



## SAN Management Tips for VMware vSphere 4

### Optimize SAN-attached storage for VMware vSphere 4 using Emulex LightPulse Fibre Channel HBAs and CNAs

#### At a Glance

VMware vSphere 4 enables data centers to leverage server virtualization to save capital and operating expenses and deliver computing resources in an agile, cost-effective way. Built on the ESX/ESXi 4.0 hypervisor, VMware vSphere 4 encompasses a comprehensive suite of products to optimize and manage the VMware virtualization environment. VMware vSphere 4 requires a scalable and reliable data infrastructure to enable the flexibility needed for today's server virtualization and tomorrow's next-generation cloud computing. Emulex LightPulse 8Gb/s Fibre Channel Host Bus Adapters and 10Gb/s Converged Network Adapters meet this requirement with high-speed, high-performance connections to storage area networks (SANs) based on the proven Fibre Channel protocol.

This Tech Brief discusses ten tips to effectively optimize and manage SAN-attached storage with VMware vSphere 4.

#### Introduction

VMware® vSphere® 4 provides a comprehensive suite of applications to manage and optimize a virtualization environment built on the ESX/ESXi 4.0 hypervisor. Networked storage is mandatory to use the most powerful features in VMware vSphere 4 and a storage area network (SAN) based on Fibre Channel technology is the best solution for most data centers.

Emulex® LightPulse® Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) deliver the performance, reliability and scalability that fully enable the benefits of a Fibre Channel SAN. To complete the solution, the HBAnyware® management application provides comprehensive management for all Emulex adapters in the network. HBAnyware 4.1 is a new release that adds online HBA and CNA management for ESXi servers to fully support the VMware environment.

Effective storage usage is a key to optimizing a virtualized data center. This tech brief provides an overview of tips and technologies to help ensure an effective implementation of VMware vSphere 4 with Emulex HBAs and CNAs.

#### 1. Installing Emulex Adapters

Drivers for Emulex HBAs and CNAs are included in the distributions for ESX 4.0 and ESXi 4.0. Drivers are loaded automatically when ESX/ESXi 4.0 detects Emulex HBAs and/or CNAs.

CNAs support the Fibre Channel over Ethernet (FCoE) protocol which allows Fibre Channel I/O to be transmitted over a 10Gb/s Ethernet network. ESX/ESXi 4.0 will see the CNA as both a Fibre Channel HBA and a network interface card (NIC). The CNA Fibre Channel function is supported with the same driver that's used for HBAs. The IP function uses a different driver.

When available, supported driver updates for both HBAs and CNAs can be downloaded from the VMware web site.

# SAN Management Tips for VMware vSphere 4

## 2. Installing HBAnyware

The HBAnyware management application provides comprehensive management for all Emulex adapters in the network. It supports both a command line interface (CLI) and a graphical user interface (GUI). In addition to ESX/ESXi support, a single HBAnyware console can manage adapters on servers running any of the supported operating systems, including versions of Windows Server, Linux and Solaris.

HBAnyware installation and setup is different for ESX 4.0 and ESXi 4.0.

**ESX 4.0**—An HBAnyware applications kit should be downloaded from the Emulex web site and installed in the ESX Service Console. The applications kit includes the components needed to run the HBAnyware CLI from the local Service Console and manage the ESX 4.0 server remotely using an HBAnyware management console.

**ESXi 4.0**—Management of Emulex adapters on ESXi servers is done with the Common Interface Model (CIM). The ESXi 4.0 distribution includes an Emulex CIM provider that communicates with the CIM client that is part of HBAnyware 4.1. This enables HBAnyware 4.1 to provide remote management of adapters on all ESXi 4.0 servers in the network from a single management console. The ESXi CIM provider supports a basic subset of HBAnyware commands, and the HBAnyware 4.1 GUI displays supported commands only.

## 3. Firmware updates

Firmware is loaded in the flash memory of an adapter and is the interface between the device driver that runs in ESX and the physical adapter. Emulex's exclusive Service Level Interface (SLI) design allows a single version of the device driver to be used for all Emulex adapters. However, each adapter requires a version of the firmware that's specific to its model, which is installed as part of the manufacturing process.

Because of the time interval between manufacturing and installation, there may be newer versions of firmware available for an adapter model. Using the most current firmware version is generally recommended and updates can be downloaded from the Emulex web site. Firmware updates should be coordinated with supplier qualification and support.

Firmware updates are done using the HBAnyware management utility GUI or CLI. The HBAnyware GUI also supports a batch mode feature to automatically update all adapters in the network that are compatible with the firmware update file.

