



SAN Management Tips for VMware vSphere 4

Optimize SAN-attached storage for VMware® vSphere™ using Emulex® LightPulse® Fibre Channel HBAs and OneConnect™ 10GbE FCoE UCNAs

At a Glance

VMware vSphere 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 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 (HBAs) and OneConnect 10Gb/s Ethernet (10GbE) FCoE Universal Converged Network Adapters (UCNAs) meet this requirement with high-speed, high-performance connections to shared storage.

This Tech Brief discusses ten tips to effectively optimize and manage Emulex adapters with VMware vSphere 4.

Products

- Emulex OneConnect Universal CNA
- Emulex LightPulse Fibre Channel HBAs
- Emulex OneCommand Manager® 5.0 (formerly HBAnyware®)

Applications

- Organizations deploying SAN-based VMware vSphere 4/4.1
- Organizations deploying converged network solutions

Introduction

VMware vSphere 4.0 and 4.1 (collectively referred to as VMware vSphere 4) provide a comprehensive suite of applications to manage and optimize a virtualization environment built on the ESX/ESXi 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 HBAs and OneConnect 10GbE FCoE UCNAs deliver the performance, reliability and scalability that fully enable the benefits of VMware vSphere 4 deployments. To complete the solution, the OneCommand Manager application provides comprehensive management for all Emulex adapters in the network.

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 adapters.

1. Installing Emulex adapters

Drivers for Emulex LightPulse HBAs are included in the distributions for ESX/ESXi 4.0 and 4.1 and drivers are loaded automatically when HBAs are installed.

Emulex OneConnect UCNAs are 10Gb/s Ethernet adapters that provide network interface card (NIC) functions with options to support the Fibre Channel over Ethernet (FCoE) and iSCSI protocols. With FCoE, Fibre Channel I/O can be transmitted over a converged 10GbE network with full backward compatibility to existing Fibre Channel switches and arrays.

FCoE drivers for OneConnect UCNAs are included with ESX/ESXi 4.1 distributions. Additional drivers for OneConnect NIC, iSCSI and FCoE functions have been qualified with the VMware I/O Vendor Program (IOVP) for ESX 4.0 and 4.1 and can be downloaded from the VMware web site. The FCoE drivers also support LightPulse HBAs. OneConnect dual-port FCoE UCNAs present two Fibre Channel adapter ports and two NIC ports to the ESX hypervisor.

SAN Management Tips for VMware vSphere 4

2. Installing OneCommand Manager

The OneCommand Manager application provides comprehensive management for all Emulex adapters in the network through a command line interface (CLI) and a graphical user interface (GUI). In addition to managing adapters in ESX and ESXi environments, OneCommand Manager can also manage adapters in servers running any of the supported operating systems, including versions of Windows Server, Linux and Oracle Solaris.

The installation and setup of OneCommand Manager is different for ESX and ESXi environments, as follows:

ESX—The OneCommand Manager application kit needs to be downloaded from the Emulex web site and installed in the ESX Service Console. The applications kit includes the components needed to run the OneCommand Manager CLI from the local Service Console and manage ESX servers remotely using the OneCommand Manager management console.

ESXi—Management of Emulex adapters on ESXi servers uses the Common Interface Model (CIM). The ESXi 4.0 and 4.1 distributions include the Emulex CIM provider that communicates with the Emulex CIM client management interface enabling centralized, on-line management of adapters using OneCommand Manager. The ESXi CIM provider supports a basic subset of OneCommand Manager commands, and the OneCommand Manager GUI displays supported commands only.

VMware vCenter Server plug-in—The OneCommand Manager plug-in for VMware vCenter Server simplifies management of Emulex LightPulse Fibre Channel HBAs and OneConnect UCNAs within VMware vSphere environments by enabling administrators to centrally monitor, manage, update and run diagnostics on Emulex adapters, directly from VMware vCenter Server. The Emulex OneCommand Manager plug-in also offers asset information, such as adapter and driver firmware versions and driver configuration settings, directly within VMware vSphere environments. This allows IT administrators to natively update adapter firmware and driver settings.

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 Fibre Channel 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 OneCommand Manager GUI or CLI. The OneCommand Manager 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

With a growing trend toward blade and rack-mount servers, systems are often configured with no local disk. Instead, the servers are booted from the SAN using an assigned Logical Unit Number (LUN).

Boot from SAN (BFS) support is different for ESX and ESXi, as follows:

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

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 can then be installed on the specified LUN.

OneCommand Manager 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—The ESXi 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.

SAN Management Tips for VMware vSphere 4

5. Choosing the file system

As shown in Figure 1, 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.

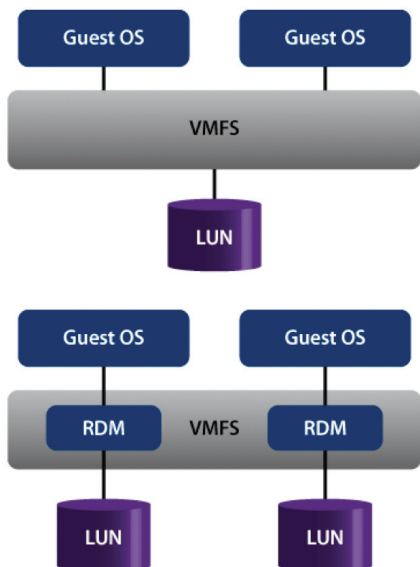


Figure 1 Storage Access with VMFS and RDM.

