its full compatibility with VMware Infrastructure 3 (VI3) and vSphere 4. VI3 is VMware's flagship server and datacenter product and vSphere 4 is its renamed next major version. These powerful software suites can optimize and manage IT environments through virtualization. They both greatly reduce operating costs and increase IT service availability, security and scalability while providing the flexibility to choose any OS, application and hardware. Building on the proven power of VI3 platform, vSphere 4 delivers many enhanced features in simplified management, application services, infrastructure services and compatibility and third-party extensibility. For more details on the new features of vSphere 4, please refer to [http://www.vmware.com/support/vsphere4/doc/vsp\\_40\\_new\\_feat.html](http://www.vmware.com/support/vsphere4/doc/vsp_40_new_feat.html).

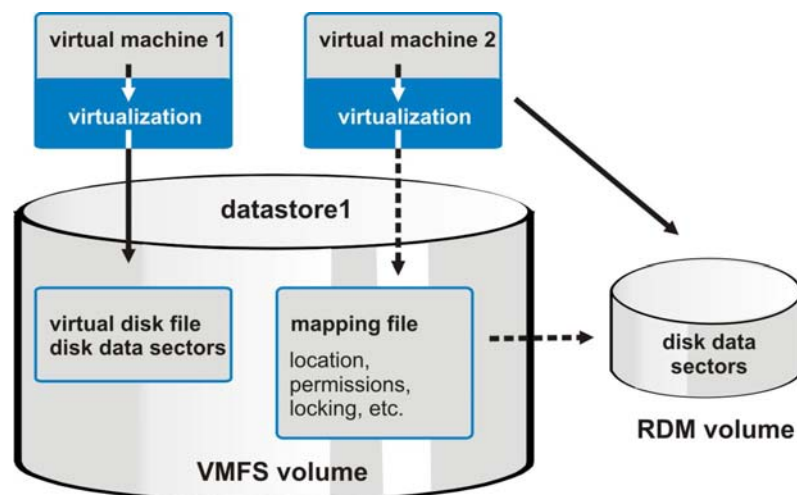
## Configuration Considerations

### Data Formats

To make the data volumes on EonStor arrays accessible to ESX servers, they have to be configured as either VMFS (Virtual Machine File System) volume or RDM (Raw Device Mapping) volume<sup>1</sup>. VMFS is VMware's proprietary clustered file system. It is the most common access method. If users would like to allow multiple virtual machines to run on and multiple physical servers to access a single volume, they should configure the volume with the VMFS format. Another alternative to make virtual machines access data volumes on the storage is RDM. Virtual machines access VMFS volumes and RDM volumes in different ways. As shown in **Figure 3**, virtual machines can directly access a virtual disk in the VMFS format but their access to the RDM volume is enabled through a mapping file in the VMFS volume. This mapping file contains metadata that redirects disk access to the physical devices.

---

<sup>1</sup> Maximum size of an RDM volume in 2TB.



**Figure 3.** Different Ways of Accessing a VMFS Volume and an RDM Volume

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

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

### Deployment of VMFS Volumes

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

1. Virtual machine boot disks and application data should be stored in separate VMFS volumes. Most I/Os issued to boot disks involves paging activities and are sensitive to response time. By separating boot disks from application data, the risk of prolonged response time due to application related I/O activities can be mitigated.
2. Database platforms for enterprise data management, such as Microsoft SQL Server or Oracle, often use active logs and/or recovery data structures to track data changes. In cases of unplanned application or operating system disruptions, these active logs or recovery data structures are critical in ensuring system recovery and data consistency. Therefore, all virtual machines supporting such database platforms should be provided with an independent VMFS volume for storing active log files and recovery data structures. Furthermore, if the files or structures are mirrored, the source and the target should be stored in separate VMFS volumes.
3. Application data, including database files, should be stored in a separate VMware file system. Furthermore, this file system should not contain any structures that are

critical for application and/or database recovery.

