# Reference Architecture for Active System 800 with VMware vSphere

*Release 1.1 for Dell PowerEdge 12th Generation Blade Servers, Dell Force10 Switches, Dell EqualLogic iSCSI SAN and Dell Active System Manager*

**Dell Virtualization Solutions Engineering**

**Revision: A02**

## Revision History

| Revision | Description |
|----------|-------------|
| A00 | Initial Version |
| A01 | Updated references |
| A02 | Updated figures. Support for Active System Manager |

# Contents

# Figures

## Tables

# 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 (Active System Manager), converged LAN/SAN fabrics, and modular server architecture for the ultimate converged infrastructure solution.  Active Infrastructure helps IT rapidly respond to dynamic business demands, maximize data center efficiency, and strengthen IT service quality.

The Active System 800 solution, a member of Dell Active Infrastructure family, is a converged infrastructure solution that has been designed and validated by Dell Engineering. It is available to be racked, cabled, and delivered to your site to speed deployment. Dell Services will deploy and configure the solution tailored for business needs, so that the solution is ready to be integrated into your datacenter. Active System 800 is offered in configurations with either VMware® vSphere® (Active System 800v) or Microsoft® Windows Server® 2012 with Hyper-V® role enabled (Active System 800m) hypervisors. This paper defines the Reference Architecture for the VMware vSphere based Active System 800v solution.

Active System 800v offers converged LAN & SAN fabric design to enable a converged infrastructure solution. The end-to-end converged network architecture in Active System 800v is based upon Data Center Bridging (DCB) technologies that enable convergence of all LAN and iSCSI SAN traffic into a single fabric. The converged fabric design of Active System 800v reduces complexity and cost while bringing greater flexibility to the infrastructure solution.

Active System 800v includes Dell PowerEdge™ M1000e blade chassis with Dell PowerEdge™ M I/O Aggregator, Dell PowerEdge™ M620 blades, Dell EqualLogic™ Storage, Dell Force10™ network switches, and VMware vSphere 5.1. The solution also includes Dell PowerEdge™ R620 servers as management servers. Dell Active System Manager, VMware vCenter Server, EqualLogic Virtual Storage Manager for VMware, and Dell OpenManage™ Essentials are included with the solution.

One of the key components of Active System 800v is Dell Active System Manager. Active System Manager simplifies complex and error-prone infrastructure lifecycle management activities like discovery, inventory, deployment, configuration, and on-going monitoring and management through automation and collapsing the management interfaces into a highly optimized guided workflow. By simplifying and automating these activities through a wizard-driven graphical user interface, Dell Active System manager enables IT to respond rapidly to business needs, maximize data center efficiency, and strengthen quality of IT service delivery.

# 2   Audience

IT administrators and IT managers — who have purchased, or are planning to purchase an Active System configuration — can use this document to understand the design elements, hardware and software components, and the overall architecture of the solution.

# 3   Solution Overview

This section provides a high-level product overview of VMware vSphere, Dell PowerEdge blade servers, Dell PowerEdge M I/O Aggregator, Dell Force10 S4810 switch, Dell Force10 S55 switch, and Dell

EqualLogic Storage, as illustrated in Figure 1. Readers can skip the sections of products with which they are familiar.

Figure 1: Active System 800v Overview

**VMware vSphere 5.1 Hypervisor**
- vMotion, Storage vMotion
- VMware HA and DRS

**Dell PowerEdge Blade Servers for Compute Cluster**
- Energy efficient PowerEdge M1000e enclosure
- 12th generation M620 blade server
- Flex Address
- CMC and iKVM for enclosure management

**Dell PowerEdge M I/O Aggregator**
- Highest Performance in a Blade Switch
- Highest Density in a Single Blade Switch
- Stackable for Simplified Management
- Scalability & Modular to Fit Your Business
- Support for converged networking with Data Center Bridging (DCB)

**Dell PowerEdge Rack Servers for Management Cluster**
- 12th generation R620 rack servers
- Concentrated computing power in 1U form factor
- Large memory and I/O capacity
- Powerful systems management with Dell iDRAC and Lifecycle Controller
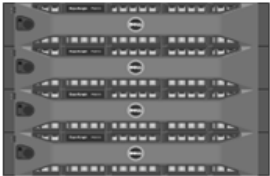
**Dell Force10 S4810 Switches for Converged Network**
- High-density 48-port 10 GbE switch with four 40 GbE uplinks
- Ultra-low-latency, non-blocking, cut-through switch for line-rate L2 and L3 performance
- Integrated network automation and virtualization tools via the Open Automation Framework
- Support for converged networking with Data Center Bridging (DCB)

**Dell Force10 S55 Switch for Management**
- High-density 48-port 1/10 GbE scalable switch
- Low-latency, non-blocking switch for line-rate L2 and L3 performance
- Integrated network automation and virtualization tools via the Open Automation Framework

**Dell EqualLogic Storage**
- 10GbE iSCSI SAN arrays with SFP+ and 10GBase-T support
- Thin Provisioning and Storage Tiering
- Support for converged networking with Data Center Bridging (DCB)
- Integration with VMware

