



Getting Started with VMware vSphere and Emulex/Cisco SANs

vmware®



EMULEX®

Subscribe
to the SAN
Virtuosity Series at
www.sanvirtuosity.com



Table of Contents

Emulex and Cisco SAN Components	1
Hardware.....	1
Emulex LightPulse® HBAs and OneConnect™ Universal CNAs.....	1
Cisco MDS 9000 Family Switches	1
Storage Devices	2
Management Software.....	2
Emulex OneCommand Manager™	2
Cisco Device Manager and Fabric Manager	3
Array Management Software.....	3
The Virtual Data Center Optimized SAN.....	4
Virtual SANs.....	4
Role-Based Access Control	4
Fibre Channel Zoning	5
Virtual Machine Management.....	6
Storage Access Options	7
Creating Virtual Machines	7
Conclusion	8

Introduction

Server virtualization has become a mainstream technology for optimizing resources and reducing costs in the data center. By consolidating workloads from underutilized and often outdated servers, substantial savings can be gained in capital expenditures (CapEx) for equipment and operating expenditures (OpEx) for maintenance, power, and management resources. VMware vSphere is the market-leading virtualization solution that provides the performance and scalability needed to virtualize business-critical applications along with the infrastructure to support cloud computing.

To gain the full benefits of VMware vSphere, most organizations use a shared Storage Area Network (SAN) based on the Fibre Channel (FC) protocol. A SAN is a dedicated network that interconnects hosts and storage devices, primarily to exchange Small Computer System Interface (SCSI) traffic for a specific user community across physical links. A set of protocols run over the SAN to handle routing, naming, and zoning. Fibre Channel SANs support critical server virtualization technologies, including high availability and streamlined virtual machine (VM) migration, and deliver the performance required for higher virtualization ratios.

As industry leaders, Emulex® and Cisco® are committed to customer success with Fibre Channel SAN technologies. As part of that effort, we are collaborating on the SAN Virtuosity series, collection of white papers and webcasts designed to help IT managers fully achieve the synergy of server virtualization and shared storage. This document is the first in the series and provides basic information for deployment of Fibre Channel SANs with VMware vSphere. To subscribe to the series, visit sanvirtuosity.com.

Emulex and Cisco SAN Components

An Emulex and Cisco SAN consists of hardware and software components.

Hardware

Emulex LightPulse® HBAs and OneConnect Universal CNAs

Application servers use one or more adapters to connect to the SAN. Emulex is currently shipping the 8-Gbp/s LPe12000 family, the 8th generation of LightPulse Fibre Channel Host Bus Adapters (HBAs) and winner of the 2009 Storage Magazine Editor's Choice Award. In addition to doubling the bandwidth from the previous generation, performance tests using the Emulex LPe12000 with VMware vSphere 4 and ESX Server 4.0 have shown a 200 % increase in I/O operations per second (IOPS) when compared to VMware ESX Server 3.5.

Emulex has also introduced the OneConnect Universal Converged Network Adapter (CNA) product family that provides block storage connectivity over a 10Gb/s Ethernet (10GbE) network. OneConnect adapters include CNAs that support Fibre Channel over Ethernet (FCoE) that is based on IEEE standards and allows Fibre Channel frames to be encapsulated into Ethernet frames and transmitted over a converged Ethernet network. The OneConnect family also includes iSCSI adapters that provide full offload of Ethernet and iSCSI protocols for optimum server efficiency. OneConnect products provide significant cost savings by supporting LAN and SAN communication over a common network infrastructure.

Cisco MDS 9000 Family Switches

Interconnection between an application server and consolidated storage devices is performed across a network (also known as a fabric) of Fibre Channel switches specifically designed for reliable SCSI I/O communication.

The smaller SAN switches that incorporate one or more single points of failure are called fabric switches. Larger switches that provide high availability without a single point of failure are called directors.

