

Reference Architecture for Active System 50 with VMware vSphere

Release 1.1 for Dell PowerEdge Rack Servers, Dell Networking Switches, Dell EqualLogic SAN, and Dell Active System Manager

Dell Virtualization Solutions Engineering

Revision: A00

July 2013



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

© 2013 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. *Dell*, the *Dell logo*, *OpenManage*, *Kace*, *EqualLogic*, *PowerVault*, and *PowerEdge* are trademarks of Dell Inc. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries. *Microsoft*, *Windows*, *Hyper-V*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *VMware*, *vSphere*, *ESXi*, *vMotion*, *vCloud*, and *vCenter* are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. *Linux* is the registered trademark of Linus Torvalds in the U. S. and other countries. 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.

July 2013

Revision History

Revision	Description	Date
A00	Initial Version	July 2013

Contents

1	Introduction	3
2	Audience	3
3	Overview	4
4	Design Principles	10
5	Prerequisites and Datacenter Planning.....	10
6	Architecture.....	11
6.1	Server to LAN Connectivity	13
6.2	Server to SAN Connectivity	14
6.3	Storage Architecture	15
7	Management Infrastructure	17
7.1	Management Components	17
7.2	Out-Of-Band Management Connectivity.....	21
8	Connecting to Datacenter Network	22
8.1	Connecting to a Datacenter Network with Dell Networking Switches	22
8.2	Connecting to a Cisco Datacenter Network	23
9	Delivery Model	24
10	Reference.....	25

Figures

Figure 1: Active System 50v Overview 4

Figure 2: Network Topology (Logical View) shown with 2 Dell PowerEdge R620 Servers..... 12

Figure 3: LAN vSwitch Configuration for Compute Servers 14

Figure 4: SAN vSwitch Configuration 15

Figure 5: PS6100 Array SAN Connectivity 16

Figure 6: Management Components..... 18

Figure 7: Connectivity of OOB management components..... 22

Figure 8: Connectivity to a Dell Network 23

Figure 9: Connectivity to Cisco Network 23

Figure 10: Active System 50v Rack and Component Overview 24

Tables

Table 1: Solution Components 5

Table 2: Traffic Type Summary 13

1 Introduction

Dell™ Active Infrastructure is a family of converged infrastructure solutions that combine servers, storage, networking, and infrastructure management into an integrated and optimized system that provides general purpose virtualized resource pools. Active Infrastructure leverages Dell innovations including unified management with Active System Manager, LAN and SAN fabrics, and rack server architecture. Active Infrastructure helps IT organizations rapidly respond to dynamic business demands, maximize datacenter efficiency, and strengthen IT service quality.

The Active System 50, a member of Dell Active Infrastructure family, is a converged infrastructure solution that is designed and validated by Dell Engineering. It is available racked, cabled, and delivered to your site to speed deployment. Dell Services deploys and configures the system tailored for business needs and ready for integration into your datacenter. Active System 50 is Dell's first Active Infrastructure System to include Dell PowerEdge™ R620 Rack Servers, iSCSI 1GbE fabric switching, Dell EqualLogic™ PS6100 Series 1GbE iSCSI storage, and Dell Active System Manager to build private cloud solutions that address key needs for small to medium business and data centers. The solution is offered in configurations with either VMware® vSphere™ or Microsoft® Windows Server® 2012 with Hyper-V® role enabled Hypervisors. The VMware vSphere solution is the Active System 50v and the Microsoft Hyper-V solution is the Active System 50m. This document defines the reference architecture for Active System 50v.

2 Audience

This document provides an overview of the Active System 50v solution. Readers, including CTOs and IT managers, can use this document to understand the overview and scope of the solution. IT administrators and managers can use this document to understand the solution architecture.

3 Overview

This section provides a high-level product overview for the major components of the Active System 50v as illustrated in Figure 1. Readers can skip the sections of products with which they are familiar.

Figure 1: Active System 50v Overview



VMware vSphere 5.1 Hypervisor

- vMotion, Storage vMotion
- VMware HA and DRS

Dell PowerEdge R620 & R420 Servers

- 12th generation R620 and R420 rack server
- Large memory and I/O capacity
- Powerful systems management with Dell iDRAC and Lifecycle Controller
- Scalable I/O capabilities that include PCIe Gen3-capable expansion slots

Dell EqualLogic Storage

- 1GbE and 10GbE iSCSI SAN arrays
- Thin Provisioning and Storage Tiering
- Integration with VMware

Dell 2700W Line Interactive Rack UPS

- LCD control panel with rich intuitive user interface
- Network Management Card (NMC) for remote management over the network
- Remote monitoring and configuration of multiple UPS using Dell UPS Management Software
- Integrates with VMware, Microsoft and Citrix

Integrated Management

- VMware vCenter Server
- Dell Active System Manager
- Dell Management plug-in for VMware vCenter
- Dell OpenManage Essentials and Dell Repository Manager
- Dell EqualLogic Virtual Storage Manager (VSM) for VMware
- Dell EqualLogic SAN HeadQuarters (HQ)
- Dell Multi-UPS Management Console (MUMC)

Cloud Enablement

- VMware vCloud Connector for Dell vCloud connectivity

Table 1 below describes the key solution components and the roles served.

Table 1: Solution Components

Component	Description	Role
Compute Cluster	PowerEdge R620 servers running embedded VMware vSphere 5.1 Update 1	Host highly-available virtual machines (VMs)
Management Server	PowerEdge R420 server with embedded VMware vSphere hosting management VMs.	Host management VMs: Dell Active system Manager, VMware vCenter Server, Dell Management Plug-in for VMware vCenter, Dell OpenManage Essentials, Dell EqualLogic Virtual Storage Manager (VSM) for VMware, Dell EqualLogic SAN HeadQuarters (SAN HQ), VMware vCloud Connector, Dell Repository Manager
Storage	Dell EqualLogic PS6100 Series controllers with 24 2.5" or 3.5" hard drive enclosures	Provide shared storage for the ESXi Hypervisor clusters
LAN Switches	Two Dell Networking 7024	Support VM, vMotion, Management and Out-of-Band (OOB) traffic
SAN Switches	Two Dell Networking 7024	Support iSCSI traffic

VMware vSphere: VMware vSphere includes the ESXi™ hypervisor as well as vCenter™ Server which is used to configure and manage VMware hosts. Key capabilities for the ESXi Enterprise Plus license level include:

- **VMware vMotion™:** VMware vMotion technology provides real-time migration of running virtual machines (VM) from one host to another with no disruption or downtime.
- **VMware High Availability (HA):** VMware HA provides high availability at the virtual machine level. Upon host failure, VMware HA automatically re-starts VMs on other physical hosts running ESXi. VMware vSphere uses Fault Domain Manager (FDM) for High Availability.
- **VMware Distributed Resource Scheduler (DRS) and VMware Distributed Power Management (DPM):** VMware DRS technology enables vMotion to automatically achieve load balancing according to resource requirements. When VMs in a DRS cluster need fewer resources, such as during nights and weekends, DPM consolidates workloads onto fewer hosts and powers off the rest to reduce power consumption.
- **VMware vCenter Update Manager:** VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.
- **VMware Storage vMotion™:** VMware Storage vMotion enables real-time migration of running VM 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 ESXi host configurations. They capture and store validated configuration information, including host compliance, networking, storage, and security settings.