**Integrated Management**
- 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 (HQ)
- Dell Repository Manager

**Cloud Enablement**
- VMware vCloud Connector for Dell vCloud connectivity

Reference Architecture for Active System 800 with VMware vSphere

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

Table 1: Solution Components

| Component | Details |
|---|---|
| Hypervisor Server | • Up to 2x Dell PowerEdge M1000e chassis with up to 32x Dell PowerEdge M620 Blade Servers and embedded VMware vSphere 5.1 |
| Converged Fabric Switch | • 2xDell Force10 S4810<br>• 2x Dell PowerEdge M I/O Aggregator in each Dell PowerEdge M1000e chassis |
| Storage | • Up to 8x Dell EqualLogic PS6110 series arrays |
| Management Infrastructure | • 2x Dell PowerEdge R620 servers with embedded VMware vSphere 5.1 hosting management VMs.<br>• 1x Dell Force10 S55 used as a 1Gb out-of-band management switch |
| Management components hosted in the management infrastructure | • Dell Active System Manager<br>• VMware vCenter Server<br>• Dell Management Plug-in for VMware vCenter<br>• Dell OpenManage Essentials<br>• Dell EqualLogic Virtual Storage Manager (VSM) for VMware<br>• Dell EqualLogic SAN HeadQuarters (HQ)<br>• VMware vCloud Connector<br>• Dell Repository Manager |

**VMware vSphere 5.1:** VMware vSphere 5.1 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 (VM) level. Upon host failure, VMware HA automatically re-starts VMs on other physical hosts running ESXi. VMware vSphere 5.1 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.

Reference Architecture for Active System 800 with VMware vSphere

- **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 & operational steps

- Infrastructure Lifecycle Mgmt — Discovery, inventory, configuration, provisioning, and ongoing management of physical and virtual infrastructure

- Resource Pooling & 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 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.  A high-speed passive mid-plane connects the server modules to the I/O modules, management, and power in the rear of the chassis. The enclosure includes a flip-out LCD screen (for local configuration), six hot-pluggable/redundant power supplies, and nine hot-pluggable N+1 redundant fan modules.

**Dell PowerEdge Blade Servers:** The PowerEdge M620 blade server is the Dell 12<sup>th</sup> generation PowerEdge half height blade 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 PowerEdge server. Each PowerEdge M620 can deploy up to 24x 32GB DIMMs, or 768GB of RAM per blade – 12TB of RAM in a single M1000e chassis.

- '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 Select Network Adapter (formerly NDC) on the PowerEdge M620 offers three modular choices for embedded fabric capability. With 10Gb CNA offerings from Broadcom, QLogic & Intel, our 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. The Broadcom and QLogic offerings offer Switch Independent Partitioning technology, developed in partnership with Dell, which allows for virtual partitioning of the 10Gb ports.

**I/O Modules:** The Dell blade chassis has three separate fabrics referred to as A, B, and C. Each fabric can have two I/O modules, for a total of six I/O module slots 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, an Ethernet pass-through module, FC switch, or FC pass-through module. InfiniBand™ switch modules are also supported. Each half-height blade server has a dual-port network daughter card (NDC) and two optional dual-port mezzanine I/O cards. The NDC connects to Fabric A. One mezzanine I/O card attaches to Fabric B, with the remaining mezzanine I/O card attached to Fabric C.

**Chassis 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 the CMC, the enclosure supports FlexAddress Plus technology, which enables the blade enclosure to lock the World Wide Names (WWN) of the FC controllers and Media Access Control (MAC) addresses of the Ethernet controllers to specific blade slots. This enables seamless swapping or upgrading of blade servers without affecting the LAN or SAN configuration.

**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 blade 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 1-to-1 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 and blade servers, see http://content.dell.com/us/en/enterprise/dcsm-embedded-management and Dell.com/blades.

**Dell PowerEdge M I/O Aggregator:** The Dell PowerEdge M I/O Aggregator (M I/OA) is a flexible 1/10GbE aggregation device that is automated and pre-configured for easy deployment into converged iSCSI and FCoE networks. The key feature of the PowerEdge M I/OA is that all VLANs are allowed as a default setting. This allows the top-of-rack (ToR) managed switch to perform all VLAN management related tasks. The external ports of the PowerEdge M I/OA are automatically all part of a single link aggregation group (LAG), and thus there is no need for Spanning-tree. The PowerEdge M I/OA can use Data Center Bridging (DCB) and Data Center Bridging Exchange (DCBX) to support converged network architecture.

The PowerEdge M I/OA provides connectivity to the CNA/Network adapters internally and externally to upstream network devices. Internally the PowerEdge M I/OA provides thirty-two (32) connections. The connections are 10 Gigabit Ethernet connections for basic Ethernet traffic, iSCSI storage traffic, or FCoE storage traffic. In a typical PowerEdge M1000e configuration with 16 half-height blade server ports, 1-16 are used and 17-32 are disabled. If quad port CAN/Network adapters or quarter-height blade servers are used, then ports 17-32 will be enabled.

The PowerEdge M I/OA includes two integrated 40Gb Ethernet ports on the base module. These ports can be used in a default configuration with a 4 X 10Gb breakout cable to provide four 10Gb links for network traffic. Alternatively these ports can be used as 40Gb links for stacking. The Dell PowerEdge M I/OA also supports three different types of add-in expansion modules, which are called FlexIO Expansion modules. The modules available are: 4-port 10Gbase-T FlexIO module, 4-port 10G SFP+ FlexIO module, and the 2-port 40G QSFP+ FlexIO module.

The PowerEdge M I/OA modules can be managed through the PowerEdge M1000e Chassis Management Controller (CMC) GUI. Also, the out-of-band management port on the PowerEdge M I/OA is reached by connection through the CMC's management port. This one management port on the CMC allows for management connections to all I/O modules within the PowerEdge M1000e chassis.