The Cisco MDS 9000 Family of multilayer directors and fabric switches offers Fibre Channel switching services that offer high performance while helping ensure high reliability. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This combination enables highly available, scalable storage networks and provides intelligent networking features. The main features include:

- Integrated SCSI over IP (iSCSI) gateway to access Fibre Channel storage with iSCSI initiators
- Fibre Channel over Internet Protocol (FC-IP) integration to provide SAN extension over long distances
- Virtual SANs to enable flexible network resource provisioning with fault domain isolation
- N-Port ID Virtualization (NPIV)
- Advanced security
- Complete set of debugging and analysis tools
- Unified SAN management

Storage Devices

The storage devices connected to a SAN are storage arrays or tape libraries. Tape libraries are used to support backup operations on physical tape volumes. A storage array is an enclosure of multiple disk drives accessible through a controller. The controller presents disk storage capacity as volumes, which are identified by a LUN (Logical Unit Number). In most storage arrays, the storage controller can present volumes of greater capacity than individual disks and can offer services as different RAID (Redundant Array of Inexpensive Disks) levels to achieve performance or reliability objectives. Other services provided by storage arrays include the capability to make point-in-time copies for backup or data analysis and replication of volumes to another array for business continuity and disaster recovery.

Emulex adapters and Cisco switches are interoperable and fully qualified with a wide range of storage arrays and tape libraries from all the leading manufacturers.

Management Software

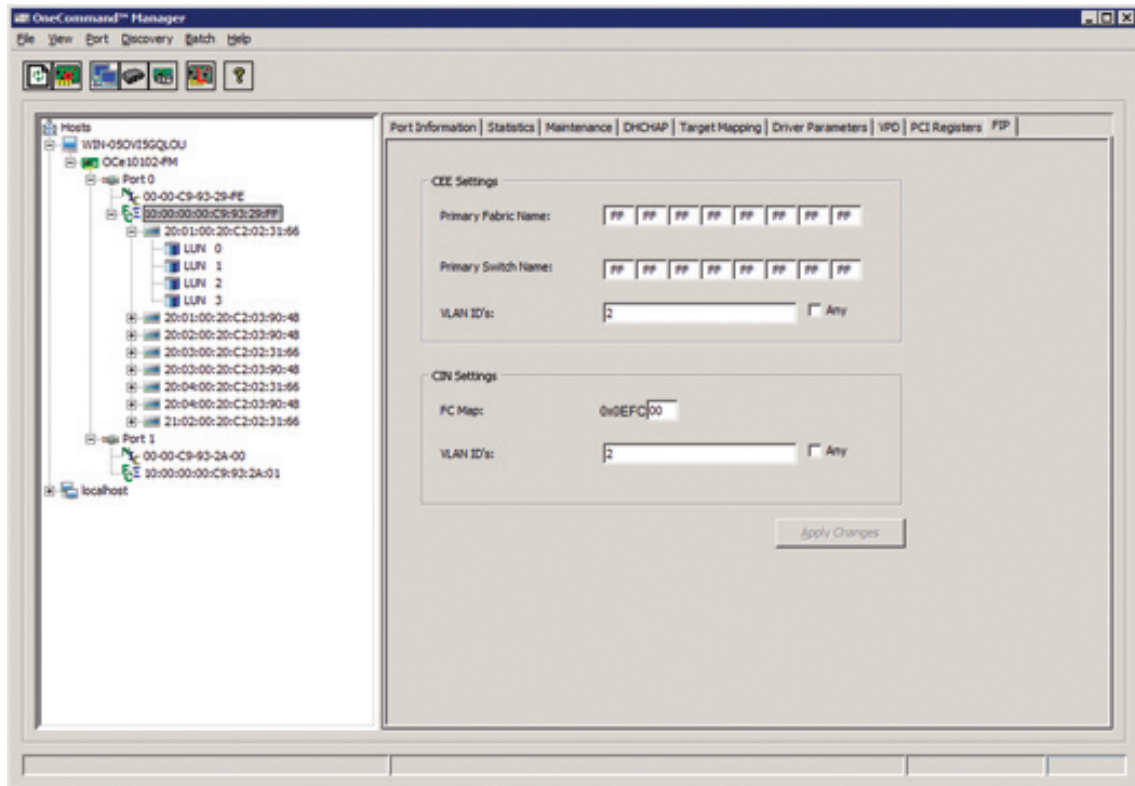
In addition to the physical deployment, a SAN requires logical management steps to be operational. Configuration includes establishing links between switches, creating groups, or zones, of devices with application servers and storage arrays that communicate with each other to establish the desired level of security.

