

BUSINESS READY CONFIGURATIONS FOR DATA CENTER VIRTUALIZATION

A SOLUTION GUIDE

FOR DELL™ POWEREDGE™ BLADE
SERVERS, DELL EQUALLOGIC™ STORAGE,
AND VMWARE vSPHERE™

July 2009

Dell Virtualization Solutions Engineering
<http://www.dell.com/virtualization/businessready>



Information in this document is subject to change without notice.

© Copyright 2009 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

This white paper is for informational purposes only and may contain typographical errors or technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

Dell, the *DELL* logo, *EqualLogic*, *PowerEdge*, *PowerConnect*, and *OpenManage* are trademarks of Dell Inc. *Microsoft* is a registered trademark of Microsoft Corporation; *VMware*, *vSphere*, *vCenter*, and *VMotion* are registered trademarks or trademarks (the "Marks") of VMware, Inc. in the United States and/or other jurisdictions. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Contents

1	Introduction	5
2	Audience and Scope	5
3	Overview	6
3.1	VMware vSphere 4	6
3.2	Dell PowerEdge Blade Server.....	7
3.3	Dell EqualLogic PS6000 Series iSCSI Arrays.....	7
3.4	Why Dell blades + VMware + Dell EqualLogic?	8
4	Specification	9
4.1	How to order	10
4.2	Dell Global Services	10
5	Reference Architecture	11
6	Network Architecture	12
6.1	LAN Architecture.....	13
6.1.1	Traffic Isolation using VLANs	14
6.1.2	Load Balancing and Failover.....	14
6.1.3	Stacking	15
6.1.4	Uplinks.....	15
6.2	iSCSI SAN Architecture	16
6.2.1	Storage Connectivity	16
6.2.2	VLAN	17
6.2.3	Load Balancing	17
6.2.4	Stacking	18
6.2.5	Uplinks.....	18
6.2.6	EqualLogic Network Requirements.....	19
6.3	Scalability.....	19
6.3.1	Scaling Compute Resources	19
6.3.2	Scaling Storage Resources.....	20
7	Storage Configuration.....	21
7.1	Volume Size Considerations	21
7.2	Array RAID Configuration	21
8	Management	22
8.1	Virtualization Management.....	22
8.2	Systems Management.....	22
8.3	Storage Management.....	23
9	Deploying and Configuring the Solution	23
9.1	Rack, Stack, and Cable.....	23
9.2	Initial Setup.....	23

10	References	26
11	Appendix A - Deployment Scripts.....	27

1 Introduction

Business Ready Configurations for Virtualization are pre-configured, easy to order, bundled solutions that include Dell PowerEdge™ blade servers, Dell EqualLogic™ SAN, and VMware vSphere™. This white paper provides a detailed reference architecture for the Business Ready Configurations. Based on extensive engineering work in architectural design and certification, customers can quickly and confidently deploy these engineering proven architectures into production environments, thereby helping to eliminate much of the costly and time consuming trial and error often encountered during complex infrastructure design and implementation. The solution is optimally configured to run virtualized workloads and is designed to provide redundancy with no single point-of-failure, as well as scalability and manageability. It includes the network architectures, storage configurations, and best practices necessary for deploying and configuring the solution. The solution will help customers achieve the full benefits of VMware Infrastructure, Dell PowerEdge blade servers, and Dell EqualLogic PS Series arrays.

2 Audience and Scope

The intended audience for this white paper is IT administrators, IT managers, and channel partners who are planning to deploy VMware virtualization using Dell blade servers and Dell EqualLogic SAN. This white paper provides an overview of the recommended servers, storage, software, and services. It can be used to plan, scope, and procure the required components to set up a virtualization infrastructure.

This white paper provides reference architecture and best practices for deploying and configuring Dell's 11th generation blade server – PowerEdge M610, Dell EqualLogic PS6000 Series arrays, and VMware vSphere 4. The solution uses iSCSI as the storage environment. Although supported, fibre channel is not discussed in this white paper. Based on customer requirements, further customization of the recommended architecture may be required. VMware Fault Tolerance feature is not discussed in this white paper. For reference architecture on high availability solution including VMware Fault Tolerance, refer to www.dell.com/virtualization.

3 Overview

This section (Figure 1) provides a high-level product overview of the VMware vSphere, Dell PowerEdge blade servers, and Dell EqualLogic PS Series arrays.

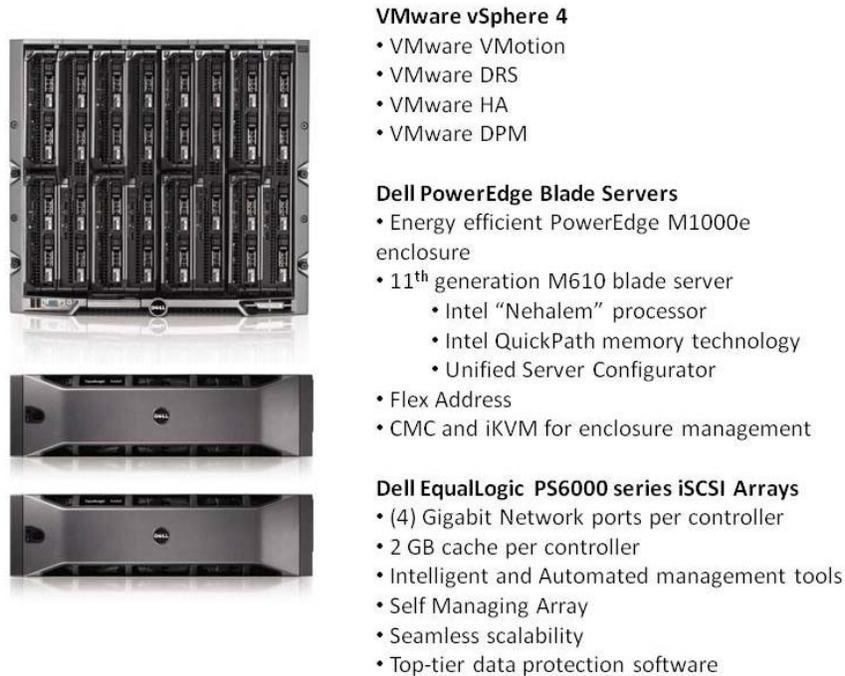


Figure 1: Overview of Dell Blade Server, Dell EqualLogic SAN, and VMware Infrastructure

3.1 VMware vSphere 4

VMware vSphere 4 comprises of VMware ESX/ESXi and vCenter Server, and includes the following capabilities:

VMware VMotion: VMware VMotion technology provides real-time migration of running virtual machines from one server to another with no disruption or downtime. VMotion can be used to achieve manual load balancing and zero downtime maintenance.

VMware Distributed Resource Scheduler (DRS): VMware DRS technology enables VMotion to automatically achieve load balancing according to resource requirements.

VMware Distributed Power Management (DPM): When virtual machines in a DRS cluster need fewer resources, such as during nights and weekends, DPM consolidates workloads onto fewer servers and powers off the rest to reduce power consumption. When virtual machine resource requirements increase (such as when users log into applications in the morning), DPM brings powered-down hosts back online to ensure service levels are met.

VMware High Availability (HA): VMware HA provides high availability at the virtual machine level. Upon server failure, VMware HA automatically re-starts virtual machines on other physical servers.

VMware vNetwork Distributed Switch: VMware’s vNetwork Distributed Switch spans many ESX hosts and aggregates networking to a centralized cluster level. vNetwork Distributed Switch abstracts configuration of individual virtual switches and enables centralized provisioning, administration and monitoring through VMware vCenter Server.

VMware Update Manager: VMware Update Manager automates patch management; enforcing compliance to patch standards for VMware ESX hosts, as well as Microsoft or Linux virtual machines. Secure patching of offline

virtual machines, as well as automated snapshots and rollback capabilities, reduce risk in the environment. Integration with DRS allows for non-disruptive VMware ESX patching with little-to-no administrative overhead.

VMware Storage VMotion: VMware Storage VMotion enables real-time migration of running virtual machine disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.

Host Profiles: Host Profiles standardize and simplify the deployment and management of VMware ESX host configurations. It captures and stores validated configuration including networking, storage and security settings. It also makes it easy to monitor host compliance.

For more information on VMware vSphere, see <http://www.vmware.com/products/vsphere/>.

3.2 Dell PowerEdge Blade Server

Blade Modular Enclosure: The Dell PowerEdge M1000e is a high-density, energy-efficient blade chassis that supports up to sixteen half-height blade servers, or eight full-height blade servers, and six I/O modules for the blade servers. A high-speed passive mid-plane connects the server modules to the I/O modules, management, and power in the rear of the enclosure. The enclosure includes a flip-out LCD screen for local configuration. The enclosure also includes six hot-pluggable and redundant power supplies and nine hot-pluggable N+1 redundant fan modules.

Blade Servers: The Dell PowerEdge M1000e supports the recently announced 11th generation Dell PowerEdge M610 and M710 blade servers based on the new Intel® *Nehalem* processors. The Dell PowerEdge M610 supports 12 DIMM slots and the Dell PowerEdge M710 supports 18 DIMM slots. The blade servers use Intel QuickPath technology to provide a high-speed link to the memory modules. The 11th generation blade servers come with the next generation system management tool: Unified Server Configurator (USC). This helps customers reduce operating costs by simplifying deployment and management. USC supports diagnostics, firmware updates, and hardware configuration. The enclosure also supports the AMD-based M605, M805, and M905 blade servers. The Intel-based M600 legacy blade server is also supported.