For more information on Dell PowerEdge M I/O Aggregator, see http://www.dell.com/us/business/p/poweredge-m-io-aggregator/pd

**Dell Force10 S4810 Switches**: The Force10 S-Series S4810 is an ultra-low-latency 10/40 GbE ToR switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 switch delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact Force10 S4810 design provides industry leading density of 48 dual-speed 1/10 GbE (SFP+) ports, as well as four 40GbE QSFP+ uplinks to conserve valuable rack space and simplify the migration to 40Gbps in the data center core. (Each 40GbE QSFP+ uplink can support four 10GbE ports with a breakout cable).

Powerful Quality of Service (QoS) features coupled with Data Center Bridging (DCB) support to make the Force10 S4810 ideally suited for iSCSI storage environments. In addition, the S4810 incorporates multiple architectural features that optimize data center network flexibility, efficiency, and availability, including Force10's stacking technology, reversible front-to-back or back-to-front airflow

for hot/cold aisle environments, and redundant, hot-swappable power supplies and fans. For more information on Force10 switches, see Dell.com/force10.

**Dell Force10 S55 Switches:** The Dell Force10 S-Series S55 1/10 GbE ToR switch is designed for high-performance data center applications. The S55 leverages a non-blocking architecture that delivers line-rate, low-latency L2 and L3 switching to eliminate network bottlenecks.  The high-density Force10 S55 design provides 48GbE access ports with up to four modular 10GbE uplinks in 1-RU to conserve valuable rack space. The Force10 S55 switch incorporates multiple architectural features that optimize data center network efficiency and reliability, including reversible front-to-back or back-to-front airflow for hot/cold aisle environments and redundant, hot-swappable power supplies and fans. For more information on Force10 switches, see Dell.com/force10.

**Dell EqualLogic PS6110 Series Storage**: The Dell EqualLogic PS6110 series arrays are 10GbE iSCSI SAN arrays. The EqualLogic PS6110 arrays provide 10GbE connectivity using SPF+ or lower-cost 10GBASE-T. A dedicated management port allows better utilization of the 10GbE ports for the storage network I/O traffic by segmenting the management traffic. The PS6110 Series 10GbE arrays can use  DCB to improve Ethernet quality of service and greatly reduce dropped packets for an end-to-end iSCSI over DCB solution, from host adapters to iSCSI target.

The key features of the EqualLogic PS6110 series arrays are:

- Dedicated 10GbE ports that enable you to use SFP+ or 10GBASE-T cabling options

- Simplified network storage management with a dedicated management port

- 2.5" drives in 2U or 3.5" drives in 4U form factors

- SAS, NL-SAS and solid state drive and hybrid options available

- Supports DCB and DCBX technologies for use in a converged LAN & iSCSI SAN network

- Efficient data protection and simplified management and operation of the EqualLogic SAN through tight integration with Microsoft®, VMware® and Linux® host operating platforms

- Includes a full-featured array monitoring and analysis tool to help strengthen your ability to analyze and optimize storage performance and resource allocation

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

**Dell PowerEdge R620 Management Server:** The Dell PowerEdge R620 uses Intel Xeon E5-2600 series processors and Intel chipset architecture in a 1U rack mount form factor. These servers support up to ten 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 24 memory DIMM slots. For more information, see the PowerEdge R620 guides at Dell.com/PowerEdge.

**Dell OpenManage™ Essentials (OME):** The Dell OpenManage™ Essentials 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.

OME 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 PowerConnect™ 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 OME, see the Data Center Systems Management page.

**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 Cloud Connectivity using 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 Dell vCloud public cloud.

- Expand your view across hybrid clouds. Use a "single pane of glass" management interface that seamlessly spans your private vSphere and public Dell vCloud environment.

- Extend your datacenter.  Move VMs, vApps, and templates from private vSphere to a Dell vCloud to free up your on-premise datacenter resources as needed.

- Consume cloud resources with confidence.  Run Development, QA, and production workloads using Dell vCloud, a VMware technology-based public cloud.

The Dell Cloud with VMware vCloud™ Datacenter is an enterprise-class, multi-tenant infrastructure-as-a-service (IaaS) public cloud solution that is hosted in secured Dell data centers. Utilizing VMware vCloud Connector, Dell Cloud provides you with unique hybrid cloud capabilities to extend your internal data center with Dell and VMware by transitioning your VMware virtualized workloads into our vCloud data center. vCloud hosting provides you with a secure, manageable, and flexible public cloud application. For more information, see Dell vCloud website.

# 4  Design Principles

The following principles are central to the design and architecture of Active System 800v Solution.

1. **Converged Network:** The infrastructure is designed to achieve end-to-end LAN and SAN convergence.

2. **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

3. **Management:** Provide integrated management using VMware vCenter, Dell Management Plug-in for VMware vCenter, Dell OpenManage Essentials, and EqualLogic Virtual Storage Manager (VSM) for VMware plug-in.

4. **Cloud Enabled:** The solution also includes connectivity to Dell vCloud using VMware vCloud Connector.

5. **Hardware configuration for virtualization:** This solution is designed for virtualization for most general cases. Each blade server is configured with appropriate processor, memory, and network adapters, as required for virtualization.

6. **Racked, Cabled, and Ready to be deployed:** Active System 800v is available racked, cabled, and delivered to the customer site, ready for deployment. Components are configured and racked to optimize airflow and thermals. Based on customer needs, different rack sizes and configurations are available to support various datacenter requirements.

7. **Flexible configurations:** Active System 800v is pre-configured to suit most customer needs for a virtualized infrastructure. The solution also supports additional options, such as configuring racks, server processors, server memory, and storage, based on customer needs.