Management is performed using a combination of applications from Emulex, Cisco, and storage array providers. A brief description is provided here, with functional details to be covered in topics throughout this series of documents.

Emulex OneCommand Manager

Emulex OneCommand Manager is a multi-protocol management application that is based on the Emulex HBAnyware management application for LightPulse HBAs. OneCommand Manager has both a Graphical User Interface (GUI) and Command-Line Interface (CLI) that support a wide variety of management functions (Figure 1).

Figure 1 Emulex OneCommand Manger GUI Display Showing Network and Storage Adapters



Key OneCommand Manager features include:

- Discovery of all SAN elements with a choice of views based on the host, fabric, and virtual ports (VPorts)
- Management of local and remote servers using the Fibre Channel or IP network
- Support for Emulex LightPulse HBAs and OneConnect FCoE CNAs, iSCSI adapters and Ethernet adapters
- Management modes to limit both local and remote management
- Batch download option to update firmware on multiple servers and adapters with a single command

Cisco Device Manager and Fabric Manager

The Cisco Device Manager and Fabric Manager management applications support secure Simple Network Management Protocol Version 3 (SNMPv3). Cisco Device Manager presents a switch view that allows point-and-click configuration of elements as interfaces and device-specific parameters.

Cisco Fabric Manager provides a GUI that displays real-time views of the network fabric and enables configuration management of Cisco MDS 9000 devices and third-party switches. For larger deployments, the optional Cisco Fabric Manager Server (FMS) performs advanced monitoring, troubleshooting and configuration for multiple fabrics with access by up to 16 Cisco Fabric Manager clients at a time.

Array Management Software

Each storage array vendor provides a management application to create and configure the volumes provided to the application server. Basic capabilities include features to:

- Create volumes from physical disks with the desired capacity and RAID level
- Expose volumes on one or more Fibre Channel interfaces
- Assign a LUN to each volume
- Open the volume to a specific server or group of servers (LUN mapping and masking)

LUN mapping and LUN masking are based on a World Wide Port Name (WWPN), a unique 16-byte value that identifies an HBA, CNA, or array port.

The Virtual Data Center Optimized SAN

Virtual SANs

The Cisco MDS 9000 uses virtual SANs (VSANs). With a VSAN, the network administrator can build a single topology containing switches, links and one or more VSANs. Each VSAN in this topology has the same behavior and property as a SAN, but the VSAN has the following additional features:

- Multiple VSANs can share the same physical fabric.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols, such as Fibre Channel routing, domain manager and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

A VSAN is the ideal container for all the servers and storage that provide VMs and VM mobility to the same user community or application group.

With a VSAN, VMware ESX servers belonging to the same VMware High Availability (HA) cluster and their associated storage devices can be deployed in the same VSAN using any port of any switch. This configuration maintains a tight isolation, but effectively shares the physical network infrastructure without the need for detailed preplanning.