I/O Modules: The enclosure provides three redundant fabrics using six I/O modules. The modules can be populated with Cisco Catalyst blade switches, Dell PowerConnect™ switches, and pass-through modules. Fibre Channel and InfiniBand® switch modules are also supported.

Management: The Dell PowerEdge M1000e has integrated management through a redundant Chassis Management Controller (CMC) module for enclosure management and integrated keyboard, video, and mouse (iKVM) modules. Through CMC, the enclosure supports FlexAddress technology. FlexAddress technology allows any M-Series blade enclosure to lock the World Wide Name (WWN) of the Fibre Channel controller and Media Access Control (MAC) of the Ethernet controller into a blade slot. This enables seamless swapping or upgrade of blade servers with Ethernet and Fibre Channel controllers without affecting the LAN/SAN.

For more information on Dell blade servers, see www.dell.com/blades.

3.3 Dell EqualLogic PS6000 Series iSCSI Arrays

The Dell EqualLogic PS6000 is the latest iSCSI storage device from Dell. Its features include four network ports per controller, faster processors, 2 GB cache per controller, support for RAID 6, increased drive capacity, and a new monitoring application, SAN HQ, at no additional cost.

In addition to the new features described above, Dell EqualLogic SAN devices provide the following capabilities:

Reliability: Dell EqualLogic PS6000 Series arrays have hot-swappable redundant components, a choice of RAID types, and hot-spare disks. They also include the Auto-Stat Disk Monitoring System (ADMS) which proactively scans disk drives in the background to help detect media anomalies and correct them.

Scalability: As each array is added to the storage group, the storage capacity and performance, in terms of both bandwidth and IOPS, are increased. This increased capacity can be utilized without downtime. Thin-provisioning permits predefinition of a group of VMFS volumes which are larger than the physical space, allowing for physical space to be added and utilized without downtime when necessary.

Self-Managing Arrays: The arrays offer many self-managing features such as automatic load balancing and storage tiering. A single storage pool can have different models that offer a range of capacity and performance parameters. In addition, different arrays in a storage pool can be configured with different RAID levels, and volumes will automatically be migrated between the RAID levels based on performance data and usage patterns. All data and volume movement can be performed online with zero downtime.

Top-Tier Data Protection Software: Advanced data protection features such as Auto Replication and Auto-Snapshot Manager come standard. The Auto-Snapshot Manager integrates with VMware vCenter and VMware's native snapshot technology to provide intelligent and consistent snapshots.

VMware Site Recovery Manager (SRM) is tightly integrated with the Dell EqualLogic PS Series to offer centralized disaster recovery management, automation, and testing for the entire virtualized datacenter using the PS Series' native Auto-Replication feature integrated directly into Site Recovery Manager using specialized storage adapter software.

For more information on Dell EqualLogic storage, see www.dell.com/equallogic.

3.4 Why Dell blades + VMware + Dell EqualLogic?

VMware vSphere on Dell blade servers and Dell EqualLogic SAN provide a compelling infrastructure that meets the demands of today's datacenters. The combined solution offers customers rapid deployment, optimized resource utilization, simplified management, seamless scalability, and energy efficiency with no compromise in performance.

Rapid Deployment

The solution offers several tools that help customers go from an out-of-the-box environment to a production environment quickly and seamlessly. Dell's USC provides a single interface to configure the server BIOS, network controller, and storage controller. USC also supports VMware ESX installations. If customers opt for VMware ESXi, Dell ships servers with integrated ESXi on an embedded SD card. Dell EqualLogic SAN provides a highly simplified deployment process. Dell EqualLogic PS Series arrays go from out-of-the-box to production deployment in under an hour with an intuitive, wizard-based GUI to integrate easily into the existing infrastructure. Using virtualization capabilities, new VMs can be created or existing VMs can be cloned or deployed from templates with just a few clicks.

Optimized Resource Utilization

In addition to increased resource utilization through server consolidation using VMware ESX/ESXi, VMware VMotion and DRS technologies dynamically distribute workloads in all servers in the cluster, ensuring each server's computing resources are optimally used. The Dell EqualLogic SAN complements this by optimally using the storage devices. The virtual machine data is intelligently distributed between all the members (storage devices) in the storage pool, ensuring that all the storage controllers and network interfaces are optimally used. This distribution is done automatically and is completely transparent from the user.

Simplified Management

One of the key benefits of virtualization is simplified management. VMware ESXi software comes integrated with Dell OpenManage™ management software, providing customers with a simplified systems management interface. The Dell EqualLogic PS Series arrays provide a single-pane management view of the entire virtualized storage pool in the SAN. The Dell EqualLogic PS6000 Series arrays intelligently balance workloads across the available arrays with no manual intervention.

Seamless Scalability

VMware virtualization and Dell servers and storage provide a highly scalable solution to meet the growing IT demands of today's datacenters. New server resources can be easily added to VMware DRS clusters. Virtual Machines will automatically migrate to utilize the newly added resources. Similarly, new storage devices can be

easily added to Dell EqualLogic storage pools. Dell EqualLogic will automatically migrate the virtual machine data to the newly added storage resource with no downtime.

Energy Efficiency

The Dell PowerEdge M1000e enclosure takes advantage of thermal design efficiencies such as ultra-efficient power supplies and dynamic power-efficient fans with optimized airflow design to efficiently cool the chassis and enable better performance in a lower power envelope.

Consolidating workloads running on older servers on the new energy-efficient Dell blade servers greatly reduces power consumption. This capability, combined with VMware’s DPM technology substantially improves the energy efficiency of the solution.

4 Specification

In the following table, we provide the specification of three pre-configured solutions. In the rest of the document, we will discuss reference architectures and best practices for deploying and using these solutions.

Table 1 Solution Specification

	Entry Configuration	Medium Configuration	Large Configuration
Solution Summary			
Solution ID	741914.1	741933.1	742026.1
Blade Chassis	(1) M1000e with (6) Cisco M3130G I/O modules	(1) M1000e with (6) Cisco M3130G I/O modules	(1) M1000e with (6) Cisco M3130X I/O modules
ESX Blade Server	(4) x M610 with 24GB memory in each blade	(6) x M610 with 48GB memory in each blade	(10) x M610 with 48GB memory in each blade
Storage Device	(1) PS6000X	(2) PS6000XV	(2) PS6000XV
Management Blade Server	(1) x M610 with VMware vCenter Server	(1) x M610 with VMware vCenter Server	(1) x M610 with VMware vCenter Server
Chassis Configuration			
I/O module for A1	Cisco M 3130G	Cisco M 3130G	Cisco M 3130X with (2) X2-CX4 10Gb Module
I/O module for B1	Cisco M 3130G with (2) 1000BAS-T SFP modules	Cisco M 3130G with (4) 1000BAS-T SFP modules	Cisco M 3130G with (4) 1000BAS-T SFP modules
I/O module for C1	Cisco M 3130G	Cisco M 3130G	Cisco M 3130X with (2) X2-CX4 10Gb Module
I/O module for A2	Cisco M 3130G	Cisco M 3130G	Cisco M 3130X with (2) X2-CX4 10Gb Module
I/O module for B2	Cisco M 3130G with (2) 1000BAS-T SFP modules	Cisco M 3130G with (4) 1000BAS-T SFP modules	Cisco M 3130G with (4) 1000BAS-T SFP modules
I/O module for C2	Cisco M 3130G	Cisco M 3130G	Cisco M 3130X with (2) X2-CX4 10Gb Module
Stacking cables	6	6	6
Management	Redundant Chassis Management Controllers (CMC)	Redundant Chassis Management Controllers (CMC)	Redundant Chassis Management Controllers (CMC)
KVM	Integrated Avocent keyboard, video, and mouse (iKVM) switch	Integrated Avocent keyboard, video, and mouse (iKVM) switch	Integrated Avocent keyboard, video, and mouse (iKVM) switch
ESX Blade Server Configuration			
Blade Server Model	M610	M610	M610
Processor	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache
Memory	24 GB (6 x 4 GB, DDR3)	48 GB (12 x 4 GB, DDR3)	48 GB (12 x 4 GB, DDR3)

Add-in Controllers	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades
Local storage and controller	Diskless configuration	Diskless configuration	Diskless configuration
Hypervisor	VMware ESXi 4.0 - Enterprise Plus	VMware ESXi 4.0 - Enterprise Plus	VMware ESXi 4.0 - Enterprise Plus
Storage Configuration			
Storage Device	(1) PS6000X	(2) PS6000XV	(2) PS6000XV
Drives	16 X 400GB, 10K SAS	16 X 450GB, 15K SAS	16 X 450GB, 15K SAS
Storage Capacity	6.4 Terabyte capacity	7.2 Terabyte per device	7.2 Terabyte per device
Management Blade Server			
Blade Server Model	(1) M610	(1) M610	(1) M610
Processor	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache
Memory	8 GB	8 GB	8 GB
Network Controller	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades	(2) Broadcom 5709 Dual Port GbE I/O Card for M-Series Blades
Virtualization Management	VMware vCenter 4.0, 3Yr Upgrade Subscription	VMware vCenter 4.0, 3Yr Upgrade Subscription	VMware vCenter 4.0, 3Yr Upgrade Subscription
Software and Services			
Additional Software	PlateSpin Migrate for Dell (P2V 5-Pack), License & 1 Year Maintenance	PlateSpin Migrate for Dell (P2V 5-Pack), License & 1 Year Maintenance	PlateSpin Migrate for Dell (P2V 5-Pack), License & 1 Year Maintenance
Services	<ul style="list-style-type: none"> 3 Year ProSupport for IT and Mission Critical 4HR 7x24 Onsite Pack EqualLogics PS Array Installation M1000e Blade Chassis On-site installation service 	<ul style="list-style-type: none"> 3 Year ProSupport for IT and Mission Critical 4HR 7x24 Onsite Pack EqualLogics PS Array Installation M1000e Blade Chassis On-site installation service 	<ul style="list-style-type: none"> 3 Year ProSupport for IT and Mission Critical 4HR 7x24 Onsite Pack EqualLogics PS Array Installation M1000e Blade Chassis On-site installation service

Note: ESXi 4.0 Patch 1 or later is required in order to be a supported configuration.