# 5 Prerequisites and Datacenter Planning

**Power, Cooling, and Weight Considerations:** Dell Active System 800v solution is configured with Power Distribution Units (PDUs) to meet the power requirements of the components, as well as regional constraints. Power consumed, cooling required, and information regarding rack weight are provided to enable customers to plan appropriate power and cooling for the solution.

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

1. An existing Ethernet infrastructure with which to integrate. 10Gb or 40Gb Ethernet infrastructure is recommended.
2. Additional components, such as Dell Force10 network cables and transceivers, are needed to uplink the solution to the customer network. The necessary components depend upon customer networking and uplink requirements.

3. Sufficient power and cooling to support the solution must be present. Detailed power, weight, and cooling requirements for the datacenter are defined in the Specification Guide for Active System 800 with VMware vSphere.

# 6   Architecture

This solution consists of a PowerEdge M1000e chassis populated with PowerEdge M620 blade servers running VMware ESXi. Figure 2 provides the high-level reference architecture for the solution.

<p align="center">Figure 2: Active System 800v Network Topology (Logical View)</p>
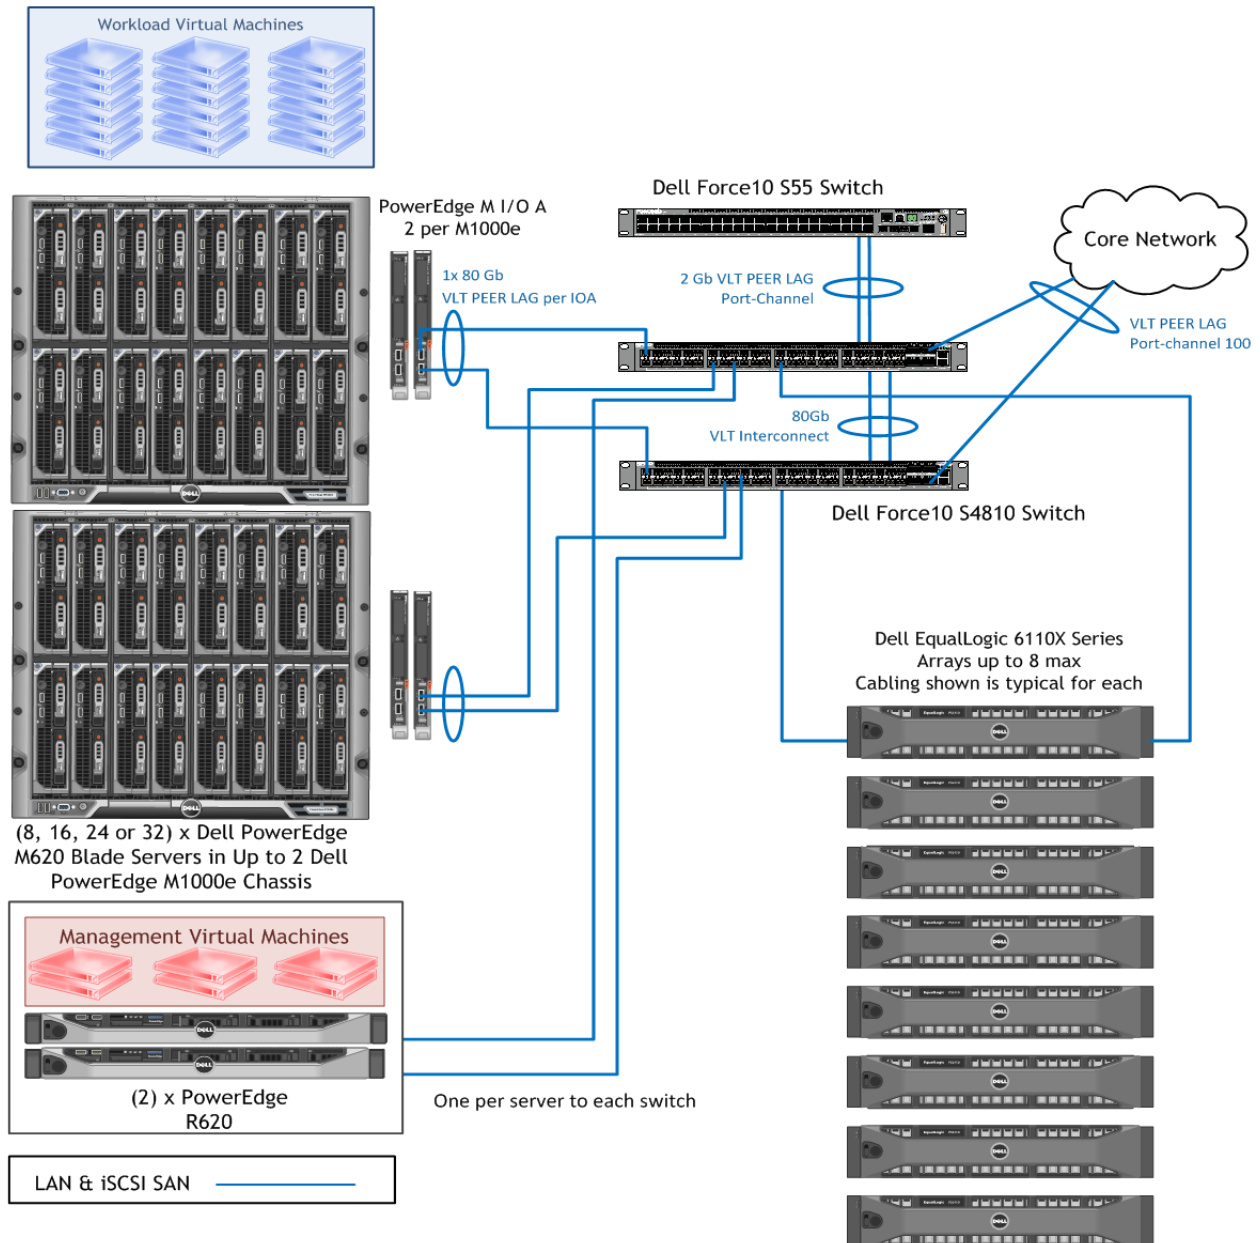