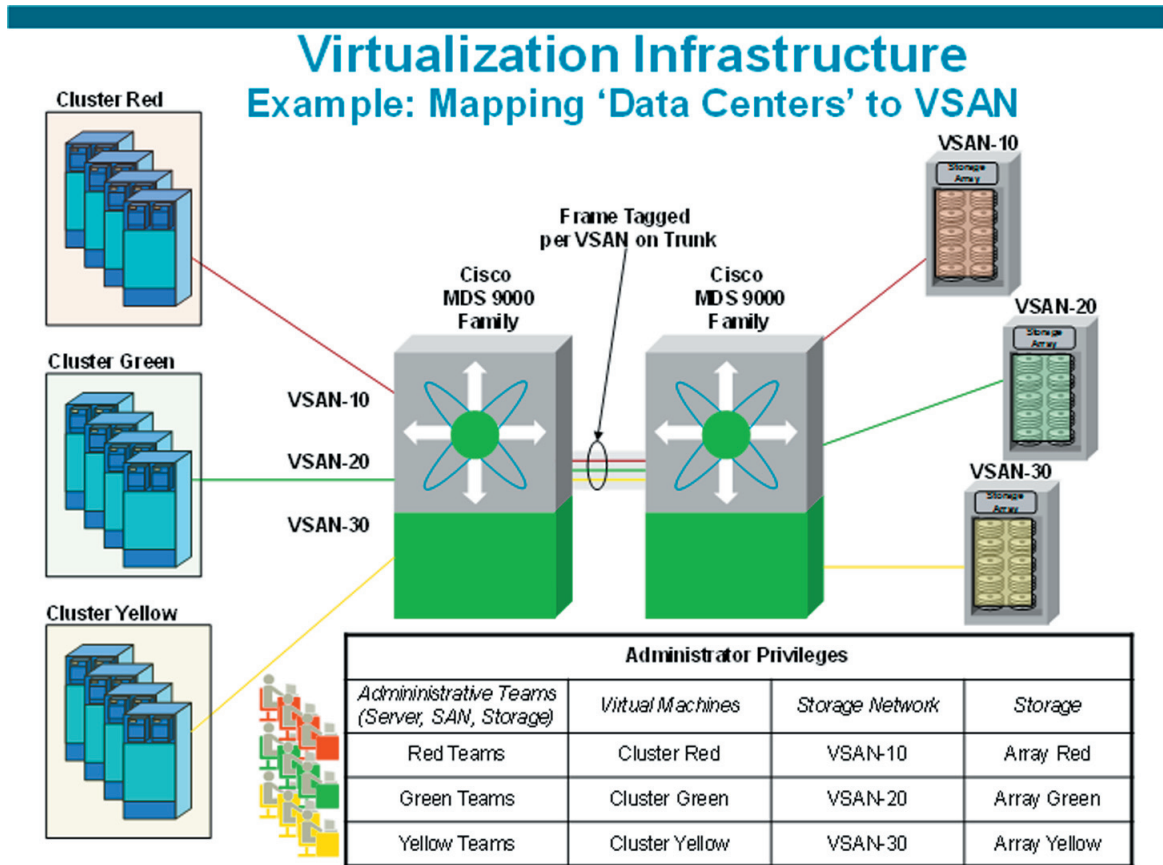
Adding or removing a physical server from a given cluster is achieved by changing the server's properties and reassigning the switch port to a different VSAN. This approach is functionally equivalent to relocating the Fibre Channel cable from one isolated physical fabric to another, but is totally virtualized. Since any fabric port can belong to any VSAN, no advance planning is needed to specify how servers will be physically associated in the fabric. The assignment can be performed on demand using software.

Role-Based Access Control

VMware vSphere management offers several levels of role-based access control (RBAC). For example, an administrator can be in charge of one or more clusters or one or more physical servers.

The Cisco MDS 9000 management architecture offers an equally sophisticated RBAC infrastructure, enabling administrators to assign VSAN administrative rights to groups or individuals in charge of a cluster of physical servers (Figure 2).

Figure 2 Administrator Privileges Assigned with RBAC



Fibre Channel Zoning

Fibre Channel zoning is the standard mechanism for grouping and controlling access between servers (hosts or initiators) and array ports (targets). Storage adapter ports and array controller ports are identified by WWPNs and must be in the same zone for servers to access storage. Zoning is used with LUN mapping and masking to control data availability and help ensure data integrity.

The zones are contained in a VSAN, with each VSAN containing one individual active zone set at a time, including the zones that are needed to provide required storage services. A zoning configuration error in a VSAN does not affect operation of other VSANs.