4. It is recommended that the VMFS volumes are no more than about 80% full. This ensures that administrators would not suddenly run out of space to accommodate user data and VMware snapshots for virtual machines.

## **RAID Level**

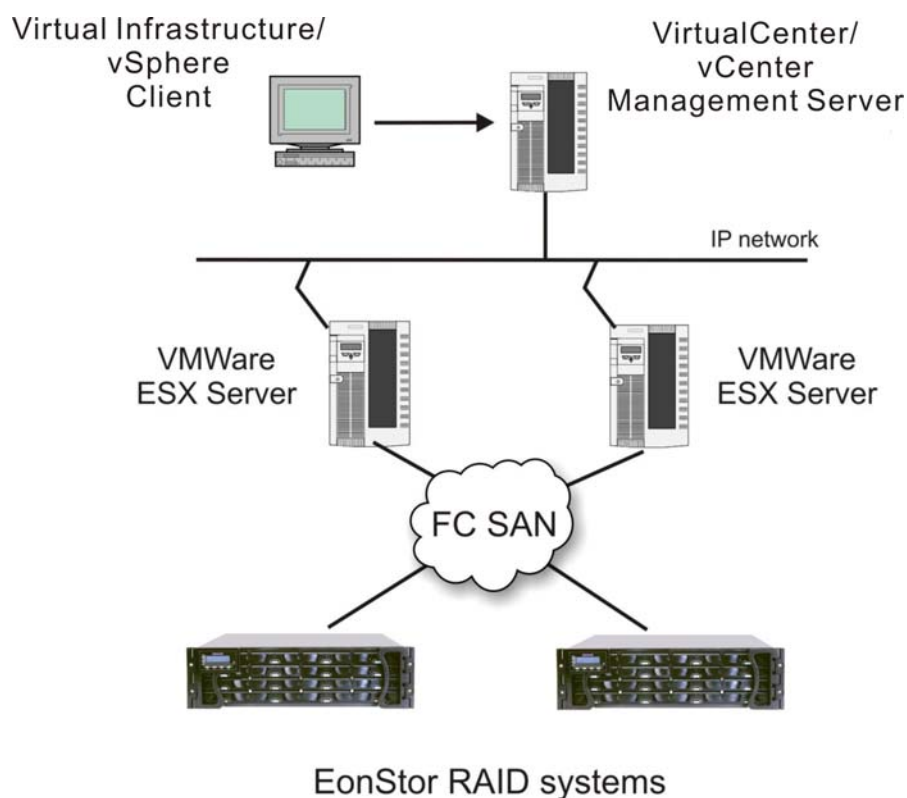
EonStor storage arrays allow users to protect their data volumes with various RAID levels, including RAID 1, RAID 3, RAID 10, RAID 5 and RAID 6. Data volumes in the same storage array can be protected with different RAID levels. The following are general guidelines for you to configure RAID levels for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

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

## **Application Example**

### **Physical Topology of the Test Environment**

As shown in **Figure 4**, our test environment consists of one management server, two computing servers, EonStor FC-host storage arrays and IP networks.



**Figure 4.** Physical Topology of the Test Environment

The vCenter/VirtualCenter Management Server is installed on an enterprise standard x86 sever that runs Windows Server 2003 Enterprise. The two computing servers which run VMware ESX Server are also enterprise standard x86 servers. They are connected to the same network and the same storage systems. For RAID storage, we use EonStor S16F-R1430, which combines four 4Gb/s Fibre Channel host ports with sixteen SAS disk drives in a 3U enclosure delivering high fault-tolerant capability, performance and reliability.