4.1 How to order

To order the solution discussed in this paper, contact your Dell Sales Representative with the *Solution ID* provided above. Customers can order any of three pre-configured solutions. Customization can be made to the configurations to fit specific needs. Based on the customization, some of the best practices discussed in this white paper may not apply.

The solution can also be directly ordered online with a few clicks. Here is the link to Dell Business Ready Configurations for Virtualization: <http://www.dell.com/virtualization/businessready>. Customization can also be made to the configurations using the Dell online ordering tool.

4.2 Dell Global Services

Dell Global Services can also be engaged to order the solution. Dell Global Services helps customers find the “better path” to virtualization to reduce total cost of ownership, speed time to ROI, increase agility, and reclaim IT resources. Today’s financial environment dictates that IT reduce its budget while still providing ever increasing infrastructure services. Dell Global Services believes that by heavily leveraging server virtualization, and increasing the rate of adoption of this technology, you can accomplish this task.

Dell Global Services virtualization consultants work with customers to design and plan around the most common bottlenecks in virtualization implementations. Our methodology, tools, and processes are designed to speed up the implementation, ease migration scheduling, automate reporting, and provide transparency to the entire process. Dell Global Services also has the ability to offer end-to-end solutions with a single point-of-contact for hardware, software, services, and on-going support. In a time when many providers aspire to do everything, we at Dell Global Services are focused on IT infrastructure services excellence.

To engage Dell Global Services, see

http://www.dell.com/content/topics/global.aspx/services/adi/virtualization_new?c=us&cs=555&l=en&s=biz.

5 Reference Architecture

This section describes the high-level reference architecture for the solution.

Design Principles

The following principles were used during the design of the architecture.

1. **Optimal hardware configuration for virtualized workloads:** The solution is designed with an optimal hardware configuration to support virtualized workloads. Each blade server is configured with sufficient memory and network adapters required for virtualization.
2. **Redundancy with no single point-of-failure:** Redundancy is incorporated in every aspect of the solution, including networking and storage.
3. **Scalability:** The solution is highly scalable and is based on the architecture of the M1000e chassis and the EqualLogic PS6000 Series arrays. Guidelines to deploy additional blade servers and storage to increase the compute and storage capacity respectively are included.
4. **Isolated, redundant, and high-performance network design:** The network architecture is designed to support isolation of various traffic types required in a virtualized environment. It is designed to have no single point-of-failure and have optimal performance through NIC teaming and load balancing.

The solution consists of a Dell M1000e chassis populated with M610 blade servers running embedded VMware ESXi. One blade is used for management (VMware vCenter) and the rest are a part of ESX Cluster. Figure 2 provides a high-level reference architecture for the solution. This is based on the large configuration described in the Specification section. Small and Medium configurations have similar reference architecture with varying number of blades and storage arrays.

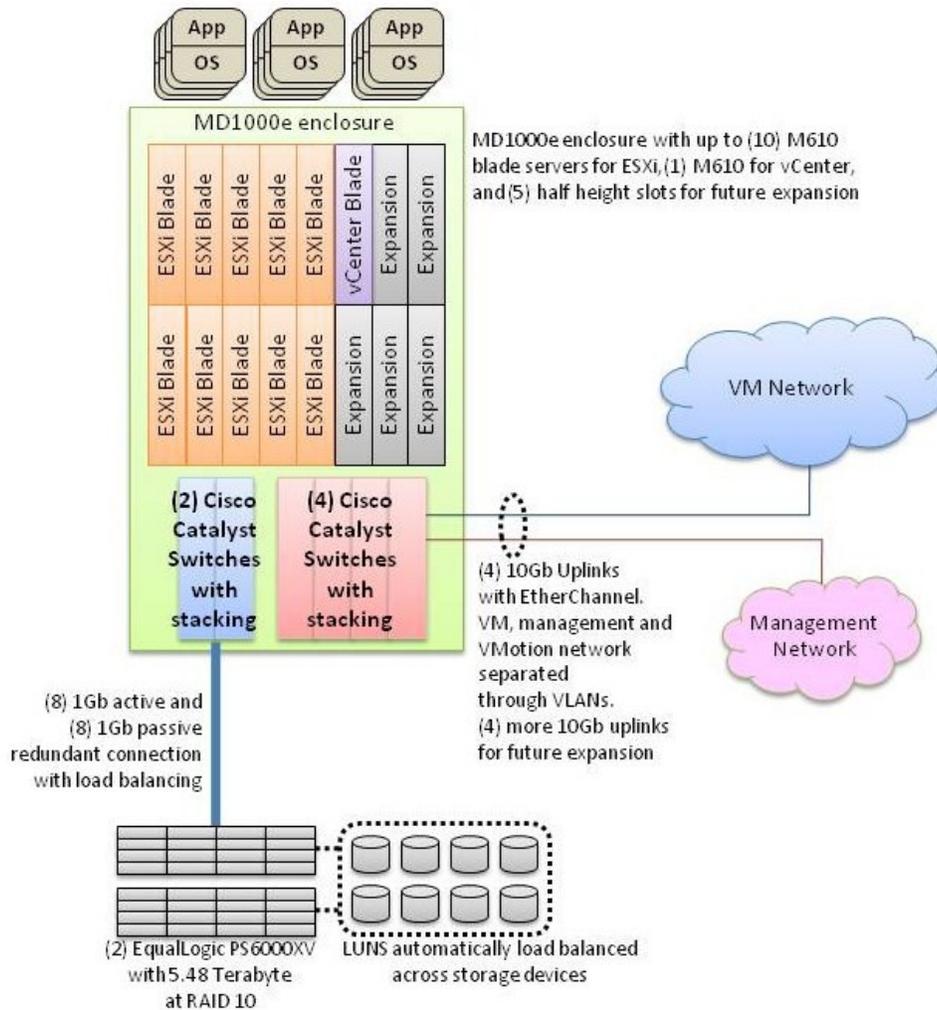


Figure 2: Reference Architecture for a Large Configuration

6 Network Architecture

The Dell blade chassis has three separate fabrics; referred to as A, B, and C¹. Each fabric has two I/O modules, making it a total of six I/O modules in the chassis. The I/O modules are A1, A2, B1, B2, C1, and C2. Each I/O module can be an Ethernet physical switch or an Ethernet pass-through module². In this solution, each I/O module contains a Cisco Catalyst 3130G or 3130X blade switch module.

VMware vSphere is typically comprised of four types of traffic: virtual machine traffic, management traffic, VMotion traffic, and iSCSI traffic. To support different traffic types in a virtualized environment, create two networks (see Figure 3). The first network is a Local Area Network (LAN), which includes management, virtual machine, and VMotion traffic. It is created by combining I/O modules A1, A2, C1, and C2 using high-speed stacking ports. The second network is dedicated to iSCSI SAN. It is created by combining I/O modules B1 and B2 using high-speed stacking ports³.

¹ <http://support.dell.com/support/edocs/systems/pem600/en/hom/HTML/about.htm#wp1212100>

² Fabrics B and C support Fibre Channel and InfiniBand, which are outside the scope of this white paper.

³ There is no special reason for choosing Fabric B for iSCSI network. Fabric C can also be chosen for iSCSI network, with Fabric A and B for LAN.

Design Principles for Network Architecture: The following design principles are used to develop the network architecture to support the above traffic types:

- **Redundancy:** Both LAN and iSCSI SAN have redundant I/O modules. Redundancy of the network adapters is achieved through NIC teaming at the virtual switch. There is no single point-of-failure.
- **Simplified management through stacking:** Switches servicing the same traffic type is combined into logical fabrics using the high-speed stacking ports on the switches.
- **iSCSI SAN physical isolation:** The iSCSI SAN should be physically separated from the LAN to avoid any network congestion.
- **Logical isolation of VMotion using VLAN:** VMotion traffic is unencrypted. It is important to logically isolate the VMotion traffic using VLANs.
- **Optimal performance:** Load balancing is used to achieve the highest throughput possible.

Based on these principles, we describe how we configure the solution to enable the various traffic types. Best practices for load balancing, failover, and VLAN configuration are provided.

This section assumes a certain understanding of networking concepts in a virtualized environment. See http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf for an overview of VMware's networking capabilities.