  The Figure 2 shows high-level logical connectivity between various components. Subsequent sections of this document provide more detailed connectivity information.

# 7 Dell Blade Network Architecture

In Active System 800v, the Fabric A in PowerEdge M1000e blade chassis contains two Dell PowerEdge M I/O Aggregator modules, one in I/O module slot A1 and the other in slot A2, and is used for converged LAN and SAN traffic. Fabric B and Fabric C (I/O Module slot B1, B2, C1, and C2) are not used.

The PowerEdge M620 blade servers use the Broadcom 57810-k Dual port 10GbE KR Blade NDC to connect to the Fabric A. Dell PowerEdge M I/O Aggregator modules uplink to Dell Force10 S4810 network switches providing LAN AND SAN connectivity.

Figure 3 below illustrates how the fabrics are populated in the PowerEdge M1000e blade server chassis and how the I/O modules are utilized.

Figure 3: I/O Connectivity for PowerEdge M620 Blade Server
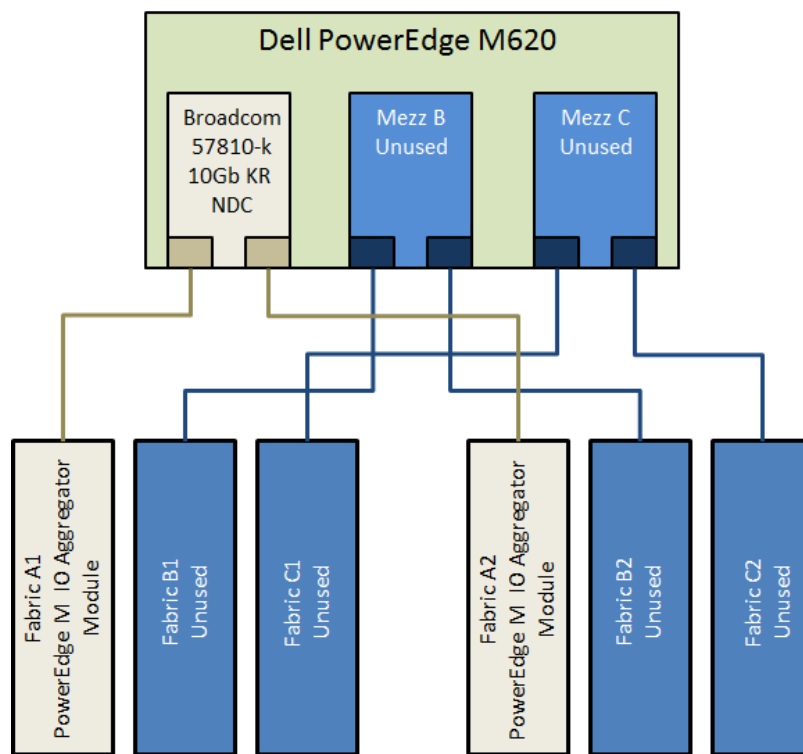


**Network Interface Card Partition (NPAR):** NPAR allows splitting the 10GbE pipe on the NDC with no specific configuration requirements in the switches. With NPAR, administrators can split each 10GbE port of an NDC into four separate partitions, or physical functions, and allocate the desired bandwidth and resources as needed. Each of these partitions is enumerated as a PCI Express function that appears as a separate physical NIC in the server, operating systems, BIOS, and hypervisor. Active System 800v solution takes advantage of NPAR. Partitions are created for various traffic types and bandwidth is allocated, as described in the following section.

## 7.1  Converged Network Architecture

One of the key attributes of the Active System 800v is the convergence of SAN and LAN over the same network infrastructure. LAN and iSCSI SAN traffic share the same physical connections from servers to storage. The converged network is designed using Data Center Bridging (IEEE 802.1) and Data Center Bridging Exchange (IEEE 802.1AB) technologies and features. The converged network design drastically reduces cost and complexity by reducing the components and physical connections and the associated efforts in deploying, configuring, and managing the infrastructure.

Data Center Bridging is a set of related standards to achieve enhance Ethernet capabilities, especially in datacenter environments, through converge network connectivity. The functionalities provided by DCB and DCBX are:

- **Priority Flow Control (PFC):** This capability provides zero packet loss under congestion by providing a link level flow control mechanism that can be controlled independently for each priority.

- **Enhanced Transmission Selection (ETS):** This capability provides a framework and mechanism for bandwidth management for different traffic types by assigning bandwidth to different frame priorities.

- **Data Center Bridging Exchange (DCBX):** This functionality is used for conveying the capabilities and configuration of the above features between neighbors to ensure consistent configuration across the network.

Dell Force10 S4810 switches, Dell PowerEdge M I/O Aggregator modules, Broadcom 57810-k Dual port 10GbE KR Blade NDCs, and EqualLogic PS6110 iSCSI SAN arrays enable Active System 800v to utilize these technologies, features, and capabilities to support converged network architecture.