### Configuration Steps

#### **Step 1: Connect EonStor Arrays and Scan Available Storage**

After connecting EonStor arrays to the ESX Server through an FC switch, we go to the *Configuration* tab in vCenter/VirtualCenter GUI, click *Storage Adapters* in the *Hardware* panel on the left, and then click *Rescan* to see the available Fibre Channel adapters. By clicking an adapter, we can see its corresponding SCSI Targets, which means the storage devices. Here, we use *vmhba2* as our Fibre Channel adapter. (See **Figure 5**)

**Storage Adapters** [Rescan...](#)

Device	Type	SAN Identifier
<b>QLA2422</b>		
vmhba2	Fibre Channel	21:00:00:e0:8b:90:56:2
vmhba3	Fibre Channel	21:01:00:e0:8b:b0:56:2
<b>AIC-8902 U320 OEM</b>		
vmhba0	SCSI	
vmhba1	SCSI	

**Details**

**vmhba2**  
 Model: QLA2422  
 WWPN: 21:00:00:e0:8b:90:56:2a  
 Targets: 1

**SCSI Target 0** [Hide LUNs](#)

Path	Canonical Path	Capacity	LUN ID
vmhba2:0:0	vmhba2:0:0	302.34 GB	0
vmhba2:0:1	vmhba2:0:1	10.00 GB	1
vmhba2:0:2	vmhba2:0:2	20.00 GB	2
vmhba2:0:3	vmhba2:0:3	30.23 GB	3

Figure 5. Choose a Fibre Channel Adapter

### Step 2: Create Data Volumes

Once we find the available SCSI targets, we click *Storage* in the *Hardware* panel on the left, and then click *Add Storage* to create data volumes. As shown in the figure below, we create a VMFS volume and name it *DataStorage300G*. (See **Figure 6**)

Summary Virtual Machines Resource Allocation Performance Configuration **Tasks & Events** Alarms Permissions Maps

**Hardware**

- Processors
- Memory
- Storage (SCSI, SAN, and NFS)
- Networking
- Storage Adapters
- Network Adapters

**Software**

- Licensed Features
- DNS and Routing
- Virtual Machine Startup/Shutdown
- Security Profile
- System Resource Allocation
- Advanced Settings


**Storage** [Refresh](#) [Remove](#) [Add Storage...](#)

Identification	Device	Capacity	Free	Type
DataStorage300G	vmhba2:0:0:1	302.25 GB	202.62 GB	vmfs3
DataStorage400G	vmhba2:0:4:1	395.50 GB	394.88 GB	vmfs3

**Details** [Properties...](#)

**DataStorage300G** Capacity

Location: /vmfs/volumes/4773b074-d...

99.63 GB  Used   
 202.62 GB  Free

Path Selection	Properties	Extents
Fixed	Volume Label: DataStorage..	vmhba2:0:0:1 302.34 ...
	Datastore Name: DataStorage..	Total Formatted Capacity 302.25 ...

Paths	Formatting
Total: 1	File System: VMFS 3.21
Broken: 0	Block Size: 1 MB
Disabled: 0	

Figure 6. Created Datastore: DataStorage300G

#### Step 4: Add Virtual Disks to the Virtual Machine

After a data volume is created, we can use it to create virtual disks for virtual machines. In the vCenter/VirtualCenter GUI, you first click *Virtual Machines* tab, right click on the virtual machine, and then select *Edit Settings*. In the *Hardware* tab, click *Add*. Then select *Hard Disk*, and click *Next* to complete the new virtual disk configuration process. In this example, in addition to the default disk (Hard disk 1), we create and map a VMFS-base virtual disk (Hard disk 2) and a raw device (Hard disk 3) to the virtual machine. (See **Figure 7**) These virtual disks are ready to be leveraged by the virtual machine.