For more information on VMware vSphere, see www.vmware.com/products/vsphere.

Dell Active System Manager: Active System Manager is an intelligent and intuitive converged infrastructure and workload manager. Active System Manager leverages templates to automate infrastructure provisioning, on-boarding, and re-configuration, which greatly simplifies and speeds up the process, and also significantly reduces errors associated with manual configuration. The result is better infrastructure and workload quality with fewer configuration errors that can be costly.

The key capabilities of Dell Active System Manager are:

- **Template-Based Provisioning** – Streamline and standardize workload deployments through centralized capture and application of best practices and operational steps
- **Infrastructure Lifecycle Mgmt** – Discovery, inventory, configuration, provisioning, and ongoing management of physical and virtual infrastructure
- **Resource Pooling and Dynamic Allocation** – Create and manage physical and virtual resource pools; efficiently schedule or allocate resources on-demand
- **End-To-End Automation** – Multi-tier automation across physical (server, storage and network) and virtual layers
- **Workflow Orchestration** – Intelligent workflow orchestration engine for rapid physical and virtual workload provisioning
- **Centralized Management** – Intuitive centralized, role-based management and access through self-service web portal

For more information on Dell Active System Manager, see [Dell Active System Manager](#).

Dell PowerEdge R620 Compute Server: The PowerEdge R620 rack server is the Dell 12th generation PowerEdge Rack server offering:

- New high-efficiency Intel® Xeon® E5-2600 family processors for more advanced processing performance, memory, and I/O bandwidth.
- Greater memory density than any previous Dell PowerEdge server. Each PowerEdge R620 can deploy up to 24 DIMMs, or 256GB of RAM per server.
- ‘Agent Free’ management with the new iDRAC7 with Lifecycle Controller allows customers to deploy, update, maintain, and monitor their systems throughout the system lifecycle without a software management agent, regardless of the operating system.
- The PowerEdge Network Daughter Card (NDC) on the PowerEdge R620 offers several choices for embedded network interface communication. With 1GbE and 10GbE CNA offerings from Broadcom, QLogic, and Intel, customers can choose the networking vendor and technology that’s right for them and their applications, and even change in the future as those needs evolve over time.

Embedded Management with Dell's Lifecycle Controller: The Lifecycle Controller is the engine for advanced embedded management and is delivered as part of iDRAC Enterprise in 12th-generation Dell PowerEdge servers. It includes 1GB of managed and persistent storage that embeds systems management features directly on the server, thus eliminating the media-based delivery of system management tools and utilities previously needed for systems management. Embedded management includes:

- Unified Server Configurator (USC) aims at local one-to-one deployment via a graphical user interface (GUI) for operating system install, updates, configuration, and for performing diagnostics on single local servers. This eliminates the need for multiple option ROMs for hardware configuration.
- Remote Services are standards-based interfaces that enable consoles to integrate, for example, bare-metal provisioning and one-to-many OS deployments, for servers located remotely. Dell's Lifecycle Controller takes advantage of the capabilities of both USC and Remote Services to deliver significant advancement and simplification of server deployment.
- Lifecycle Controller Serviceability aims at simplifying server re-provisioning and/or replacing failed parts, and thus reduces maintenance downtime.

For more information on Dell Lifecycle Controllers, see <http://content.dell.com/us/en/enterprise/dcsm-embedded-management>.

Dell OpenManage Essentials: The Dell OpenManage™ Essentials (OME) Console provides a single, easy-to-use, one-to-many interface through which to manage resources in multivendor operating system and hypervisor environments. It automates basic repetitive hardware management tasks – like discovery, inventory, and monitoring – for Dell servers, storage, and network systems. OME employs the embedded management of PowerEdge™ servers – Integrated Dell Remote Access Controller 7 (iDRAC7) with Lifecycle Controller – to enable agent-free remote management and monitoring of server hardware components like storage, networking, processors, and memory.

OpenManage Essentials helps you maximize IT performance and uptime with capabilities like:

- Automated discovery, inventory, and monitoring of Dell PowerEdge™ servers, Dell EqualLogic™ and Dell PowerVault™ storage, and Dell Networking switches
- Server health monitoring, as well as BIOS, firmware, and driver updates for Dell PowerEdge servers, blade systems, and internal storage
- Control of PowerEdge servers within Microsoft® Windows®, Linux®, VMware®, and Hyper-V® environments

For more information on OpenManage Essentials, see Dell.com/openmanageessentials.

Dell Networking 7024 Switches: At the heart of the solutions network configuration are four Dell Networking 7024 switches. These managed Layer 3 Gigabit Ethernet switches offer the enterprise-class level of performance required for this configuration. LAN and SAN switches are physically and logically separated per best practices to support security and network traffic segmentation. VLANs are implemented to support solution management, security, and network traffic segmentation, and routing is leveraged to provide flexible connectivity.

Dell EqualLogic PS6100 Series Storage: The Dell EqualLogic PS6100 is a virtualized iSCSI SAN that combines intelligence and automation with fault tolerance to provide simplified administration, rapid deployment, enterprise performance and reliability, and seamless scalability. The storage architecture delivers a self-optimizing SAN that is simple to manage and has an all-inclusive software suite to help reduce Total Cost of Ownership (TCO).

With a 24 drive chassis full of 300GB SAS drives, the PS6100X array delivers 7.2 Terabyte (TB) of iSCSI-based storage built on fully-redundant, hot-swappable enterprise hardware. Scale out capacity and performance is provided by adding additional arrays. Built-in software functionality includes automated load balancing, snapshots and replication, multi-path I/O (MPIO), and consistency sets. SAN HQ is also available for Multi-SAN historical performance monitoring.

Advanced data protection features such as Auto Replication and Auto-Snapshot Manager (ASM) also come standard. The ASM integrates with VMware vCenter and VMware's native snapshot technology to provide intelligent and consistent snapshots.

Dell PowerEdge R420 Management Server: The Dell PowerEdge R420 uses Intel® Xeon® E5-2400 series processors and Intel chipset architecture in a 1U rack mount form factor. These servers support up to eight 2.5" drives and provide the option for an LCD located in the front of the server for system health monitoring, alerting, and basic management configuration. An AC power meter and ambient temperature thermometer are built into the server, both of which can be monitored on this display without any software tools. The server features two CPU sockets and 12 memory DIMM slots.

Dell Management Plug-in for VMware vCenter: Dell Management Plug-in for VMware vCenter is included in the solution. This enables customers to:

- Get deep-level detail from Dell servers for inventory, monitoring, and alerting – all from within vCenter
- Apply BIOS and Firmware updates to Dell servers from within vCenter
- Automatically perform Dell-recommended vCenter actions based on Dell hardware alerts
- Access Dell hardware warranty information online
- Rapidly deploy new bare metal hosts using Profile features

For more information, see the web page for [Dell Management Plug-in for VMware vCenter](#).