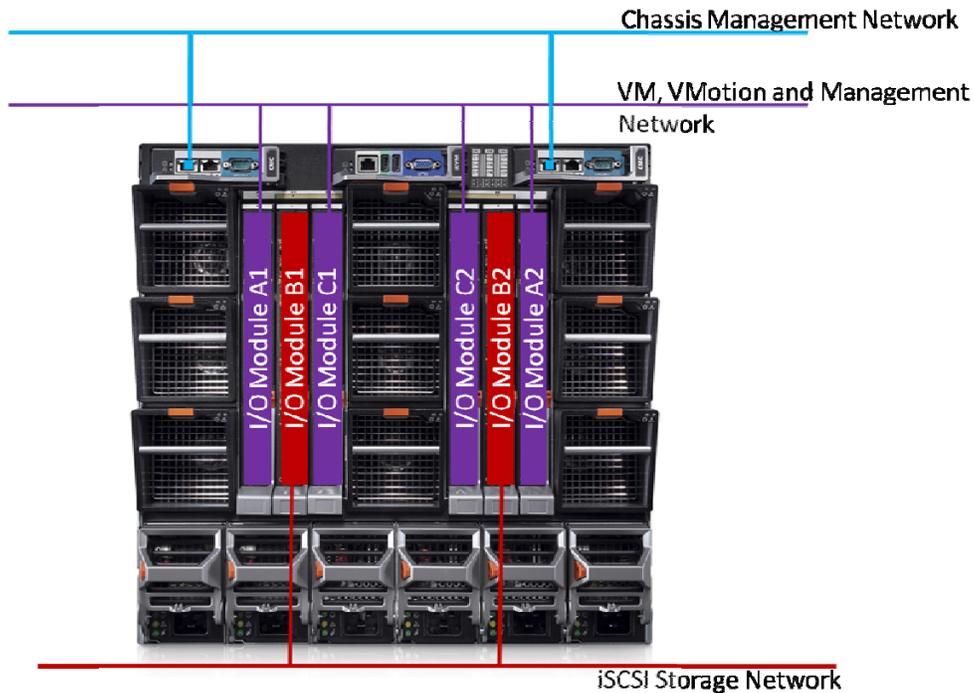


Figure 3: Overview of Network Configuration

6.1 LAN Architecture

The LAN comprises of virtual machine traffic, management traffic, and VMotion traffic. A virtual switch (vSwitch0) is created and four 1 Gb physical network adapters are connected as uplinks to this virtual switch. This creates a team of four network adapters, enabling NIC failover and load balancing. These network adapters are in turn connected to Cisco Catalyst 3130 I/O modules in slots A1, A2, C1, and C2. Three port groups are created on vSwitch0: one for the virtual machine network, another for the VMotion network, and the third for the service console (for ESX) or management (for ESXi). When virtual machines are created, the virtual network adapters are connected to the virtual machine network port group.

Figure 4 illustrates the virtual switch configuration with port groups, and how the virtual switch is connected to physical network adapters and in turn to the I/O modules.

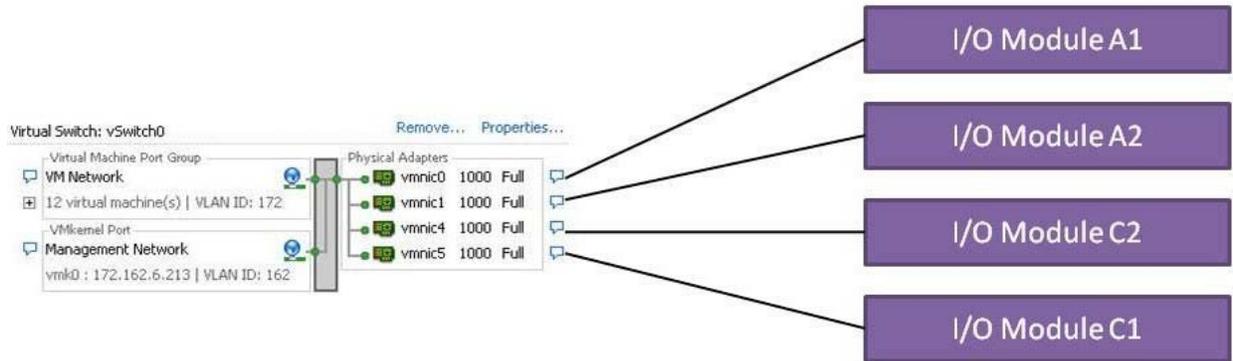


Figure 4: LAN Configuration

6.1.1 Traffic Isolation using VLANs

The traffic on the LAN is separated into three VLANs; one VLAN each for management, VMotion, and virtual machine traffic. Network traffic is tagged with the respective VLAN ID for each traffic type in the virtual switch. This is achieved through the Virtual Switch Tagging (VST) mode. In this mode, a VLAN is assigned to each of the three port groups. The virtual switch port group tags all outbound frames and removes tags for all inbound frames. Routing between VLANs is dependent on the specific customer requirements and is not included in the reference architecture. If desired, the Cisco Catalyst 3130 I/O modules can be easily configured to provide the routing.

Trunking must be used so that all the VLANs can share the same physical connection. To achieve this, all the internal ports in the Cisco I/O modules are configured to be in the trunk mode.

VMotion traffic is unencrypted, and so it is highly recommended to isolate VMotion traffic⁴. Using the above VLAN configuration, we achieve traffic isolation between various traffic types, including the VMotion traffic. The four network adapters provide sufficient bandwidth for all the traffic types.

6.1.2 Load Balancing and Failover

Load balancing enables sharing of network traffic between the network adapters in a team, thereby generating higher throughput. The VMware virtual switch has three options to configure load balancing:

- **Route based on the originating virtual switch port ID (default configuration):** For outbound traffic, the physical adapter is selected based on the hash of the virtual port. This means that a given virtual network adapter will use only one physical adapter at any given time to transmit network packets. Packets are received on the same physical adapter.
- **Route based on source MAC hash:** Here, the physical adapter is selected for transmit based on the hash on the source MAC address. This means that a given virtual network adapter will use only one physical adapter at any given time to transmit network packets. Packets are received on the same physical adapter.
- **Route based on IP hash:** Here, the physical adapter is selected for transmit based on the hash of the source and destination IP address. EtherChannel can be used with this setting. The physical switch load balancing algorithm will then determine which physical adapter is selected to receive the packets. To achieve proper load balancing, both the virtual switches and the physical switches need to be configured. Note that the virtual switch and the physical switch hashing algorithms work independent of each other.

The default load balancing configuration of route based on the originating virtual switch port ID can be used. Multiple VMs will use different physical adapters to transmit and receive traffic. Service console (for ESX) or

⁴ <http://blogs.vmware.com/security/2008/02/keeping-your-vm.html>

management network (for ESXi) has only one port ID or MAC address. Hence, the service console or management network will use only one physical adapter for communicating, unless there is a failover to another physical adapter. The same is the case for VMotion traffic. VMotion will use only one physical adapter for communicating, unless there is a failover to another physical adapter.

Note: All the physical network adapters will be configured as *active* by default when load balancing is enabled.

6.1.3 Stacking

Each switch I/O module is a standalone switch that can be managed independently. Switches can also be stacked together to form a single virtual blade switch (VBS) using high-speed stacking cables. One of the switch modules acts as the master and the entire VBS can be managed by connecting to the master switch. Stacking provides several advantages. By stacking switch modules together, we manage the switches as a single entity; thereby simplifying management. Stacking also provides for a high-speed data path between the switch modules. The Cisco Catalyst switch modules support stacking and up to eight Cisco switches can be stacked to form a single VBS.

Switch modules A1, A2, C1, and C2 can be stacked together to form a VBS. Four stacking cables are used to form a *ring* as shown in Figure 5. This provides redundancy in case of a single-stack link failure. Stacking for the iSCSI SAN is described in the next subsection.

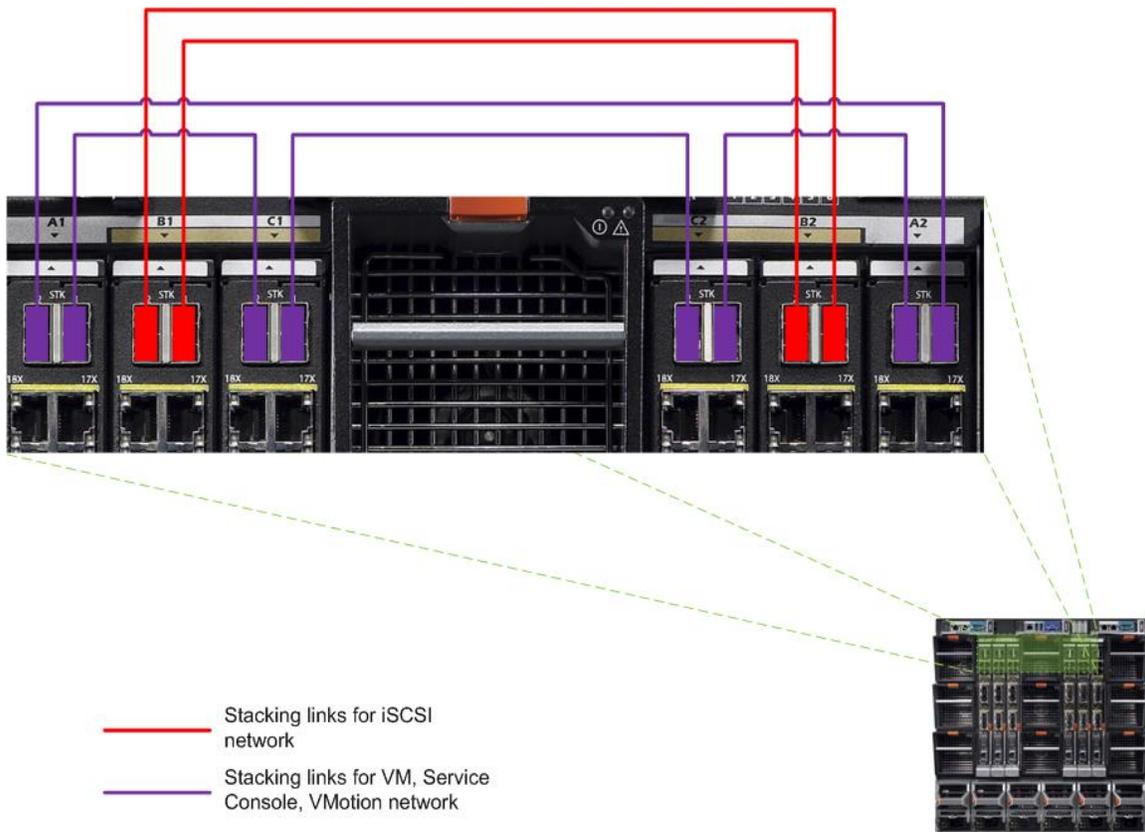


Figure 5: Stacking of I/O modules

6.1.4 Uplinks

The blade chassis can be connected to the existing network infrastructure in the datacenter with various options. Choosing the right option depends on the configuration of the existing infrastructure. In this section, we will discuss the options available and our recommendations for the solution.