A zone consists of multiple zone members, and only members in the zone can access each other. If zoning is not activated, all devices are members of the default zone. If zoning is activated, any device that is not in an active zone is a member of the default zone. Members of the default zone may or may not communicate with each other, with the best practice being to deny communication.

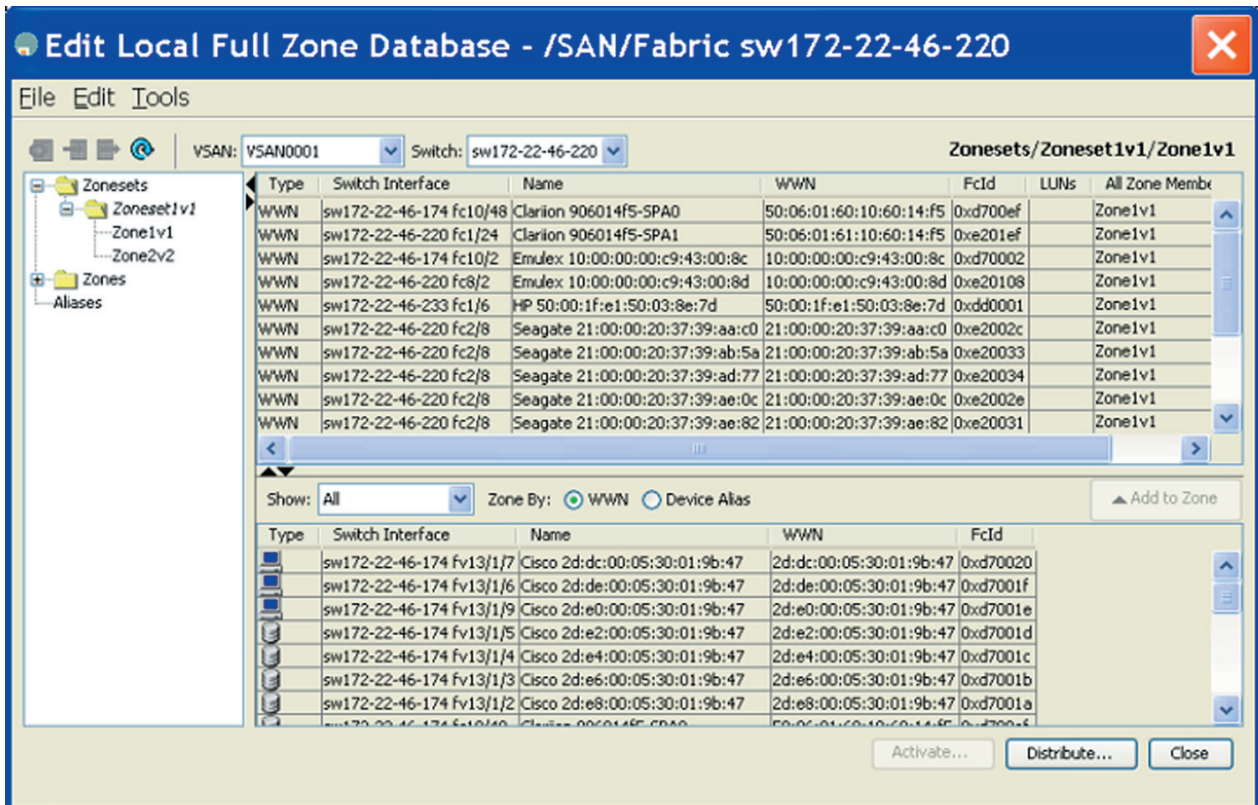
All servers in the same VMware HA cluster must be part of the same zone with a common target. They must be able to read and write to the same LUNs to access VM datastores.

Zoning can be administered from any switch in the fabric. When you activate a zone, all switches in the fabric receive the active zone set. If a new switch is added to an existing fabric, zone sets are acquired by the new switch. New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices. The best practice is to avoid configuring a large number of targets or a large number of initiators in a single zone, keeping the zone scope to servers and storage devices that are part of the same cluster.