## 4. Boot from SAN (BFS)

With a growing trend toward blade and rack-mount servers, systems are often configured with no local disk and the operating system installed on and booted from a logical unit number (LUN) on a storage array.

BFS support is different for ESX 4.0 and ESXi 4.0.

**ESX 4.0**—BFS support for ESX 4.0 is the same as other server operating systems (Windows, Linux, Solaris and others) and uses boot code that is typically loaded in the adapter's flash memory at the time of manufacturing. Most adapters ship with a Universal Boot package with boot code for three system types. The x86BootBIOS boot code is used for ESX 4.0.

The only BFS setup requirement is specifying the LUN that will be the boot device. This is done by entering <CTRL> E or <ALT> E while the server is initializing. The x86BootBIOS setup is started automatically and a utility is run to specify LUNs that can be used as boot devices and to set other boot parameters. ESX 4.0 can then be installed on the specified LUN.

HBAnyware 4.0 and later also supports specifying BFS settings while the server is online. The system will use the new settings when the server is rebooted. This feature works well when servers need to be repurposed for different uses.

**ESXi 4.0**—The ESXi 4.0 operating system is always booted from a local drive, typically a USB drive. As a result, there is no requirement or support for BFS with ESXi 4.0.

## 5. Choosing the File System

As shown in Figure 2, ESX provides two options for virtual machines (VMs) to access storage—VMware Virtual Machine File System (VMFS) and Raw Device Mapping (RDM). VMFS is typically recommended, but there are scenarios where RDM may be the better option.

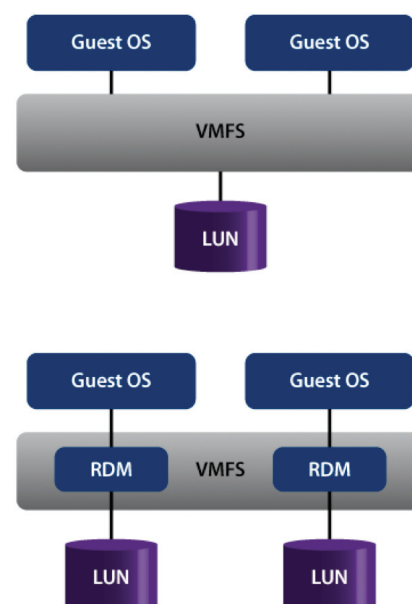


Figure 1 Storage access with VMFS and RDM.

# SAN Management Tips for VMware vSphere 4

**VMFS**—VMFS is a cluster file system. With VMFS, a single LUN is formatted as a datastore and files for VMs are stored in directories in the datastore. With this architecture, multiple VMs on multiple ESX servers can share the same datastore and associated LUN. On-disk locking is used to ensure that a VM can only run on one server at a time. With multiple VMs using a common LUN, the effort required for LUN management is minimized. VMFS is used by VMware VMotion, High Availability (HA), Distributed Resource Scheduler (DRS), and Distributed Power Management (DPM) to enable a dynamic, automated, self-optimized data center

**RDM**—RDM uses a mapping file inside VMFS which allows VMs to have direct access to LUNs. This adds more overhead to manage individual LUNs for VMs, but does provide support for applications that require direct access to storage. Examples would be NPIV support, snapshot applications running in a VM and Microsoft failover clustering scenarios. Supported by the underlying VMFS file system, RDM can also be used with VMware VMotion, HA, DRS and DPM.

## 6. Virtual Ports

Virtual ports (VPorts) use N\_Port ID Virtualization (NPIV), an ANSI T11 standard that enables a physical adapter port to be logged in to a fabric switch with multiple connections. Each VPort has a unique virtual worldwide port name (WWPN) address and identity on the SAN.

Using VMware vCenter, VMs can be created with an NPIV/VPort connection to the SAN. This allows VMs to have unique support for key SAN management features, including zoning, LUN masking, storage chargeback, backup and quality of service (QoS). Without NPIV, VMs use the shared identity of the host's physical adapters. RDM-based storage is required for NPIV/VPort support.




Figure 2 VMs and VPorts with unique access to LUNs.

The HBAnyware 4.1 management console provides new detail for VMs that are associated with VPorts for ESX 3.5 and ESX 4.0. This allows administrators to easily see the relationship between physical ports, VPorts and VMs to facilitate load balancing and SAN troubleshooting.