Each Cisco 3130X I/O module has four built-in 1 Gbports. Each I/O module can also be populated with two 10 Gb Ethernet uplink ports. This gives us a total of 24 Gb uplink throughput per I/O module. With four I/O modules, this gives us a total of 96 Gb uplink capability.

If the existing infrastructure supports a 10 Gb network, one of the two 10 Gb ports in each Cisco I/O module can be populated. The 10 Gb links from the four switches are joined together into a single EtherChannel to provide a high-speed data path into the blade chassis. The additional 10 Gb port and four 1 Gb ports for each I/O module can be used for future expansion.

If the existing infrastructure does not support 10 Gb, the solution can be configured with Cisco 3130G I/O modules with 1 Gb uplinks. The 3130X I/O modules can also be used. They can be populated with 4 additional 1Gb ports instead of the two 10Gb ports and upgraded to 10Gb later.

VLAN and routing must be configured on these uplinks depending on the existing infrastructure requirements.

6.2 iSCSI SAN Architecture

The iSCSI SAN comprises of iSCSI traffic to the Dell EqualLogic PS Series arrays. In this section, we discuss the best practices for configuring the iSCSI SAN using the software iSCSI initiator in ESX/ESXi. Software initiator in the guest Oses is also supported, but is outside the scope of this document.

A virtual switch (vSwitch1) is created and the two physical network adapters connected are connected as uplinks to the virtual switch. These network adapters are in turn connected to Cisco Catalyst 3130G switches in I/O modules – B1 and B2. Two port groups are created, one for each VMkernel interface. Each VMkernel interface is associated with each physical network adapter. The two VMkernel interfaces are used by the software iSCSI initiator to connect to the storage arrays. This is illustrated in Figure 6, with the two VMkernel interfaces named as “iSCSI VMkernel” and “iSCSI VMkernel2”. Two VMkernel interfaces are created so that both physical network adapters can be utilized by the iSCSI initiator using *Round Robin* load balancing. Detailed instructions for configuring the iSCSI connections are included in section 9.

An optional third port group can also be created (not shown in the figure). A virtual machine can be connected to this portgroup and the VM can be used to monitor or manage the isolated iSCSI network. For example, SAN HQ can be installed in the virtual machine and used to collect statistics from your Dell EqualLogic PS Series arrays.



Figure 6: iSCSI SAN Configuration

6.2.1 Storage Connectivity

Each Cisco Catalyst 3130G I/O module contains four fixed 1 Gb ports. In addition to this, each I/O module has two TwinGig Converter modules placed in the X2 ports and populated with a total of four 1 Gb SFP modules. This gives us a total of sixteen 1 Gb ports for connecting the storage arrays. Each controller on the Dell EqualLogic PS6000 Series array has eight ports (four active and four passive connections). The four active ports are split between two different physical switches to provide redundancy. The passive ports are connected in a similar fashion.

Figure 7 illustrates how two Dell EqualLogic PS6000 Series arrays are connected to the Dell PowerEdge M1000e chassis.

For connecting more than two arrays to the solution, refer to the **Scalability** section.

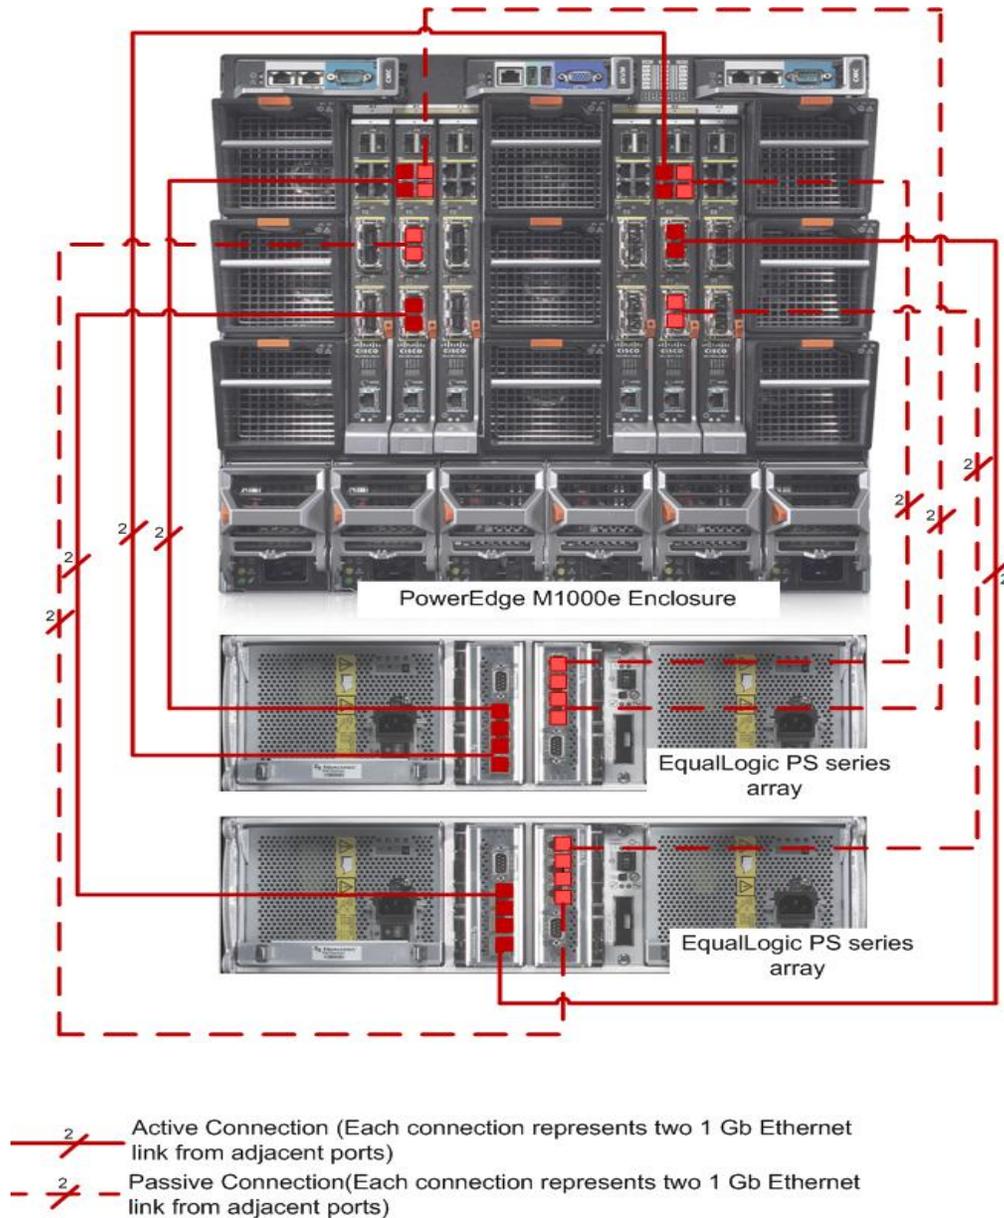


Figure 7: Storage Connectivity

6.2.2 VLAN

The storage traffic is completely isolated from other networks. There is no requirement to enable multiple VLANs for this architecture. The ports on the physical switch can be configured in the access mode. The virtual switch does not have VLANs defined in the port groups.

6.2.3 Load Balancing

Multipathing is a technique that allows more than one physical path to be used to transfer data between a host and an external storage device. In version 4.0 of ESXi, VMware provides the VMware Native Multipathing Plugin. This

plugin supports three Path Selection Plugins: Most Recently Used (MRU), Fixed, and Round Robin (RR). The default selection for the PS6000 is *Fixed*, but this can be changed to *Round Robin* for better performance. In *Round Robin* mode, multiple paths or multiple physical network adapters can be used to access a single iSCSI volume simultaneously, thereby enabling load balancing. Note that in order to have multiple paths, multiple VMKernel Ports for iSCSI must be created, each associated with one physical network adapter, as discussed in this section.