### 7.1.1  Converged Network Connectivity

The Active System 800v design is based upon a converged network. All LAN and iSCSI traffic within the solution share the same physical connections. The following section describes the converged network architecture of Active System 800v.

**Connectivity between hypervisor hosts and converged network switches:** The compute cluster hypervisor hosts, PowerEdge M620 blade servers, connect to the Force10 S4810 switches through the PowerEdge M I/O Aggregator I/O Modules in the PowerEdge M1000e blade chassis. The management cluster hypervisor hosts, PowerEdge R620 rack servers, directly connect to the Force10 S4810 switches.

- **Connectivity between the Dell PowerEdge M620 blade servers and Dell PowerEdge M I/O Aggregators:** The internal architecture of PowerEdge M1000e chassis provides connectivity between the Broadcom 57810-k Dual port 10GbE KR Blade NDC in each PowerEdge M620 blade server and the internal ports of the PowerEdge M I/O Aggregator. The PowerEdge M I/O Aggregator has 32 x 10GbE internal ports. With one Broadcom 57810-k Dual port 10GbE KR Blade NDC in each PowerEdge M620 blade, blade servers 1-16 connect to the internal ports 1-16 of each of the two PowerEdge M I/O Aggregator. Internal ports 17-32 of each PowerEdge M I/O Aggregator are disabled and not used.

- **Connectivity between the Dell PowerEdge M I/O Aggregator and Force10 S4810 switches:** The two PowerEdge M I/O Aggregator modules are configured to operate as a port aggregator for aggregating 16 internal ports to eight external ports.

  The two fixed 40GbE QSFP+ ports on each PowerEdge M I/O Aggregator are used for network connectivity to the two Force10 S4810 switches. These two 40GbE ports on each PowerEdge M I/O Aggregator are used with a 4 x 10Gb breakout cable to provide four 10Gb links for network traffic from each 40GbE port. Out of the 4 x 10Gb links from each 40GbE port on each PowerEdge M I/O Aggregator, two links connect to one of the Force10 S4810 switches and the other two links connect to the other Force10 S4810 switch. Due to this design, each PowerEdge M1000e chassis with two PowerEdge M I/O Aggregator modules will have total of 16 x 10Gb links to the two Force10 S4810 switches. This design ensures load balancing while maintaining redundancy.

- **Connectivity between the Dell PowerEdge R620 rack servers and Force10 S4810 switches:** Both of the PowerEdge R620 servers have two 10Gb connections to the Force10 S4810 switches through one Broadcom 57810 Dual Port 10Gb Network Adapter in each of the PowerEdge R620 servers.

**Connectivity between the two converged network switches:** The two Force10 S4810 switches are connected using Inter Switch Links (ISLs) using two 40 Gbps QSFP+ links. Virtual Link Trunking (VLT) is configured between the two Force10 S4810 switches. This design eliminates the need for Spanning Tree-based networks; and also provides redundancy as well as active-active full bandwidth utilization on all links.

**Connectivity between the converged network switches and iSCSI storage arrays:** Each EqualLogic PS6110 array in Active System 800v uses two controllers. The 10Gb SFP+ port on each EqualLogic controller is connected to the Force10 S4810 switches. This dual controller configuration provides high availability and load balancing.

Figure 4 below illustrates the resultant logical converged network connectivity within the Active System 800v solution.

Figure 4: Converged Network Logical Connectivity
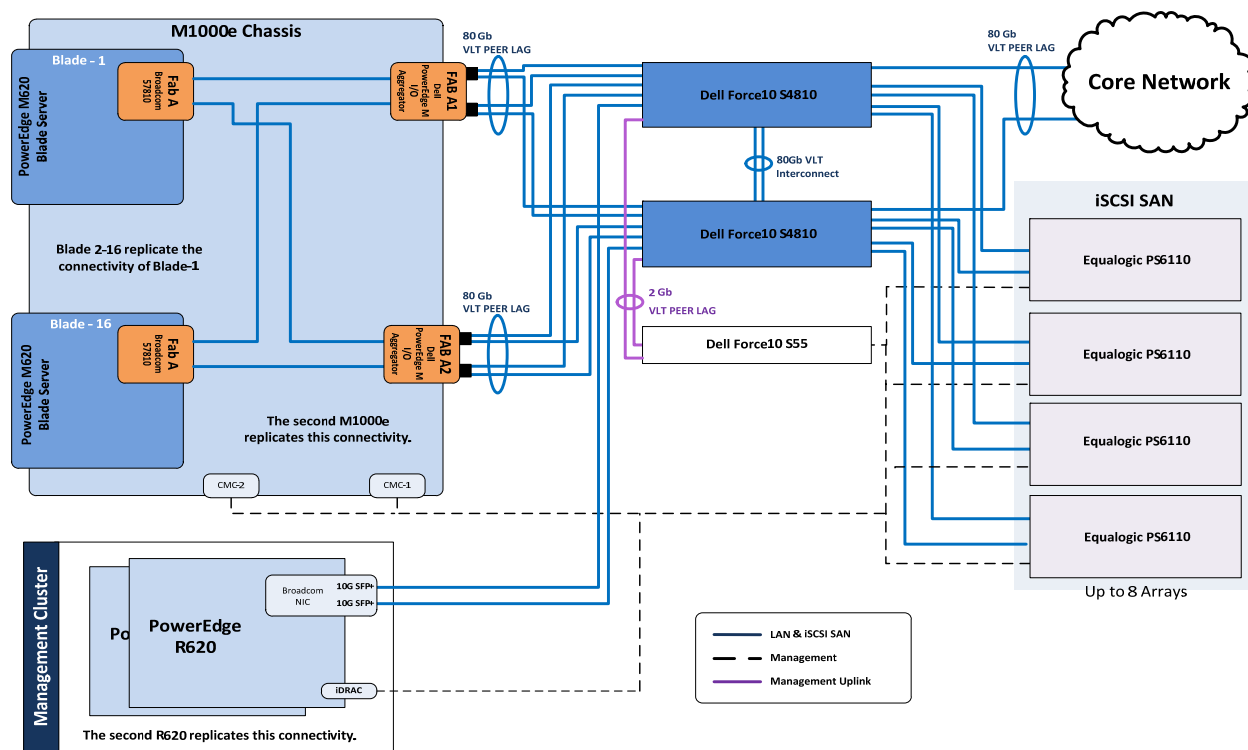