Dell 2700W line Interactive Rack UPS: Monitor your power consumption, power events, available battery runtime, alarm conditions, and more with free software and a range of control options.

- **The Dell Multi-UPS Management Console (MUMC):** Designed for multi-host server environments, Dell's console facilitates easy and versatile monitoring and management across the network from a single interface. MUMC also provides agentless integration into multi-host managers.
- **The Dell UPS Local Node Manager (ULNM):** Created for single-host environments, Dell's node manager provides graceful automatic shutdown of any UPS, load segment or connected device during prolonged power disruptions, preventing data loss and saving work-in-progress.
- **Free multi-platform software:** Manage your UPS systems with software that works intuitively across most popular virtualization platforms.

Reference Architecture for Active System 50 with VMware vSphere

- **Large on-unit display:** Change control settings easily with a large, backlit LCD control panel offering the same color-coded system as Dell servers.
- **Virtual Environments:** Integrates with VMware, Microsoft, and Citrix. Manage virtual machine start/stop profiles through plug-ins like VMware vCenter. Use advanced virtualization functions like vMotion™ and LiveMigration to move VMs.

For more UPS and PDU information, see the web page for [Dell UPS and PDU Information](#).

VMware vCloud Connector: VMware vCloud Connector lets you view, operate on, and transfer your computing resources across vSphere and vCloud Director in your private cloud environment, as well as to public cloud.

- Expand your view across hybrid clouds. Use a "single pane of glass" management interface that seamlessly spans your private vSphere and public cloud environment.
- Extend your datacenter. Move VMs, vApps, and templates from private vSphere to a public cloud to free up your on-premise datacenter resources as needed.

For more information, see the web page for [VMware vCloud Connector](#).

4 Design Principles

The following principles are central to the design and architecture of Active System 50v Solution.

1. **Redundancy to minimize single point of failure:** The system is designed to mitigate failure points. NIC teaming and MPIO are used to provide failover across the redundant network interfaces. iSCSI storage redundancy is achieved with multiple ports and storage controllers. For network traffic, NIC ports are teamed in such a way to distribute traffic across separate ports. The solution also includes redundant power supplies connected to separate PDUs. The solution is further optimized to reduce components such as NICs, cables, and IO Modules, then utilizes hypervisor based high-availability to provide virtual machine failover. Out-of-Band Management is not architected with this level of redundancy since mission critical workloads will continue to operate in the event of an OOB management failure.
2. **Management:** Provide integrated management using Active System Manager, VMware vCenter, Dell Management Plug-in for VMware vCenter, Dell OpenManage Essentials, and EqualLogic Virtual Storage Manager (VSM) plug-in for VMware.
3. **Dedicated iSCSI Network:** Dedicated NICs and switches are provided for iSCSI storage traffic to isolate the storage traffic from LAN traffic. This ensures minimal latency for storage I/O operations.
4. **Cloud Enabled:** The solution also includes connectivity to public cloud using VMware vCloud Connector.
5. **Hardware configuration for virtualization:** This solution is designed for general virtualization use cases. Each server is configured with appropriate processor, memory, and network adapters, as required for virtualization.
6. **Racked, Cabled, and Ready to be deployed:** Active System 50v is available racked, cabled, and delivered to the customer site, ready for deployment. Components are configured and racked to optimize airflow and thermals.
7. **Flexible configurations:** Active System 50v is pre-configured to suit most customer needs for a virtualized infrastructure. The solution also supports additional options, such as configuring server processors, server memory, and storage, based on customer needs.

5 Prerequisites and Datacenter Planning

To support the architecture, the following components are required to be present in the customer environment:

- An existing Ethernet infrastructure with which to integrate. 1Gb or 10Gb Ethernet infrastructure is recommended. Additional components, such as Dell then *10GbE Uplink Modules for SFP+* may be needed to uplink the solution to the customer network. The necessary components depend upon customer networking and uplink requirements.
- Domain Name System (DNS) and Network Time Protocol (NTP) Server must be available on the management network.
- Database - Required for VMware vCenter Server.
- Sufficient power and cooling to support the solution must be present.

6 Architecture

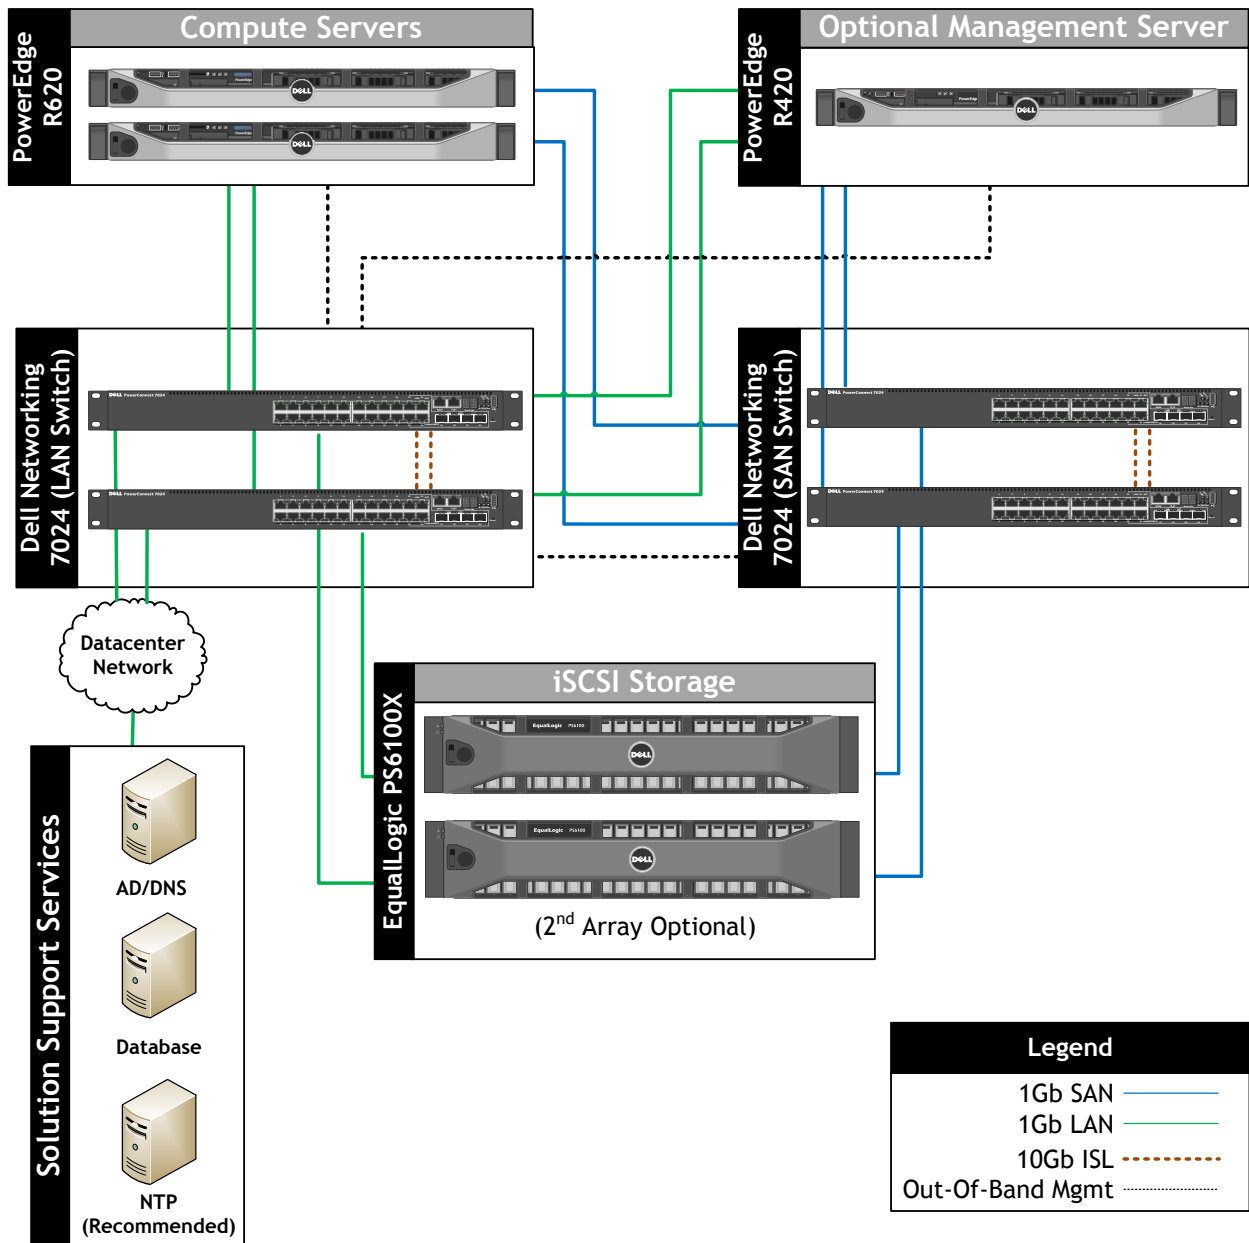
This solution consists of PowerEdge R620 and R420 rack servers running VMware ESXi connected to 7024 network switches.