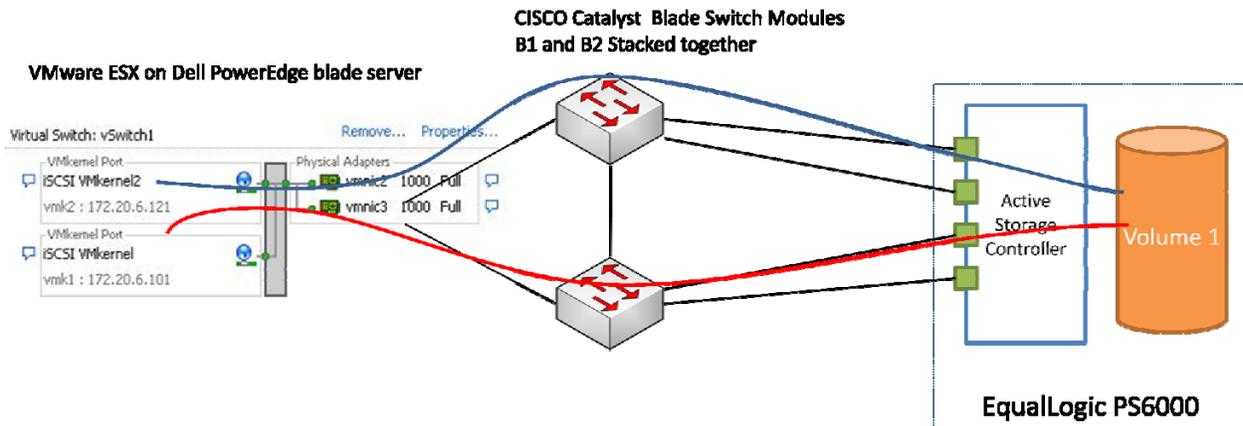


Figure 8: Multipathing using Round Robin

Dell EqualLogic SAN automatically distributes iSCSI connections between the available four network interfaces in each controller. It also automatically distributes volumes between different storage devices in the storage pool. No additional configuration is required.

For a detailed discussion of iSCSI connections, see the VMWare ESX/ESXi 4.0 “iSCSI SAN Configuration Guide”. For a detailed description of how to configure the iSCSI connection, see the section titled “Deploying and Configuring the Solution”.

6.2.4 Stacking

For the iSCSI SAN, switch modules B1 and B2 can be stacked together to form a VBS. Two stacking cables are used to form a *ring* as shown in Figure 5. This provides redundancy in case of a single-stack link failure. Stacking also acts as inter switch link (ISL). The stacking cables provide sufficient bandwidth to meet the ISL best practices for the storage arrays.

6.2.5 Uplinks

The reference architecture connects the blade chassis to two EqualLogic storage arrays directly. This is designed for a standalone configuration and there are no available ports to connect the iSCSI network to an existing infrastructure. The EqualLogic storage should be managed from the management blade or a management virtual machine in the chassis.

If uplinks are required, or the storage needs to be integrated with an external management infrastructure, intermediate switches need to be used. The storage arrays can be connected with a two-tiered architecture. See the section 6.3 for more information on scalability.

6.2.6 EqualLogic Network Requirements

In addition to the above guidelines, EqualLogic has specific recommendations for connecting PS Series arrays to the network. We have highlighted some of the important recommendations below:

- Do not use Spanning-Tree (STP) on switch ports that connect end nodes (iSCSI initiators or array network interfaces). However, if you want to use STP or Rapid STP (preferable to STP), you should enable the port settings available on some switches that let the port immediately transition into the STP forwarding state upon link up. This functionality can reduce network interruptions that occur when devices restart and should only be enabled on switch ports that connect end nodes.

Note: The use of Spanning-Tree for a single-cable connection between switches is encouraged, as is the use of trunking for multi-cable connections between switches.

- Enable Flow Control on each switch port and NIC that handles iSCSI traffic. PS Series arrays will correctly respond to Flow Control.
- Disable unicast storm control on each switch that handles iSCSI traffic if the switch provides this feature. However, the use of broadcast and multicast storm control is encouraged on switches.
- Enable Jumbo Frames on the switches, and VMware kernel interfaces. See section 10 for details on how this is configured.

For more and latest information, see the *Dell EqualLogic PS Quick Start Guide* at <https://www.equallogic.com/support/> (Account registration may be required).

6.3 Scalability

As workloads increase, the reference configuration can be scaled to provide additional resources. In this section we provide guidelines for adding blades servers, blade chassis and storage arrays to the reference configuration.

6.3.1 Scaling Compute Resources

Adding blade servers: If the cluster has reached or is about to exceed its acceptable maximum resource utilization, additional blades can be added to the cluster. For quick addition of a new blade and for the cluster to be uniform, host profiles can be used. A host profile is created from an existing ESXi blade server on the same cluster. It will capture the networking, storage and security settings from the validated configuration. An ESXi blade server with basic vSphere installation can be added to the cluster first. A host profile is then applied to this newly added blade server. To make this process more efficient and to make additional changes not supported in the host profile, a sample script is provided in the Appendix.

Adding blade chassis: If the chassis is full of blades and more compute resources are needed, an additional blade chassis can be added to the solution. To make the two chassis work seamlessly, the network switches could be stacked between the chassis as shown in Figure 9. The two chassis can then be uplinked to the external network infrastructure.

Note that the maximum number of Cisco I/O modules that can be stacked together is eight. To scale beyond two chassis, additional pairs of two chassis can be stacked together and uplinked to the external network infrastructure.

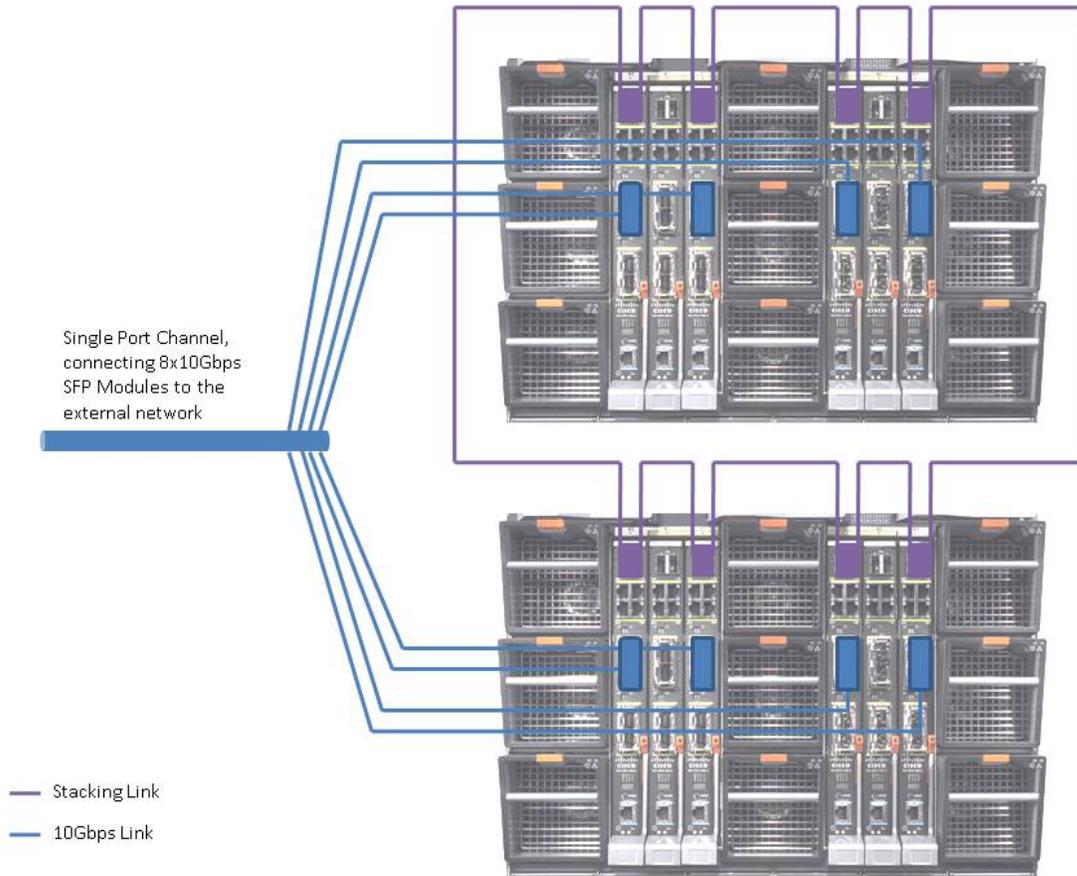


Figure 9: Scaling the compute resource

6.3.2 Scaling Storage Resources

When connecting multiple chassis, the stacking ports can be used to link the iSCSI I/O modules. Up to four blade chassis with two CISCO I/O modules each may be linked together using this method. Each chassis that contains two I/O modules provides enough external ports to add two additional PS6000 arrays.

To scale beyond 4 chassis, or two PS6000 arrays per chassis, you will need to use intermediate switches in the iSCSI fabric. When using intermediate switches the external switches should be stacked to allow for the communication between the EqualLogic arrays. However, the internal I/O module switches should not be stacked because a port channel must be used to connect the internal switches to the external switches, and the port-channel is limited to eight network ports. Unstacking the internal switches allows for the creation of two separate port channels, one for each internal switch. See Figure 10 for setting a scalable configuration with external Ethernet switches.