## 7.1.2   Converged Network Configuration

This section provides details of the different configurations in the Active System 800v that enable the converged network in the solution.

**DCB Configuration:** Data Center Bridging (DCB) and Data Center Bridging Exchange (DCBX) technologies are used in Active System 800v to enable converged networking. The Force10 S4810 switches, PowerEdge M I/O Aggregator modules, Broadcom 57810-k Dual port 10GbE KR Blade NDCs, Broadcom 57810 Dual Port 10Gb Network Adapters, and EqualLogic PS6110 iSCSI SAN arrays support DCB and DCBX.

Within the Active System 800v environment, DCB settings are configured within the Force10 S4810 switches.  Utilizing the DCBX protocol, these settings are then automatically propagated to the PowerEdge M I/O Aggregator modules. Additionally, the DCB settings are also propagated to the network end nodes, including the Broadcom Network Adapters in PowerEdge R620 rack servers, the Broadcom NDCs in the PowerEdge M620 blade servers, and the EqualLogic PS6110 storage controllers. The DCB settings are not propagated to the Force10 S55 out-of-band management switch and the associated out-of-band management ports but the out-of-band management traffic going to the core from Force10 S55 switch traverses through the Force10 S4810 switches. When the out-of-band management traffic traverses through the Force10 S4810 switches, it obeys the DCB settings.

DCB technologies enable each switch-port and each network device-port in the converged network to simultaneously carry multiple traffic classes, while guaranteeing performance and QoS. In case of Active System 800v, DCB settings are used for the two traffic classes: (i) Traffic class for iSCSI traffic, and (ii) Traffic class for all non-iSCSI traffic, which, in the case of Active System 800v, are different LAN traffic types. DCB ETS settings are configured to assign bandwidth limits to the two traffic classes.

These bandwidth limitations are effective during periods of contention between the two traffic classes. The iSCSI traffic class is also configured with Priority Flow Control (PFC), which guarantees lossless iSCSI traffic.

The Broadcom Network Adapters and the Broadcom NDCs support DCB and DCBX. This capability, along with iSCSI hardware offload, allows Active System 800v solution to include an end-to-end converged network design, without requiring support from the VMware vSphere hypervisor for DCB.

Figure 5 below provides a conceptual view of converged traffic with DCB in Active System 800v.

Figure 5: Conceptual View of Converged Traffic Using DCB
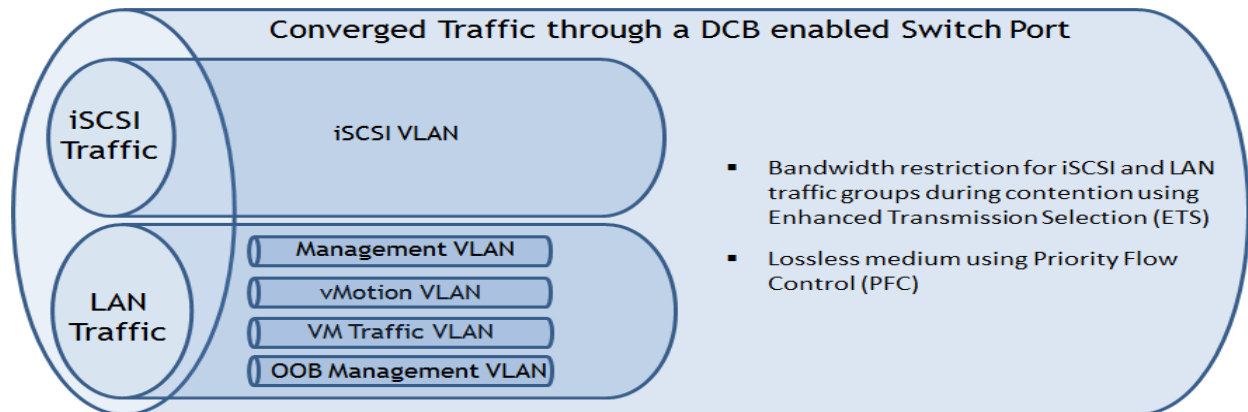
**Virtual Link Trunking (VLT) for S4810s:** Inside each Active System 800v, a Virtual Link Trunking interconnect (VLTi) is configured between the two Force10 S4810 switches using the Virtual Link Trunking (VLT) technology. VLT peer LAGs are configured between the PowerEdge M I/O Aggregator modules and Force10 S4810 switches, and also between the Force10 S4810 switch and the Force10 S4810 switches.