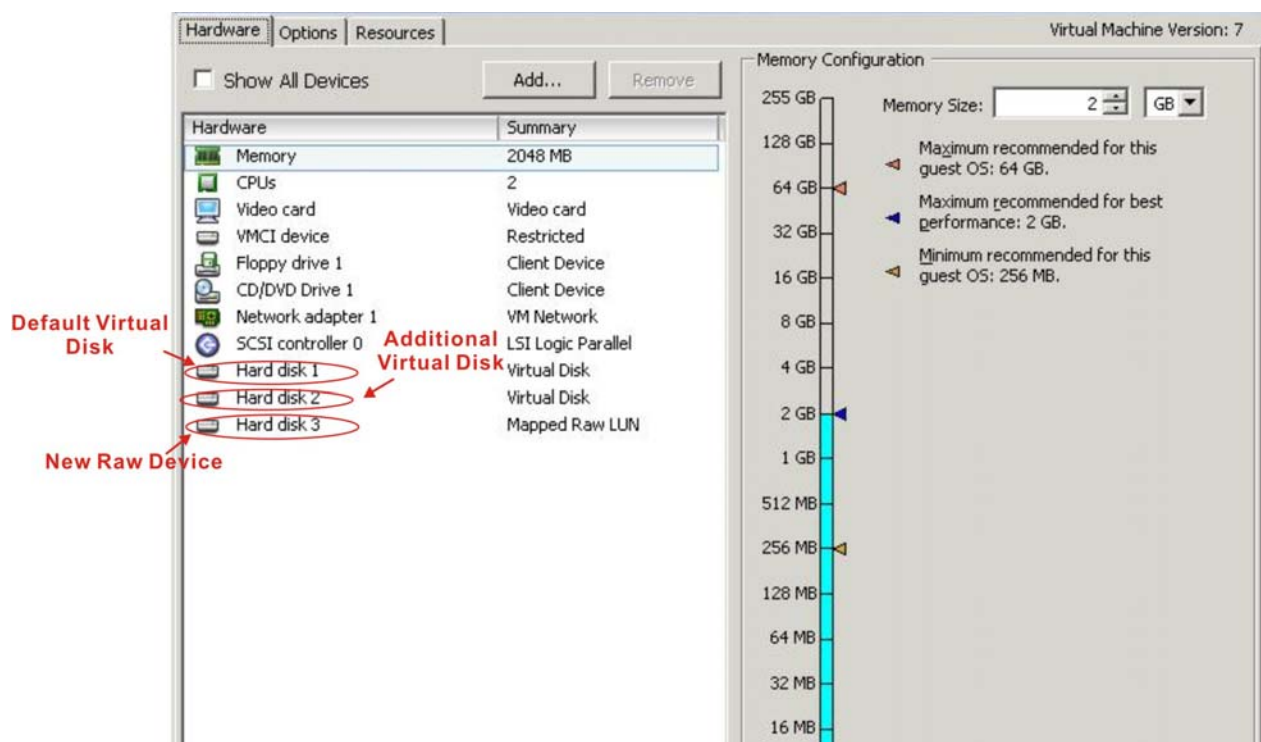
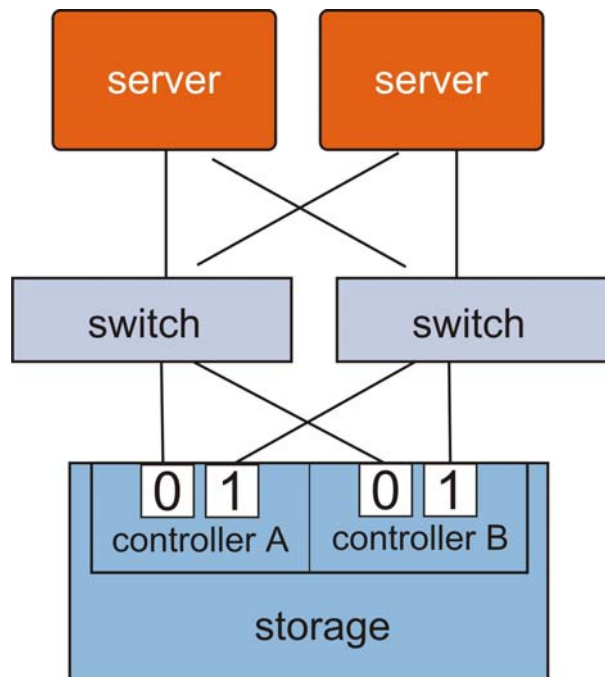