For small deployments, the Quick Config Wizard can be used to add or remove zone members in a VSAN. The Quick Config Wizard is supported with the Cisco MDS 9124 Fabric Switch, Cisco MDS 9134 Fabric Switch, Cisco Fabric Switch for the HP c-Class BladeSystem and Cisco Fabric Switch for the IBM BladeCenter. The Quick Config Wizard can be used only with standalone switches that have no existing zoning defined on the switch.

Cisco Fabric Manager’s Edit Local Full Zone Database tool can be used to manage zones across multiple switches. All zoning features are available through the Edit Local Full Zone Database dialog box (Figure 3).

Figure 3 Managing Zones across Multiple Switches with Cisco Fabric Manager



Virtual Machine Management

This document describes VM management based on VMware vCenter Server, which provides centralized management of VMware vSphere environments. This approach is consistent with that of Emulex and Cisco management tools that provide management for adapters and switches throughout the network from a single management console.

Storage Access Options

VMware ESX Server offers two options for managing block-level disk access:

- VMware Virtual Machine File System (VMFS)
- Raw Device Mapping (RDM)

VMFS is a virtual, clustered file system that allows multiple VMware ESX servers to access the same storage. With this architecture, multiple virtual VMs on multiple VMware ESX servers can share the same datastore and associated LUN. On-disk locking is used to help ensure that a VM can run on only one server at a time. With multiple VMs using a common LUN, the effort required for LUN management is reduced. VMware recommends VMFS datastores for most virtual disk storage.

RDM maps a specific LUN to a single VM by using a mapping file inside VMFS that acts as a proxy for a raw device, allowing direct block-level access from the VM to the LUN. The mapping file is presented to the management software as an ordinary disk file available for the usual file system operations.

The two file systems can be used in parallel on the same VMware ESX server, and both are supported with critical VMware vSphere features such as VMware HA, VMware vMotion, and VMware Distributed Resources Scheduler (DRS).

Creating Virtual Machines

Use the following basic steps to prepare SAN-based storage and create a VM:

1. Create a LUN.

The requirements of the LUN should be carefully evaluated prior to its creation. For example, the RAID level for the LUN will affect both performance and reliability, and the choice of disk drives and cache allocation will affect performance. With VMFS, the LUN likely will be shared by multiple VMs, so the LUN size and properties will affect multiple applications.

LUNs are created by the storage administration team through the storage array management software, and fabric zoning and LUN masking can be used to isolate the storage as appropriate.

2. Create a datastore.

Typically, the virtual infrastructure administrator uses VMware vCenter to create the datastore, using any available LUN that is unassigned.

3. Create a VM.

VMware vCenter is also used to create VMs. VMs can be added to a newly created datastore as described in the previous step or to an existing datastore that already includes VMs. The exception is creation of a VM that uses RDM, which requires addition of the raw disk map file in the new or preexisting datastore plus assignment of the raw LUN to the VM.

Conclusion

The combination of VMware vSphere and Fibre Channel SANs provides substantial CapEx and OpEx savings. Cisco and Emulex are leaders in Fibre Channel technologies and provide scalable, high-performance hardware and software components that fully enable server virtualization in the data center. This document explains some of the components and concepts involved in deploying VMware vSphere on a SAN in conjunction with the Cisco MDS 9000 and Emulex HBAs and Universal CNAs.

Cisco Systems, Inc. 170 West Tasman Drive, San Jose, CA 95134-1706 USA Tel 408-526-4000 www.cisco.com

©2010 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.

Emulex Corporation. 3333 Susan Street, Costa Mesa, CA 92626 USA Tel 714-662-5600 www.emulex.com

Copyright © 2010 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

Information furnished by Emulex is believed to be accurate and reliable. However, no responsibility is assumed by Emulex for its use; or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or related rights of Emulex. Emulex, the Emulex logo, LightPulse, OneConnect, OneCommand and SLI are trademarks of Emulex.

VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-427-5000 www.vmware.com

© 2010 VMware, Inc. All rights reserved. VMware, the VMware “boxes” logo and design, vSphere, Virtual SMP and vMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.