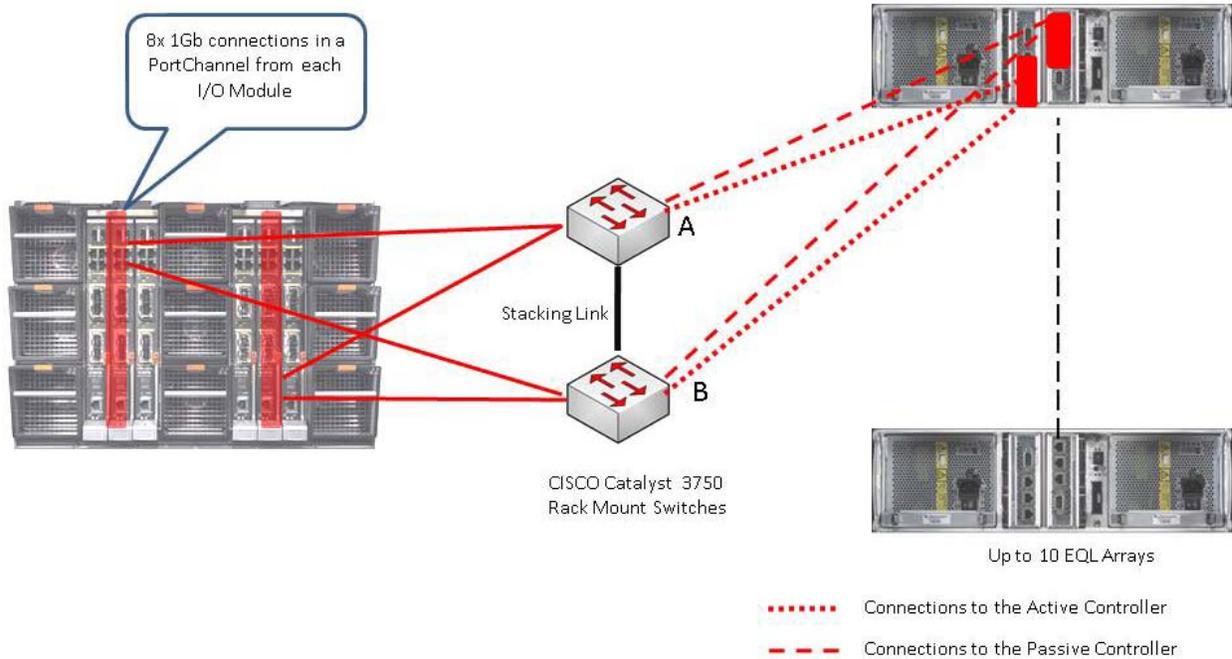


Figure 10: Storage setup for a scalable configuration

7 Storage Configuration

The section provides recommendations for configuring the Storage Array. For storage networking and stacking guidelines, see Section 6.2.4.

7.1 Volume Size Considerations

Volume sizes would depend on the customer environment and the type of workloads. With the changes in the iSCSI load balancing architecture for ESX 4.0, the previous recommendation of using multiple volumes has been removed. Using multiple VMkernels in a “round-robin” setting allows ESX to establish multiple connections to a single volume. Depending on the environment, you may decide to create multiple ~500 GB volumes with multiple VMs. It is important to include space for Guest OS memory cache, snapshots, and VMware configuration files when sizing the volumes. Additionally, you can configure thin-provisioned volumes to grow on demand only when additional storage is needed for those volumes. Thin-provisioning can increase the efficiency of storage utilization. For a detailed discussion of VMFS considerations, see the VMware “iSCSI SAN Configuration Guide”, http://www.vmware.com/pdf/vsphere4/r40/vsp_40_iscsi_san_cfg.pdf

7.2 Array RAID Configuration

The storage array RAID configuration is highly dependent on the workload in your virtual environment. The Dell EqualLogic PS Series arrays support four RAID types – RAID 5, RAID 6, RAID 10, and RAID 50. It is recommended to configure the array in RAID 10 for best throughput and performance. In general, RAID 10 provides the best performance, at the expense of storage capacity. RAID 50 generally provides more usable storage, but has less performance than RAID 10 in random I/O situations and requires additional overhead in the case of a drive failure scenario. RAID 5 provides the most storage capacity at the expense of a little lower performance and availability.

8 Management

In this section, we describe the various components used to manage the solution.

Figure 11 displays the various management components.

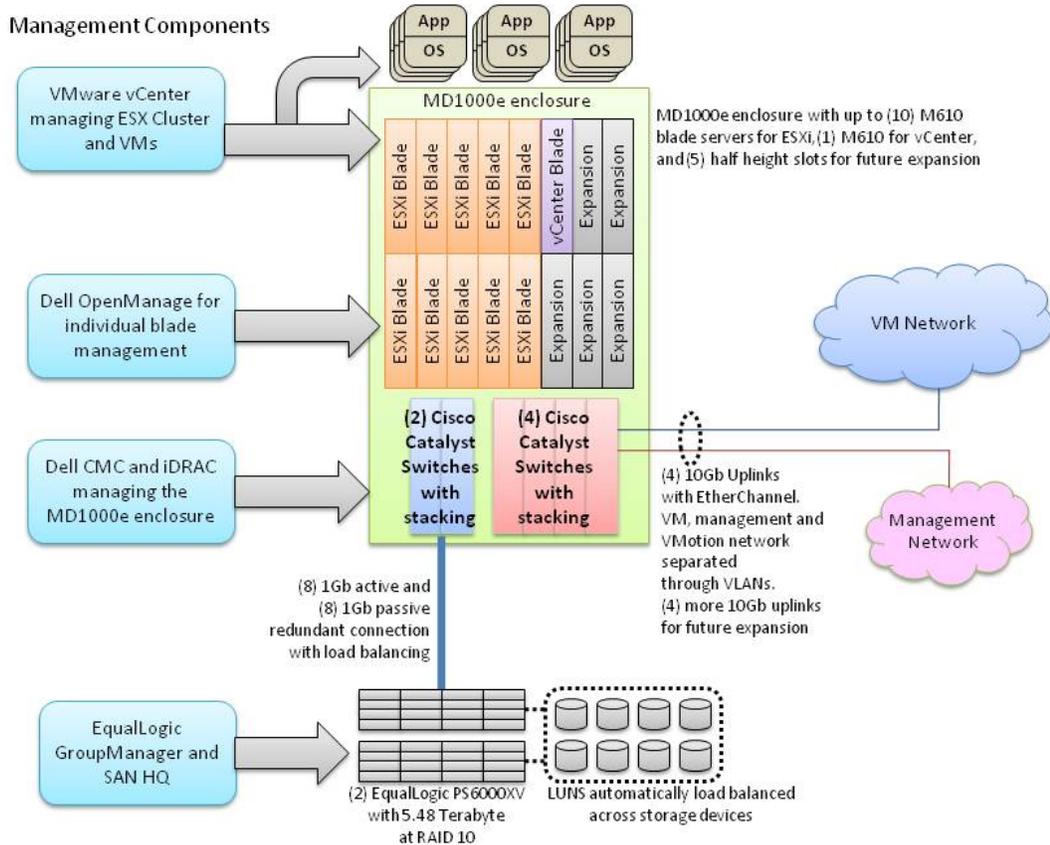


Figure 11: Management components

8.1 Virtualization Management

The VMware vCenter provides a single control and management center for VMware ESX Servers. VMware vCenter allows you to manage all ESX Servers and virtual machines, utilize advanced features such as VMotion, HA, and DRS, monitor performance and resource optimization, and use standard templates to provision new servers in minutes. For more information on vCenter, see <http://www.vmware.com/products/vcenter/>.

PlateSpin® Workload Management allows you to profile, migrate, protect, and manage server workloads between physical and virtual infrastructures in heterogeneous IT environments. The PlateSpin Workload Management portfolio for Dell servers gives Dell customers a powerful data center management platform for Dell's 32- and 64-bit Windows and Linux servers with support for leading hypervisors. For more information on PlateSpin migration tools and other products, see www.dell.com/platespin.

8.2 Systems Management

The following tools and software can be used to manage the hardware:

Dell OpenManage: Dell OpenManage Server Administrator (OMSA) can be installed on VMware ESX and can be used to manage the hardware. VMware ESXi comes integrated with Dell OpenManage Server Administrator. For more information on OpenManage and its capabilities, see www.dell.com/openmanage.

Deployment and change management using Unified Server Configurator (USC): The Dell PowerEdge M610 blade servers come with USC. This helps customers reduce operating costs by simplifying deployment and management. Key features include: Diagnostics, self-update (UEFI, Driver Pack update), firmware updates (BIOS, NIC FW, RAID Controllers), and hardware configuration.

Out-of-band CMC and iDRAC: CMC provides a single, secure interface to manage the inventory, configuration, monitoring, and alerting for chassis components (iKVM, CMC), I/O modules, servers, and iDRAC. It also provides excellent real-time power management, monitoring, and alerting capabilities. The Dell chassis also provides the users with system-level power limiting, slot-based prioritization, and dynamic power engagement functionalities. iDRAC on each server provides the flexibility to remotely manage the server through Console redirection and Virtual CD-ROM/DVD/Floppy/Flash capabilities.

8.3 Storage Management

Dell EqualLogic provides a rich set of management features that is available at no additional cost and come with exceptionally easy-to-use management tools. For more information on Dell EqualLogic features and management capabilities, see www.dell.com/equallogic.

9 Deploying and Configuring the Solution

This section provides guidelines for the initial deployment and configuration of the solution.

9.1 Rack, Stack, and Cable

1. Rack the chassis with blade servers and I/O modules.
2. Rack the Dell EqualLogic storage.
3. Connect power to the chassis and storage.
4. Consider PDUs. This document provides detailed information on how to configure the PDUs: http://www.dell.com/downloads/global/products/pedg/en/pe_m1000e_selection_whitepaper.pdf
5. Cable servers and storage.
 - a. Cable the server and storage for Ethernet connectivity as per the architecture.
 - b. Connect the switches with the given stacking configuration. Cable all the Ethernet switches with the stacking cables before you power on the chassis for the first time.

9.2 Initial Setup

Configuring the Chassis

Configure the Chassis Management Controller using the web interface or the front LCD panel on the chassis. By default the CMC will get a DHCP address, which you can view from the LCD.

- Use static chassis management IP address (preferable).
- Enable FlexAddressing.
- Configure the iDRAC IP address for the blades.
- Configure the Cisco Switch Management IP.
- PowerThreshold.