Figure 7. Virtual machine Property View

## Basic Troubleshooting and FAQ

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



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

### Networking: example

The screenshot shows the VMware vCenter Configuration console for a virtual switch. The 'View' is set to 'Distributed Virtual Switch'. The 'Networking' section is expanded, showing a 'Virtual Machine Port Group' named 'VM Network' connected to a 'Physical Adapters' named 'vmnic0' with a speed of '1000' and 'Full' duplex. Below this, a list of 11 virtual machines is shown, each with a network icon and a refresh button. The VMs listed are: rh-bus-local, win-lsi-local, win-bus-local-a, rr1-lsi-tmp, rh5.2, rh-bus-local-a, rh-lsi-local-a, win-lsi-local-a, rh-lsi-local, rr3-lsi, and rh5.2-tmp. Below the VM list, the 'VMkernel Port' section shows 'VMkernel' with IP address 'vmk0 : 192.168.177.253'. The 'Service Console Port' section shows 'Service Console' with IP address 'vswif0 : 192.168.200.60'.

## Storage Adapters: example

The screenshot shows the VMware Configuration console for a host. The 'Storage Adapters' tab is active, displaying a table of installed adapters:

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

Below the table, the 'Details' section for the selected 'vmhba33' adapter is shown:

**vmhba33** Properties...

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

View: **Devices** Paths

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

## Maps: example

The screenshot shows the VMware Maps console. The main area displays a network map with the following components and connections:

- VM Network** (Network icon) connected to **private** (Network icon).
- private** connected to **c-w2k3-b** (Host icon).
- c-w2k3-b** connected to **esx15** (Host icon).
- c-w2k3-b** connected to **192.168.140.15** (IP address icon).

On the right, the 'Map Relationships' panel is visible, showing the following settings:

**Virtual Machine Resources**

Host Options

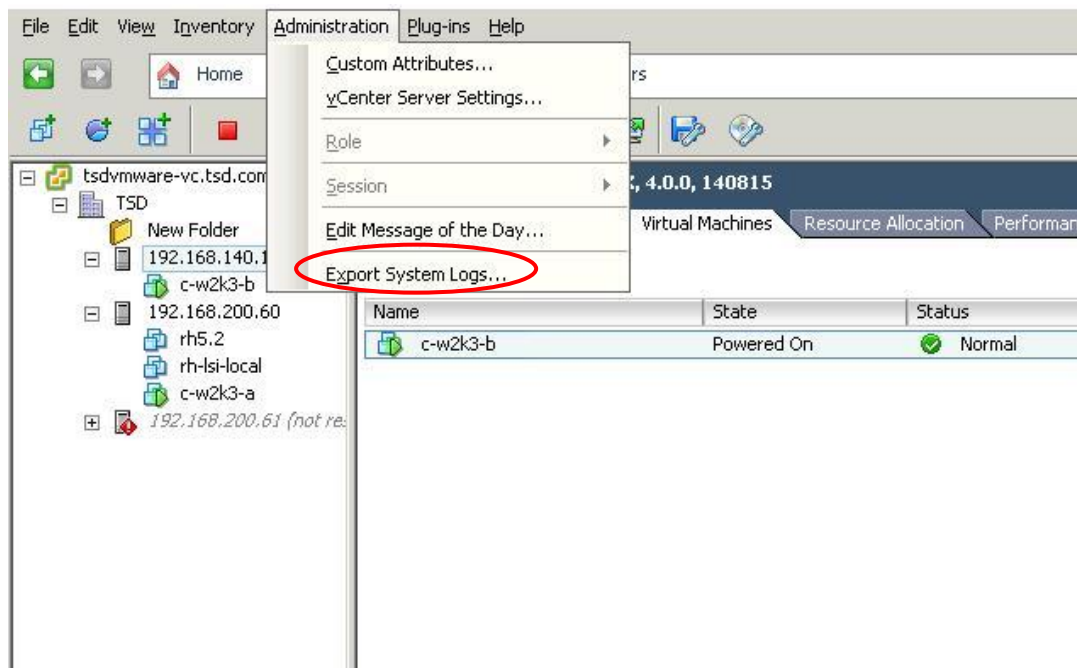
- Host to VM
- Host to Network
- Host to Datastore