vSphere network traffic in this solution is comprised of four distinct types: Virtual Machine (VM), Management, vMotion, and iSCSI. In addition, support for out-of-band management is included. Two separate networks are created to support different traffic types:

- **LAN** - This network supports management, VM, vMotion, and out-of-band management. In addition, uplinks to core infrastructure provide connectivity to the solution support services (DNS, NTP, and Database for vCenter Server).
- **SAN** - This network supports iSCSI traffic.

Figure 2 provides the high-level reference architecture for the solution. The figure shows high-level logical connectivity between various components.

Figure 2: Network Topology (Logical View) shown with 2 Dell PowerEdge R620 Servers



The table below summarizes the use of each traffic type.

Table 2: Traffic Type Summary

Traffic Type	Use
Management	Supports virtualization management traffic and communication between the ESXi servers in the cluster, and configuration and monitoring of the storage arrays and LAN switches.
vMotion	Supports migration of VMs between ESXi servers in the cluster.
VM	Supports communication between the VMs hosted on the ESXi cluster and external systems.
Out-of-Band Management (OOB)	Supports configuration and monitoring of the servers through the iDRAC management interface. In addition, the SAN switches can be managed on this network.
iSCSI Data	Supports iSCSI traffic between the servers and storage array(s).

Subsequent sections of this document provide more detailed connectivity information for these traffic types.

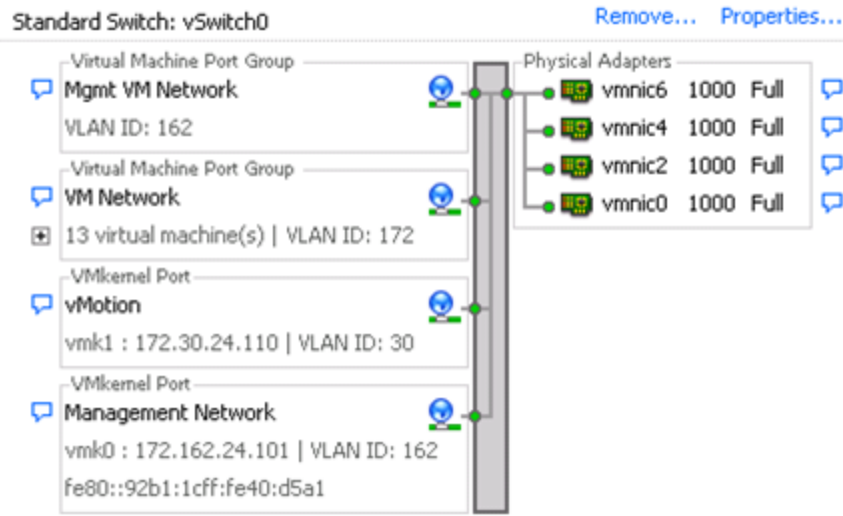
6.1 Server to LAN Connectivity

One of the key attributes of the Dell Active System 50v design is that there is no single point of failure and redundancy is incorporated into all mission critical components of the solution. Management applications, however, are not designed with this level of redundancy because the mission critical workloads continue to operate in the event of a management application failure. Network redundancy for the mission critical components is achieved with redundant network interface controllers (NICs) and redundant switches. iSCSI storage redundancy is achieved with redundant NICs, switches, and storage controllers. VMware High Availability (HA) provides virtual machine fail over.

Each Dell PowerEdge R620 rack server is configured with a Broadcom BCM5720 rack Network Daughter Card (rNDC) with four 1Gb ports and a Broadcom BCM5719 add-in network adapter providing four 1Gb ports. Two ports on each network adapter (four total) are connected to the LAN 7024 switches.

For each PowerEdge R620, a single ESXi virtual switch is created to support all LAN traffic (as showing in Figure 3). Unique port groups are then created for Management VMs and VM traffic. In addition, VMkernel interfaces are defined for management and vMotion traffic. Each port group and VMkernel interface is configured to tag with the appropriate VLAN. The Management port group is utilized by the Dell Management Plug-in and the EqualLogic VSM/VMWARE Plug-in for VMware vCenter.

Figure 3: LAN vSwitch Configuration for Compute Servers



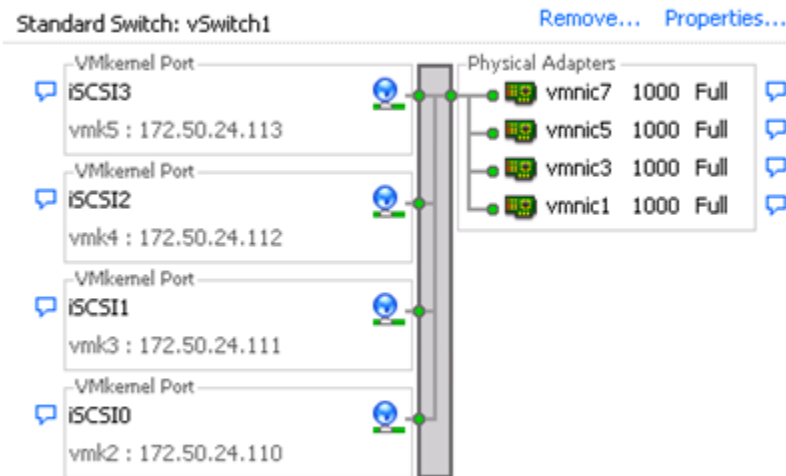
PowerEdge R420 rack management server is configured with a Broadcom BCM5720 rNDC with two 1Gb ports and a Broadcom BCM5720 add-in network adapter providing two 1Gb ports. One port on each network adapter (two total) are connected to the LAN 7024 switches. A single ESXi virtual switch is created to support all LAN traffic, similar to R620 compute servers.