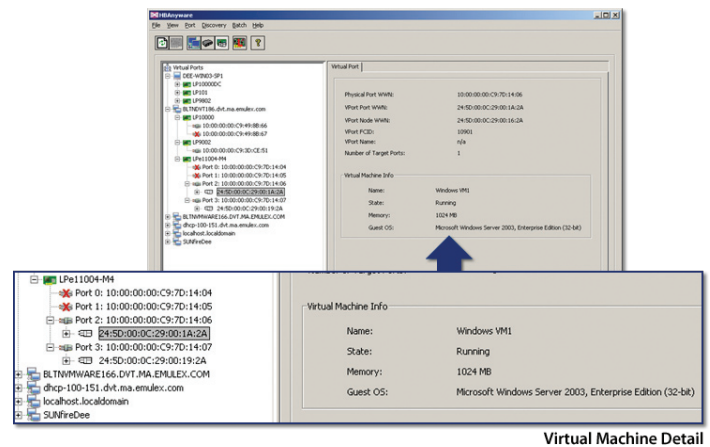


Figure 3 HBAnyware showing VPort and virtual machine detail.

## 7. High Availability

As server virtualization moves to business-critical applications, high availability support becomes mandatory. This is particularly true when any server downtime can affect multiple VMs and applications.

There are many hardware and software capabilities that support high availability. These include:

**Redundant Data Paths**—Redundant data paths are a critical best practice for Fibre Channel SANs. This includes multiple adapters in servers, multiple switches in the network, multiple ports on storage arrays and the use of RAID technology in the array.

ESX supports failover with redundant adapters that have access to the same LUN. One adapter is specified as the active path and is used for all I/O to the LUN. ESX does not support load balancing with I/O on multiple paths. If an adapter goes offline, I/O is automatically re-routed to an alternate path.

**VMware HA**—VMware High Availability (HA) uses a pool of server resources to provide failover for VMs. An agent runs on each server and maintains a “heartbeat” with other servers in the pool. If the heartbeat is lost for a specified time period, all of the VMs are restarted on other servers.

VMware HA requires shared storage to allow all servers in the resource pool to have access to VM disk files and other files used by applications. The adapters for all servers in the pool should be in the same zone. Adapters on servers outside the pool should be in a different zone.

## SAN Management Tips for VMware vSphere 4

### 8. Resource Mobility

VMware vSphere includes many technologies to optimize virtualized servers and networked storage. Two key elements are VMware VMotion and VMware Storage VMotion:

**VMware VMotion**—VMware VMotion provides the flexibility to move running VMs to different physical servers using storage as a pivot point. Adapters on both the originating and destination server must have access to the same storage as controlled by zoning or LUN masking.

VMware VMotion has many uses, including matching the application workload with the appropriate server resource, performing scheduled maintenance and retiring low-performance systems, all done with no disruption in service.

VMware VMotion also supports VPorts and NPIV. VMs will restart on the destination server using the existing VPort connection and SAN identity.

**VMware Storage VMotion**—VMware Storage VMotion provides flexibility with shared storage. VM disks files can be migrated between heterogeneous storage arrays with no disruption in service using the VM and physical host as a pivot point.

Use cases for Storage VMotion include transitioning to new arrays and migrating VMs to larger capacity or better performing LUNs. Fibre Channel zoning and LUN masking must be configured to insure the VM and host server have access to the datastore after the migration is completed.

### 9. Optimized Disk Utilization

VMware vSphere 4 introduces two new storage features to optimize disk usage:

**Thin Provisioning**—With VMware thin provisioning, storage is only used as needed. This eliminates the need to pre-allocate storage for both the virtual machine OS and the applications that it will be running, which typically results in significant amounts of committed, but unused disk space. There is a potential for free disk space to be unavailable when needed, so alarms and reports are provided to track usage and plan for storage expansion.

**Volume Grow**—Volume Grow allows a VMFS datastore to be expanded without affecting virtual machines that are running. A simple two-step process is used. First, the array management utility is used to increase the space allocated for the datastore LUN. Volume Grow is then used to expand the datastore up to the increased LUN size. The newly available space appears as a larger VMFS volume and a growth event is logged to VMware vCenter.

### 10. High Performance I/O

VMDirectPath is supported as an experimental technology in VMware vSphere 4 with Emulex LightPulse LPe12000 8Gb/s HBAs. VMDirectPath provides a direct connection from a guest OS to an I/O device, bypassing storage virtualization to deliver high-performance access to storage. An HBA port must be fully dedicated to a single VM, so VMDirectPath should be reserved for applications that require maximum I/O.



**World Headquarters** 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600  
**Wokingham, UK** +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177  
**Paris, France** +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547

[www.emulex.com](http://www.emulex.com)