VM Options

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

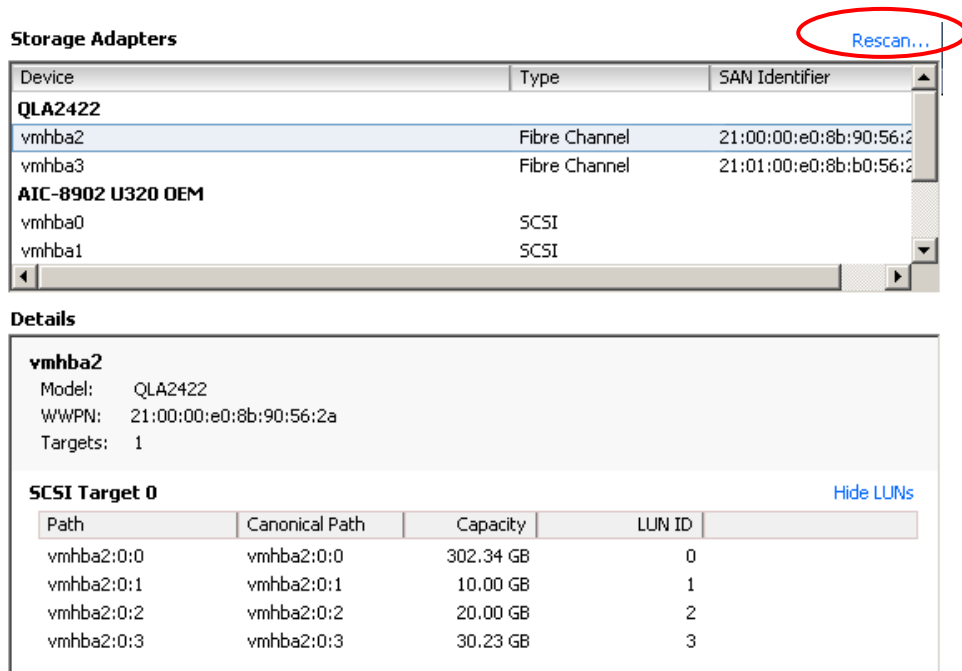
Apply Relationships

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



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

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



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

3. What are the storage configuration limitations in a VMware virtualized environment? Please check [http://www.vmware.com/pdf/vsphere4/r40/vsp\\_40\\_config\\_max.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_config_max.pdf) for details.
  
4. How can I make data paths successfully failover when redundant controllers failover?  
If you are using ESX 4.0 with EonStor storage arrays installed with firmware ver.3.64 or later, system would automatically handle this without any manual configuration. However, if you are using ESX 3.x, please add a footnote for redundant controller storage following the steps in the application note: [http://www.infortrend.com/doc/appNote/APP\\_VMware\\_footnote\\_1117.pdf](http://www.infortrend.com/doc/appNote/APP_VMware_footnote_1117.pdf).
  
5. Can virtual machines be migrated to a different data volume without interruptions?  
Yes, virtual machines can be migrated online to another data volume. Please check the following links for details:  
[http://www.vmware.com/products/vi/storage\\_vmotion.html](http://www.vmware.com/products/vi/storage_vmotion.html)  
<http://blogs.vmware.com/vi/2008/06/storage-vmotion.html>
  
6. If I would like to implement multi-pathing, is there any special settings I should do on my storage? Should I install EonPath?  
No, there is no special settings for storage. You can just follow general multiple-pathing configurations. Moreover, since VMware supports native multiple-pathing functions, you need not install other drivers, including EonPath, for multiple-pathing implementations.