Traffic Isolation using VLANs: The traffic on the LAN is separated into four VLANs: one VLAN each for Management, vMotion, VM, and Out-of-Band Management (OOB) traffic. VLAN tagging for the OOB traffic is performed by the switches.

6.2 Server to SAN Connectivity

The SAN includes two 7024 switches which support iSCSI traffic. The four remaining ports in each Dell PowerEdge R620 rack server are connected to the SAN 7024 switches. A single virtual switch is created to support all SAN traffic. The four 1Gb physical networks that are dedicated for iSCSI traffic are connected to this virtual switch (as shown in Figure 4). Four VMkernel interfaces are created and each VMkernel interface is associated with only one physical network port.

Figure 4: SAN vSwitch Configuration



The creation and configuration of the vSwitches and iSCSI VMkernel interfaces are accomplished through the EqualLogic Multipathing Extension Module (MEM) installation process, as discussed below.

EqualLogic Multipathing Extension Module: EqualLogic MEM provides advanced multipathing capabilities with path redundancy and improved performance by intelligently routing data. The EqualLogic MEM consists of two components: a Path Selection Plug-in (PSP) driver that is used by the VMware Native Multipathing Plug-in (NMP) to select the best path for each I/O to the EqualLogic storage devices and an EqualLogic Host Connection Manager (EHCM) that manages iSCSI sessions to the EqualLogic storage devices.

By utilizing the PSP in the EqualLogic MEM, storage I/O is improved with increased bandwidth and reduced network latency. This is because the PSP can use its knowledge about the exact distribution of volumes in the storage arrays. The EHCM provides automatic connection management by distributing iSCSI sessions across all configured adapters and PS Series group Ethernet ports. It also automates load balancing across multiple active paths to the storage array.

PowerEdge R420 management server SAN connectivity is similar to PowerEdge R620 servers, with the exception that the R420 uses two network ports for SAN connectivity.

6.3 Storage Architecture


EqualLogic PS6100 storage provides capabilities essential to the Active System 50v design, such as flexibility in configuring RAID arrays and creating volumes, thin provisioning, and storage tiering, while providing tight integration with VMware vSphere for better performance and manageability through the use of VSM and SAN HQ.

6.3.1 Storage Connectivity

The EqualLogic PS6100X contains redundant storage controllers. Each storage controller has four 1Gb connections supporting iSCSI data and a 100Mb dedicated management traffic port. The four iSCSI data connections on each controller are split between the two SAN switches for redundancy and support EqualLogic vertical port failover.

Figure 5 shows how the two controllers on storage array are connected.

Figure 5: PS6100 Array SAN Connectivity



6.3.2 EqualLogic Group and Pool Configuration

Each EqualLogic array (or member) is assigned to a particular group. Groups help in simplifying management by enabling management of all members in a group from a single interface. Each group contains one or more storage pools. Each pool must contain one or more members and each member is associated with only one storage pool.

The iSCSI volumes are created at the pool level. In the case where multiple members are placed in a single pool, the data is distributed amongst the members of the pool. With data being distributed over a larger number of disks, the potential performance of iSCSI volumes within the pool is increased with each member added.

6.3.3 Storage Array RAID Configuration

The storage array RAID configuration is highly dependent on the workload in your virtual environment. The EqualLogic PS series storage arrays support three RAID types: RAID 6, RAID 10, and RAID 50. The RAID configuration will depend on workloads and customer requirements. In general, RAID 10 provides the best performance at the expense of storage capacity.

RAID 10 generally provides better performance in random I/O situations, and requires additional overhead in the case of a drive failure scenario. RAID 50 generally provides more usable storage, but has less performance than RAID 10. RAID 6 provides better data protection than RAID 50.

6.3.4 Volume Size Considerations

Volumes are created in the storage pools. Volume sizes depend on the customer environment and the type of workloads. Volumes must be sized to accommodate not only the VM virtual hard drive, but also the size of the virtual memory of the VM and additional capacity for any snapshots of the VM.

It is important to include space for the guest operating system memory cache, snapshots, and VM configuration files when sizing these volumes. Additionally, one 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.

With each volume created and presented to the servers, additional iSCSI sessions are initiated. When planning the solution, it is important to understand that group and pool limits exist for the number of simultaneous iSCSI sessions. For more information, refer to the current EqualLogic Firmware (FW) Release Notes available at the [EqualLogic Support site](#) (requires registration).

6.3.5 Storage Access Control

Access to the created volumes can be restricted to a subset of the servers that have physical connectivity to the EqualLogic arrays. For each volume, access can be restricted based on IP address, iSCSI qualified name (IQN), and/or Challenge Handshake Authentication Protocol (CHAP). When creating a volume for the servers in the ESXi cluster, access must be granted to all servers in the cluster.

6.3.6 Jumbo Frames

In this Dell Active System configuration, jumbo frames are enabled for all devices of the iSCSI SAN fabric. This includes the server network interface ports, the network switch interfaces, and the Dell EqualLogic PS6100 Series interfaces.

6.3.7 Thin Provisioning

Particularly in virtualization environments, thin provisioning is a common practice. This allows for efficient use of the available storage capacity. The volume and corresponding partition may grow as needed, typically in an automated fashion to ensure availability of the volume. However, as storage becomes over-provisioned in this scenario, careful management and capacity planning is critical.

Dell EqualLogic PS6100 Series offers thin provisioning and eliminates pre-allocated but unused capacity.

7 Management Infrastructure

Active System 50v solution includes an optional Dell PowerEdge R420 server which is used to host the management virtual machines. The EqualLogic storage is shared between the management server and the compute cluster. The EqualLogic storage must be sized so that sufficient capacity and bandwidth are allocated for both the management VMs and compute VMs. The PowerEdge R420 server runs VMware ESXi hypervisor.

7.1 Management Components

The Active System 50v solution includes the necessary management components required to manage the solution infrastructure. The following management components are included in the Active System 50v solution.