Follow the procedures in the CMC User Guide

<http://support.dell.com/support/edocs/software/smdrac3/cmc/cmc2.0/en/ug/pdf/cmcug.pdf>) and iDrac User Guide (<http://support.dell.com/support/edocs/software/smdrac3/idrac/idrac20modular/en/ug/pdf/ug.pdf>)

Configuring Networking

1. Configure network switch modules. It is recommended that you create a configuration file for the network switches. Sample switch configuration files are included in appendix A. These files are written in a sparse format using the range command for interfaces. These scripts are also available for download here: <http://www.delltechcenter.com/page/Virtualization+Scripts>
Look for the file named BRC_40_Networking.zip

Two ways of loading the switch file onto the switches are recommended:

Serial: You can connect a serial cable from a Microsoft Windows® or Linux computer to the master switch of a VBS, which can be identified by the blue light at the bottom of the switch. You will need a terminal program such as HyperTerminal.

TFTP: You can use the Chassis CMC to configure an IP address for each VBS, then use telnet or ssh to connect to the switch and then tftp to copy the config file to the switch. A TFTP server is required for this method.

Configuring the Server

Configure the server either locally or by using the iDRAC Console Redirection remotely. Perform the following:

- Enable VT for BIOS.
- For ESXi, it is recommended to have diskless servers.
- Configure Unified Server Configurator for self-update and firmware repository.
- Assign a password, management IP address, and DNS name to ESXi using the DCUI (Direct Console User Interface).

See the “ESXi Embedded and vCenter Server Setup Guide”,

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxi_e_vc_setup_guide.pdf

See the Unified Server Guide (http://www.dell.com/downloads/global/solutions/unified_server_overview.pdf) for more information.

Configuring Storage

Configure the Dell EqualLogic PS Series. For information on configuring storage, see the *Dell EqualLogic PS QuickStart Guide* at <https://www.equallogic.com/support/> (Account registration may be required).

1. Create volumes
2. Create access list
3. Ensure that you add the VMkernel iSCSI IPs to the access list. (Use * to denote the IP address range.)
4. Create volumes to be shared between multiple servers (cluster).

Configuring vCenter

You will need to install SQL database before installing VMware vCenter. See the “ESXi Embedded and vCenter Server Setup Guide”, http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxi_e_vc_setup_guide.pdf to install vCenter on the management server.

Deployment Steps

Note that ESXi 4.0 Patch 1 or later is required in order to be a supported configuration.

1. Create a cluster in vCenter with DRS and DPM features enabled. Add the ESX servers to the cluster.
2. Place the servers in maintenance mode.
3. Download the VMA appliance from VMware and deploy it on an existing server.
4. Sample deployment scripts are included in Appendix A.
5. Create the deployment scripts in the VMA. Ensure the scripts have sufficient permissions for execution.
6. In the script named doiscsi.sh, change the server names and both vmkernel IP addresses to match your network.
7. In the script named configiSCSI.sh:
 - Change the target address to the GroupAddress of your EqualLogic array.
 - Change the username and password for the esxcli swiscsi commands.
 - Make sure that the iSCSI software adapter name, vmhba35, matches your hardware configuration. (You can use vSphere client to check this)
 - Change the temp path to one of your iSCSI luns. Make sure the directory on your lun already exists. (This is required so you have enough space to apply Patch 1), and
8. Make sure that the VMA can properly resolve your host names, for example, ping blade1.sb.lab.
9. Initialize VMA
./doiscsi.sh myinit.sh
The sudo command may prompt for your vi-admin password for your vma appliance.

10. Perform the initial configuration of your iSCSI interfaces. (vifpinit only has to be done once per login)
vifpinit
./doiscsi.sh deploy.sh
11. Download the file name BRC_40_profile.zip from <http://www.delltechcenter.com/page/Virtualization+Scripts>, it contains a file named profile4.zip
12. In vCenter, import the host profile, bundle4.vpf, and name it bundle4.
13. Customize the profile for your environment, for example, change the NTP server address.
14. In vCenter, attach and apply the bundle4 profile to each of your hosts. When you apply the profile, you may get an error communicating with the remote host.. After a couple of minutes, run a compliance check again.
The profile should make the following changes on each server:
 - ReCreate all the switches and port groups
 - Update the DN Configuration
 - Configure the Time Servers
15. Run the script to configure iSCSI and load balancing.(vifpinit only has to be done once per login)
vifpinit
./doiscsi.sh configiSCSI.sh
16. Reboot the server to enable changes to scratchconfig
17. Apply Patch 1, the script assumes that the .zip patch bundle is in the same directory
vifpinit
./doiscsi.sh myupdate.sh
18. Reboot the server again to enable patch 1.
19. Power Monitoring and Management:
 - a. Use CMC to perform power monitoring of Dell blades when workload VMs are running.
 - b. Use DMC to set up SNMP alerts for physical machines for easy management.

10 References

iSCSI overview - A "Multivendor Post" to help our mutual iSCSI customers using VMware

http://virtualgeek.typepad.com/virtual_geek/2009/01/a-multivendor-post-to-help-our-mutual-iscsi-customers-using-vmware.html

Integrating Blade Solutions with EqualLogic SANs

http://www.dell.com/downloads/global/partnerdirect/apj/Integrating_Blades_to_EqualLogic_SAN.pdf

Cisco Products

http://www.cisco.com/en/US/products/ps6746/Products_Sub_Category_Home.html

Cisco 3130 Product Page

<http://www.cisco.com/en/US/products/ps8764/index.html>

Cisco Catalyst 3750 and 2970 Switches: Using Switches with a PS Series Group

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5269>

Common issues when using the ESX SW iSCSI Initiator

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5271>

Deploying Pools and Tiered Storage in a PS Series SAN

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5239>

Deploying Thin Provisioning in a PS Series SAN

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5245>

Microsoft Windows: Aligning Disk Sectors for Optimal Performance

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5223>

Network Connection and Performance Guidelines

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5311>

Network Connection Guidelines

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5331>

Protecting the Virtual Environment Using Auto-Snapshot Manager/VMware Edition

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=7369>

PS Series Array Network Performance Guidelines

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

Snapshots in a VMware Virtual Infrastructure

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5259>

VMware Virtual Infrastructure 3.x Considerations, Configuration, and Operation Using an EqualLogic PS Series SAN

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5235>

Windows Disk Sector Alignment

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5323>

VMware vSphere 4 iSCSI SAN Configuration Guide

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_iscsi_san_cfg.pdf

11 Appendix - Deployment Scripts

In this section we list the scripts referred in section 9. These sample scripts are provided “as-is” and can be modified to suit your existing environment. It is highly recommended to test these scripts in a pre-production environment.

deploy.sh

```
#!/bin/sh
# This script creates a virtual switch in ESXi server
# and assigns IP address
echo $1
/usr/bin/esxcfg-vswitch --server $1 --add vSwitch1
/usr/bin/esxcfg-vswitch --server $1 --mtu 9000 vSwitch1
/usr/bin/esxcfg-vswitch --server $1 --link vmnic2 vSwitch1
/usr/bin/esxcfg-vswitch --server $1 --link vmnic3 vSwitch1
/usr/bin/esxcfg-vswitch --server $1 --add-pg "iSCSI VMkernel" vSwitch1
/usr/bin/esxcfg-vmknic --server $1 --add "iSCSI VMkernel" --mtu 9000 --
ip=$2 --netmask=255.255.0.0
/usr/bin/esxcfg-vswitch --server $1 --add-pg "iSCSI VMkernel2" vSwitch1
/usr/bin/esxcfg-vmknic --server $1 --add "iSCSI VMkernel2" --mtu 9000 --
ip=$3 --netmask=255.255.0.0
```

configiSCSI.sh

```
#!/bin/sh
# This script enables iSCSI and configures load balancing for iSCSI LUNs
echo $1
/usr/bin/vicfg-iscsi --server $1 --swiscsi --enable
/usr/bin/esxcli --server $1 --username root --password password swiscsi
nic add -n vmk1 -d vmhba35
/usr/bin/esxcli --server $1 --username root --password password swiscsi
nic add -n vmk2 -d vmhba35
/usr/bin/vicfg-iscsi --server $1 --discovery --add --ip 172.20.218.110
vmhba35
/usr/bin/vicfg-rescan --server $1 vmhba35
/usr/bin/esxcli --server $1 nmp device list | grep EQLOGIC |awk '{print
substr($7,2,index($7,"))-2})}' > llist.txt
for lun in $(cat llist.txt); do /usr/bin/esxcli --server $1 nmp device
setpolicy --device $lun --psp VMW_PSP_RR; done
# set up scratch directory
/usr/bin/vicfg-advcfg --server $1 --set /vmfs/volumes/lun1/$1/
ScratchConfig.ConfiguredScratchLocation
```

myupdate.sh

```
#!/bin/sh
# This script applies Patch 1
echo $1
vihostupdate --server $1 --bundle ESXi400-200906001.zip --install
```

doiSCSI.sh

```
#!/bin/sh
# This is wrapper script used to execute the above scripts one at a time
# for each host
# Example invocation ./doiSCSI.sh deploy.sh
# sample ./ $1 <ESX host name> <VMk1 IP address> <VMk2 IP address>
```

Business Ready Configurations for Data Center Virtualization

```
./$1 blade2-1.sb.lab 172.20.6.201 172.20.6.221  
./$1 blade2-2.sb.lab 172.20.6.202 172.20.6.222  
./$1 blade2-3.sb.lab 172.20.6.203 172.20.6.223  
./$1 blade2-4.sb.lab 172.20.6.204 172.20.6.224
```

myinit.sh

```
#!/bin/sh  
# This initializes fast pass in VMA for authentication  
echo $1  
sudo vifp addserver $1 --username root --password <password>  
/opt/vmware/vima/bin/vifpinit $1
```