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 N_Port ID Virtualization (NPIV) support (described in #6 below), 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 NPIV, an ANSI T11 standard that enables a physical adapter port to be logged into a fabric switch with multiple connections. Each vPort has a unique virtual worldwide port name (WWPN) and identity on the SAN (see Figure 2).

VMware vCenter™ Server—Using VMware vCenter Server, 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.

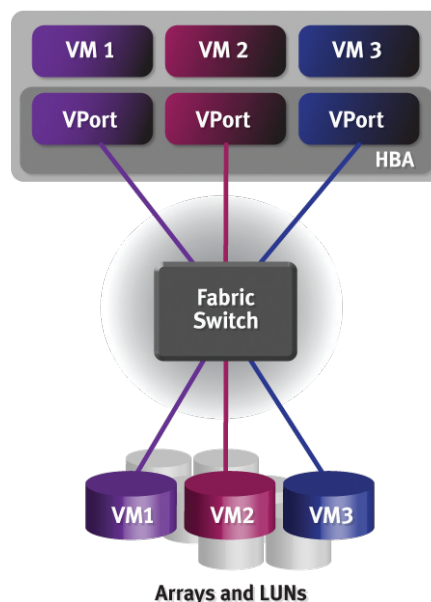


Figure 2 VMs and vPorts with unique access to LUNs.

SAN Management Tips for VMware vSphere 4

vPort mapping—OneCommand Manager also provides vPort to VM associations, allowing administrators to easily see the relationship between physical ports, vPorts and VMs. vPort mapping places critical information within easy reach, helping IT administrators reduce response time for corrective actions (see Figure 3).

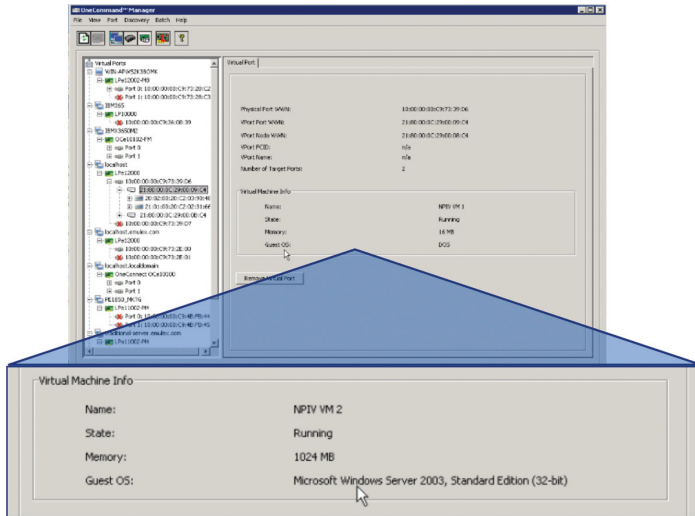


Figure 3 OneCommand Manager showing vPort to VM associations.

7. High availability

As server virtualization moves to business-critical applications, high availability support becomes mandatory. This is particularly true when server downtime affects 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 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.

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.

SAN Management Tips for VMware vSphere 4

9. Optimized disk utilization with vStorage

VMware vSphere 4.0 introduced a new technology called VMware vStorage with new 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 VM operating system 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 VMs 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.

VMware vSphere 4.1 added another important new capability:

Storage I/O Control—Using Storage I/O Control, rules and policies for storage QoS can be configured to set the business priority for each VM. When I/O congestion is detected, I/O resources are dynamically allocated to VMs based on QoS priorities.

10. High performance I/O

VMDirectPath—This product 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/UCNA port must be fully dedicated to a single VM, so VMDirectPath should be reserved for applications that require maximum I/O.

vNetwork Distributed Switch—Users can leverage VMware vNetwork Distributed Switch to dynamically adjust workloads across a team of physical adapters with robust Emulex OneConnect teaming capabilities that include assignment of NIC ports and NIC teams to multiple Virtual Local Area Networks (VLANs) and link aggregation of Ethernet traffic (FCoE traffic unteamed).

VMware NetQueue—This product is a multi-queue performance technology that significantly improves TCP/IP performance in 10GbE virtualized environments. The NIC function in OneConnect UCNAs supports VMware NetQueue, enabling improved server consolidation, reduced hypervisor overhead and an increase in the number of VMs that can run on virtualized servers.

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
Tokyo, Japan +81 3 5325 3261 | **Bangalore, India** +91 80 40156789

Connect with Emulex

twitter.com/emulex [friendfeed.com/emulex](https://www.facebook.com/emulex) bit.ly/emulexlinks [bit.ly/emulexfb](https://www.linkedin.com/company/emulex)



www.emulex.com

©2011 Emulex, Inc. All rights reserved. This document refers to various companies and products by their trade names. In most, if not all cases, their respective companies claim these designations as trademarks or registered trademarks. This information is provided for reference only. Although this information is believed to be accurate and reliable at the time of publication, Emulex assumes no responsibility for errors or omissions. Emulex reserves the right to make changes or corrections without notice. This report is the property of Emulex and may not be duplicated without permission from the Company.