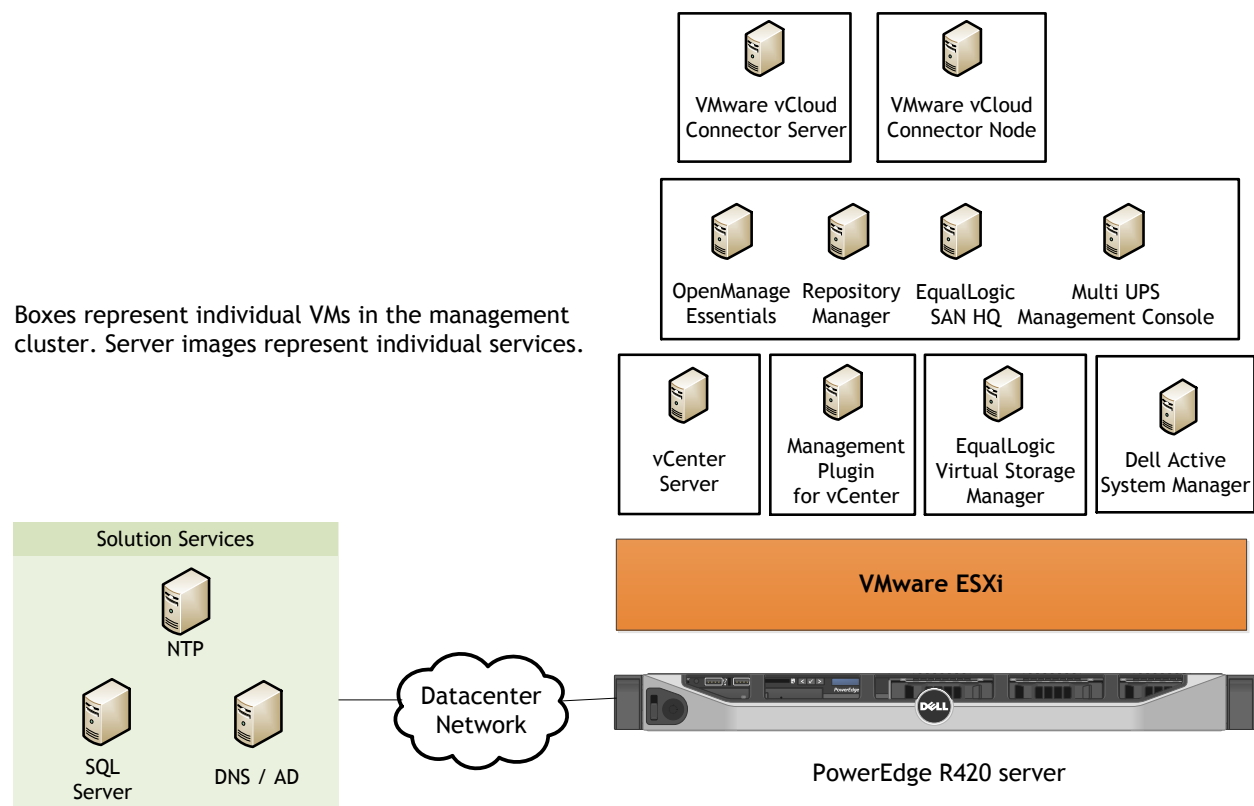
- Dell Active system Manager
- VMware vCenter Server
- Dell Management Plugin for vCenter
- Dell EqualLogic Virtual Storage Manager for VMware

Reference Architecture for Active System 50 with VMware vSphere

- Dell OpenManage Essentials
- Dell Repository Manager
- Dell EqualLogic SAN HQ.
- Dell Multi-UPS Management Console (MUMC)
- VMware vCloud™ Connector Server
- VMware vCloud™ Connector Node

These components are installed in virtual machines in the management infrastructure as illustrated in Figure 6.

Figure 6: Management Components



The remainder of this section will provide an introduction to each component and how they are integrated into the Active System 50v solution.

Dell Active System Manager: Dell Active System Manager is the Active Infrastructure management software that is part of the Active Systems. The Dell Active System Manager virtual appliance is deployed on the management cluster. For fullest functionality, direct internet access, or access through a proxy, is recommended.

Active System Manager addresses key factors that impact service levels, namely infrastructure configuration errors, incorrect problem troubleshooting, and slow recovery from failures. Active System

Manager dramatically improves the accuracy of infrastructure configuration by reducing manual touch points.

Dell OpenManage Essentials (OME): In the Active System 50v, Dell OpenManage Essentials (OME) is sized and configured to monitor the solution components. It is deployed on a Windows 2008 R2 virtual machine within the management cluster. High availability of the OME virtual machine is provided by VMware High Availability service. OME utilizes a local SQL Express database. For fullest OME functionality, direct internet access, or through a proxy, is recommended.

OME is utilized for discovery, inventory, and hardware level monitoring of rack servers, blade chassis, EqualLogic storage, Dell UPS, and network switches. Each of these components are configured to send SNMP traps to the centralized OME console to provide a “single pane of glass” monitoring interface for major hardware components. OME provides a comprehensive inventory of solution components through WS-MAN and SNMP inventory calls. For instance, reporting is available to provide blade and rack server firmware versions or solution warranty status. OME can be used as the single point of monitoring for all hardware components within an enterprise.

For more information on OpenManage Essentials, see the [Data Center Systems Management](#) page.

Dell Repository Manager (DRM): Within the Active System 50v solution, Dell Repository Manager (DRM) is installed on the same Windows 2008 R2 VM as Dell OpenManage Essentials. DRM is an application that allows IT Administrators to more easily manage system updates. DRM provides a searchable interface used to create custom collections known as bundles and repositories of Dell Update Packages (DUPs). These bundles and repositories allow for the deployment of multiple firmware, BIOS, driver, and software updates at once. Additionally, Dell Repository Manager makes it easier to locate specific updates for a particular platform, which saves you time. For example, in Repository Manager you can create a bundle with the latest updates for a Dell PowerEdge R620 or R420 server. DRM can be used in conjunction with other OpenManage tools to help ensure that your PowerEdge server is kept up to date.

For more information on Dell Repository Manager, see <http://content.dell.com/us/en/enterprise/d/solutions/repository-manager>.

Dell Management Plug-in for VMware vCenter (DMPVV): Dell Management Plug-in for VMware vCenter is deployed as a virtual appliance within the management cluster, and is attached to the VMware vCenter Server within the Active System 50v stack. DMPVV communicates with the VMware vCenter Server, the hypervisor management interfaces, and server out-of-band management interfaces (iDRAC). For ease of appliance, firmware updates, and warranty information, it is recommended that the DMPVV appliance has access to an internet connection either directly, or through a proxy. Dell Management Plug-in for VMware vCenter enables customers to:

- Get deep-level detail from Dell servers for inventory, monitoring, and alerting – all from within vCenter
- Apply BIOS and Firmware updates to Dell servers from within vCenter
- Automatically perform Dell-recommended vCenter actions based on Dell hardware alerts
- Access Dell hardware warranty information online
- Rapidly deploy new bare metal hosts using Profile features

For more information, see the web page for [Dell Management Plug-in for VMware vCenter](#).

Dell EqualLogic Virtual Storage Manager (VSM) for VMware: Within Active System 50v, the Dell EqualLogic Virtual Storage Manager (VSM) for VMware is deployed as a virtual appliance within the management cluster and is registered to the VMware vCenter Server. VSM communicates with the dedicated management interfaces of the EqualLogic storage enclosures over the out-of-band network. VSM enables customers to perform many storage administrative tasks from vSphere client including:

- Create Smart Copy snapshots, replicas, and clones of various types of VMware Infrastructure (VI) objects.
- Restore the state of virtual machines using saved Smart Copy snapshots and replicas.
- Setup replication of data stores and sets of data stores stored on one EqualLogic PS Series group to a secondary PS Series group (potentially at a remote location) for disaster tolerance.
- Recover from replicas on the secondary site, including failover and failback of virtual machines and their data.
- Create Virtual Desktop Infrastructure (VDI) manual desktop pools.
- Provision of data stores on EqualLogic iSCSI volumes.

Dell EqualLogic SAN HQ: Within the Active System 50v, Dell EqualLogic SAN HQ is installed on the same Windows 2008 R2 VM as OpenManage Essentials. SAN HQ communicates with the dedicated management interface of the EqualLogic storage enclosure to gather performance and event logs.

Dell EqualLogic SAN HQ provides consolidated performance and robust event monitoring across multiple groups. The key benefits of EqualLogic SAN HQ include:

- **Multi-Group Management:** EqualLogic SAN HQ enables centralized monitoring of multiple EqualLogic PS Series groups from a single graphical interface.
- **Comprehensive information about the EqualLogic PS Series arrays:** EqualLogic SAN HQ provides comprehensive information on configuration, capacity, I/O performance, and network performance for EqualLogic PS Series groups, pools, members, disks, volumes, and volume collections. These in depth analytical tools enable flexible, granular views of SAN resources and provide quick notification of hardware, capacity, and performance-related problems.
- **Experimental analysis:** EqualLogic SAN HQ collects information on current hardware configuration and distribution of reads and writes and provides information about EqualLogic PS Series group performance, relative to a specific workload. Customers can perform experimental analysis to determine if a group has reached its full capabilities, or whether they can increase the group workload with no impact on performance. This helps in identifying requirements for storage growth and future planning.
- **Events and alerts:** EqualLogic SAN HQ provides performance related and email alerts and hardware alarms on multiple parameters. This feature ensures users take timely action to make data more available and more secure.
- **Formatted reports, graphs, and archives:** Customizable reports and graphs are available on performance, capacity utilization and trending, group configuration with alerts, replication, status, host connections, and more.

Dell Multi-UPS Management Console (MUMC): Designed for multi-host server environments, Dell's console facilitates easy and versatile monitoring and management across the network from a single interface. MUMC also provides agentless integration into multi-host managers. Other features include:

- Ability to monitor and manage multiple power and environmental devices from a Web browser
- Allows virtual machine shutdown for VMware® high availability (HA) clusters using Remote Virtual Machine Shutdown (RVMS)
- Creates virtual composite groups of UPSs to manage redundant installations as one device
- User-definable tree structure enables grouping, access, and management of multiple devices across multiple locations
- Agentless control of multi-host VMware vCenter™ and Citrix® XenServer® installations
- Auto discovery of Dell UPS and Dell intelligent PDU systems on the network
- Power device asset management tracks moves, adds, and changes of equipment
- All the functionality of an Enterprise-class monitoring solution is included at no charge or at a fraction of the cost
- Support for up to 32 UPS and PDU systems is included at no charge

For more information, see the web page for [Dell UPS Management Software](#).

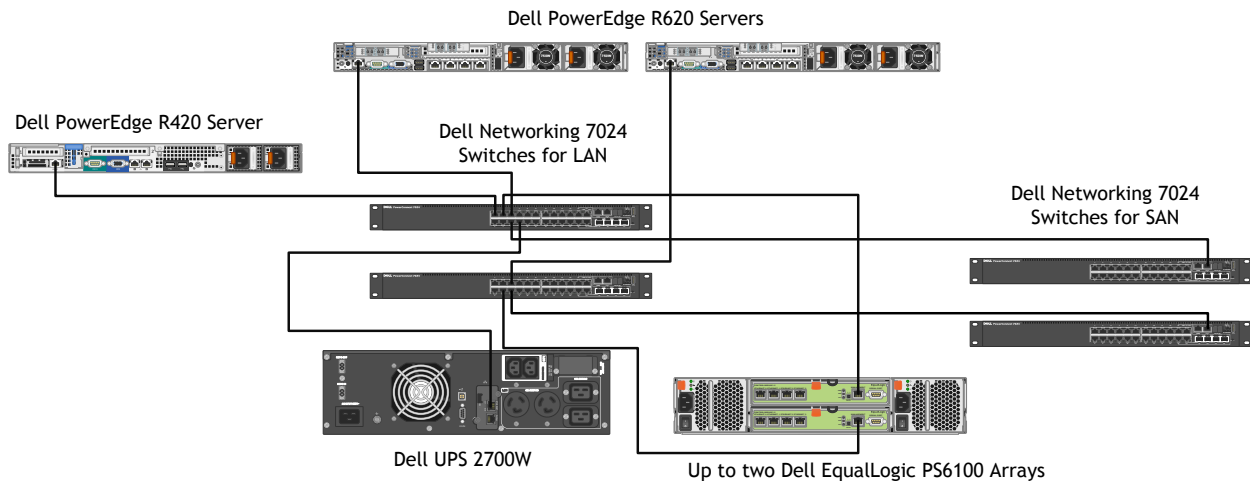
VMware vCloud Connector: VMware vCloud Connector is an optional component of the Active System 50v solution. When included, it is deployed upon the management stack, alongside other management VMs. For the base functionality, three VMs are necessary, a single 'server' VM and two 'node' VMs. The node VMs are responsible for the physical transfer of VM workloads. Within the Active System 50v, two of these components, the server and the local node, are installed. The third component, 'remote' node VM, should be installed outside of the Active System 50v solution, near the infrastructure to which it provides connectivity.

After deploying the VMware vCloud Connector 'node' VMs, the size of the virtual disk may have to be increased based on the size of expected VM to be transferred and the number of concurrent transfers anticipated.

7.2 Out-Of-Band Management Connectivity

The Dell Networking 7024 LAN switches are used as out-of-band management switches. Each of the Active System 50 component OOB management ports are connected to the 7024 switches, as shown in Figure 7.

Figure 7: Connectivity of OOB management components.



8 Connecting to Datacenter Network

It's likely that an Active System 50 is connected into a data center infrastructure that consists of Dell Networking switches, Cisco switches, or those of some other vendor.

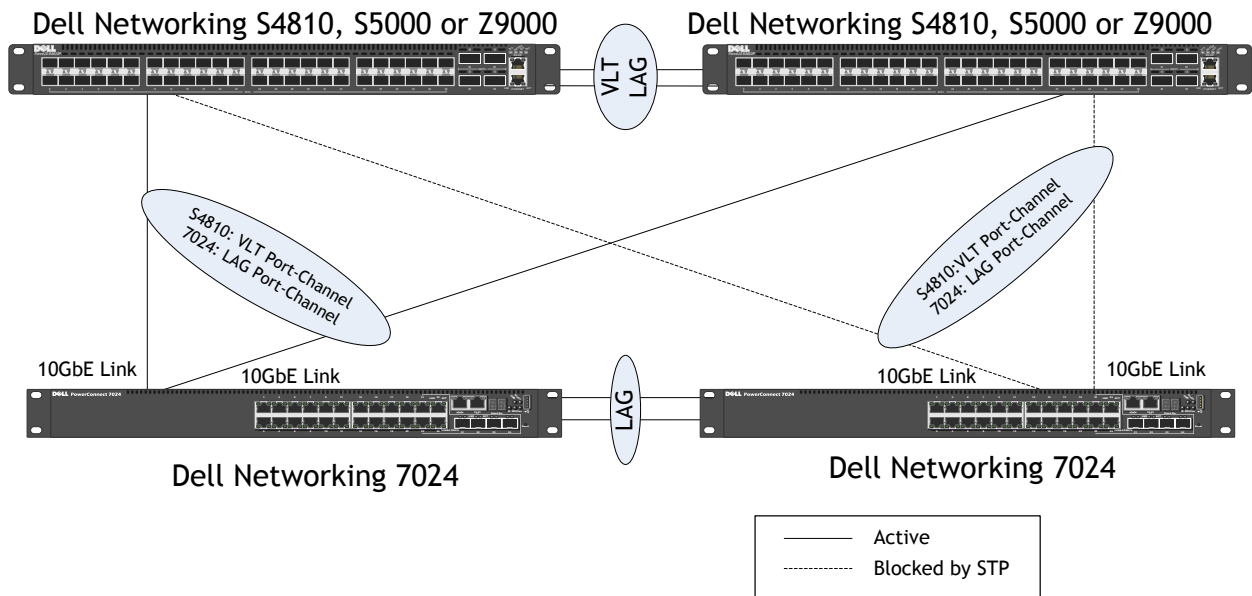
There are several ways to uplink the switches to the core network. Multiple uplinks can be created on each switch connecting to the core network switches. The solution provides four 1Gb uplinks from each switch to link into an existing core network infrastructure. If the core network infrastructure supports 10Gb Ethernet, then 10Gb uplink modules must be added to each switch.

In this section, we provide examples to show how Active System is connected to the Datacenter network.

8.1 Connecting to a Datacenter Network with Dell Networking Switches

Figure 8 below illustrates an example with Dell Networking S-series/Z-series 10Gb switches. The two Dell Networking S-series/Z-series switches can be connected together using VLTi. Each of the 7024 switches in the Active System use a two 10GbE port LAG connecting to the two S-Series/Z-series switches. One of the LAG will be active, while the other will be blocked by STP. Multiple 1GbE ports may also be used.

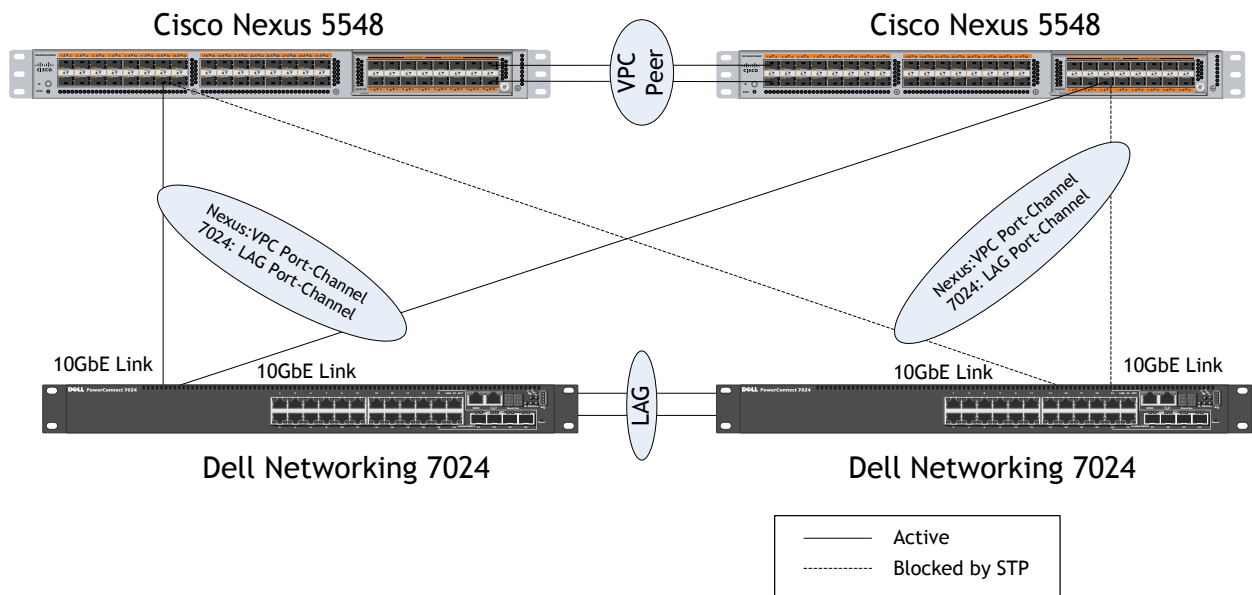
Figure 8: Connectivity to a Dell Network



8.2 Connecting to a Cisco Datacenter Network

Figure 9 below illustrates an example with Cisco switches. The two Cisco switches can be connected together using VPC. Each 7024 switches in the Active System use a two 10GbE port LAG connecting to the two Cisco switches. One of the LAG will be active, while the other will be blocked by STP. Multiple 1GbE ports may also be used.

Figure 9: Connectivity to Cisco Network



9 Delivery Model

The Dell Active System 50v reference architecture can be purchased as a complete solution. This solution is available to be racked, cabled, and delivered to the customer site, to speed deployment. Dell Services deploys and configures the solution tailored to the business needs of the customer and the architecture is developed and validated by Dell Engineering. For more details or questions about the delivery model, please consult with your Dell Sales representative.

Figure 10 below shows the Dell Active System 50v solution in the two compute host configuration with two EqualLogic PS6100 storage arrays.

Note that switches are shown mounted forward for representation. In actual use, ports face the back of the rack. Additional PDUs are utilized within the rack.

Figure 10: Active System 50v Rack and Component Overview

(1) KMM (optional)

(2) PC7024 LAN

(2) PC7024 SAN

(1) R420 Mgmt Host (optional)

(2) R620 Compute

(Up to 2) PS6100 2U, 4U
or mixed 2U and 4U


(2) 2700W UPS (optional)

(2) PDU

(2) C19-C20 power cable

(2) IEC309-16 power cable

(2) L6-20P power cable



(2) CX4 Module

(2) CX4 Module

10 Reference

VMware references:

- [VMware vSphere Edition Comparisons](#)
- [VMware vSphere Compatibility Matrixes](#)
- [VMware High Availability \(HA\): Deployment Best Practices](#)
- [VMware Virtual Networking Concepts](#)

Dell EqualLogic references:

- [EqualLogic Technical Content](#)
- [Dell EqualLogic PS Series Architecture Whitepaper](#)
- [Configuring iSCSI Connectivity with VMware vSphere 5 and Dell EqualLogic PS Series Storage](#)
- [Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 5.1, 5.0 and 4.1 and PS Series SANs](#)
- [How to Select the Correct RAID for an EqualLogic SAN](#)
- [Using Tiered Storage in a PS Series SAN](#)
- [Monitoring your PS Series SAN with SAN HQ](#)

Dell Management reference:

- [Dell Management Plug-In for VMware vCenter references - Solution Brief](#)