Virtual Link Trunking technology allows a server or bridge to uplink a single trunk into more than one Force10 S4810 switch, and to remain unaware of the fact that the single trunk is connected to two different switches. The switches, a VLT-pair, make themselves appear as a single switch for a connecting bridge or server. Both links from the bridge network can actively forward and receive traffic. VLT provides a replacement for Spanning Tree-based networks by providing both redundancy and active-active full bandwidth utilization.

Major benefits of VLT technology are:

1. Dual control plane on the access side that lends resiliency.

2. Full utilization of the active LAG interfaces.

3. Rack-level maintenance is hitless and one switch can be kept active at all times.

Note that the two switches can also be stacked together. However, this is not recommended, as this configuration will incur downtime during firmware updates of the switch or failure of stack links.

**NPAR configuration:**

In Active System 800v, each port of the Broadcom 57810-k Dual port 10GbE KR Blade NDCs in the PowerEdge M620 blade servers, and the Broadcom 57810 Dual Port 10Gb Network Adapters in

PowerEdge R620 rack servers is partitioned into four ports using NPAR to obtain a total of eight I/O ports on each server. As detailed in the subsequent sections, one partition each on every physical I/O port is assigned to management traffic, vMotion traffic, VM traffic and iSCSI traffic.

The Broadcom NDC and the Broadcom Network Adapter allow setting a maximum bandwidth limitation to each partition. Setting maximum bandwidth at 100 will prevent the artificial capping of any individual traffic type during periods of non-contention. For customers with specific requirements, NPAR maximum bandwidth settings may be modified to limit the maximum bandwidth available to a specific traffic type, regardless of contention.

The Broadcom NDC and the Broadcom Network Adapter also allow setting relative bandwidth assignments for each partition. While utilizing NPAR in conjunction with Data Center Bridging (DCB) and Data Center Bridging Exchange (DCBX), the relative bandwidth settings of the partitions are not enforced. Due this fact, the relative bandwidth capability of the Broadcom NDCs and the Broadcom Network Adapters are not utilized in Active System 800v.

**iSCSI hardware offload:** In Active System 800v, iSCSI hardware offload functionality is used in the Broadcom 57810-k Dual port 10GbE KR Blade NDCs in the PowerEdge M620 blade servers, and also in the Broadcom 57810 Dual Port 10Gb Network Adapters in the PowerEdge R620 rack servers. The iSCSI offload protocol is enabled on one of the partitions on each port of the NDC or the Network Adapter. With iSCSI hardware offload, all iSCSI sessions are terminated on the Broadcom NDC or on the Broadcom Network Adapter.

**Traffic isolation using VLANs:** Within the converged network, the LAN traffic is separated into four unique VLANs; one VLAN each for management, vMotion, VM traffic, and out-of-band management. The iSCSI traffic also uses a unique VM. Network traffic is tagged with the respective VLAN ID for each traffic type in the virtual switch. Routing between the management and out-of-band management VLANs is required to be configured in the core or the Force10 S4810 switches. Additionally, the Force10 S4810 switch ports that connect to the blade servers are configured in VLAN trunk mode to pass traffic with different VLANs on a given physical port. The table 2 below provides an overview of different traffic types segregated by VLANs in the Active System 800v, and which edge devices with which they are associated.

Table 2: VLAN Overview

| Traffic Type (VLAN segregation) | Description | Associated Network Device |
|---|---|---|
| Management | vSphere management traffic and Active System 800v management services | Broadcom NDC and Broadcom Network Adapter |
| vMotion | VMware vMotion traffic | Broadcom NDC and Broadcom Network Adapter |
| VM | LAN traffic generated by compute cluster VMs | Broadcom NDC and Broadcom Network Adapter |
| iSCSI | iSCSI SAN traffic | Broadcom NDC and Broadcom Network Adapter |
| Out-of-Band Management | Out-of-Band Management traffic | iDRAC, CMC, and EqualLogic Management Ports |

**Hypervisor network configuration for LAN and iSCSI SAN traffic:** VMware ESXi hypervisor is configured for the LAN and iSCSI SAN traffic that are associated with the blade servers. LAN traffic in Active System 800v solution is categorized into four traffic types: VM traffic, management traffic, vMotion traffic, and Out-of-Band (OOB) management traffic. OOB management traffic is associated with CMC, iDRAC, and EqualLogic SAN management traffic. VM traffic, management traffic, and vMotion traffic are associated with the blade servers in the compute cluster and the rack servers in the management servers. Similarly, iSCSI SAN traffic is also associated with the blade servers and the rack servers. On each hypervisor host within the compute cluster and the management cluster, a virtual switch is created for each of the three LAN traffic types associated with the blade and the rack servers, and also for the iSCSI traffic.

On the compute cluster hosts (the PowerEdge M620 blade servers), one vSwitch each is created for VM traffic, vSphere management traffic, vMotion traffic, and iSCSI traffic. Two partitions, one from each physical network port, are connected as uplinks to each of the virtual switches. This creates a team of two network ports, enabling NIC failover and load balancing for each vSwitch. On the management cluster hosts (the PowerEdge R620 rack servers), one vSwitch each is created for management traffic, vMotion traffic, and iSCSI traffic. In this case, all VMs are management VMs, so the VM traffic and the vSphere management traffic are on the same management VLAN. Due to this fact, the VM traffic port group and the vSphere management traffic port group are on the same vSwitch.

The resultant compute cluster and management cluster hypervisor host configuration is illustrated in Figure 6.

Figure 6: vSwitch and NPAR Configuration for the Hypervisor Hosts
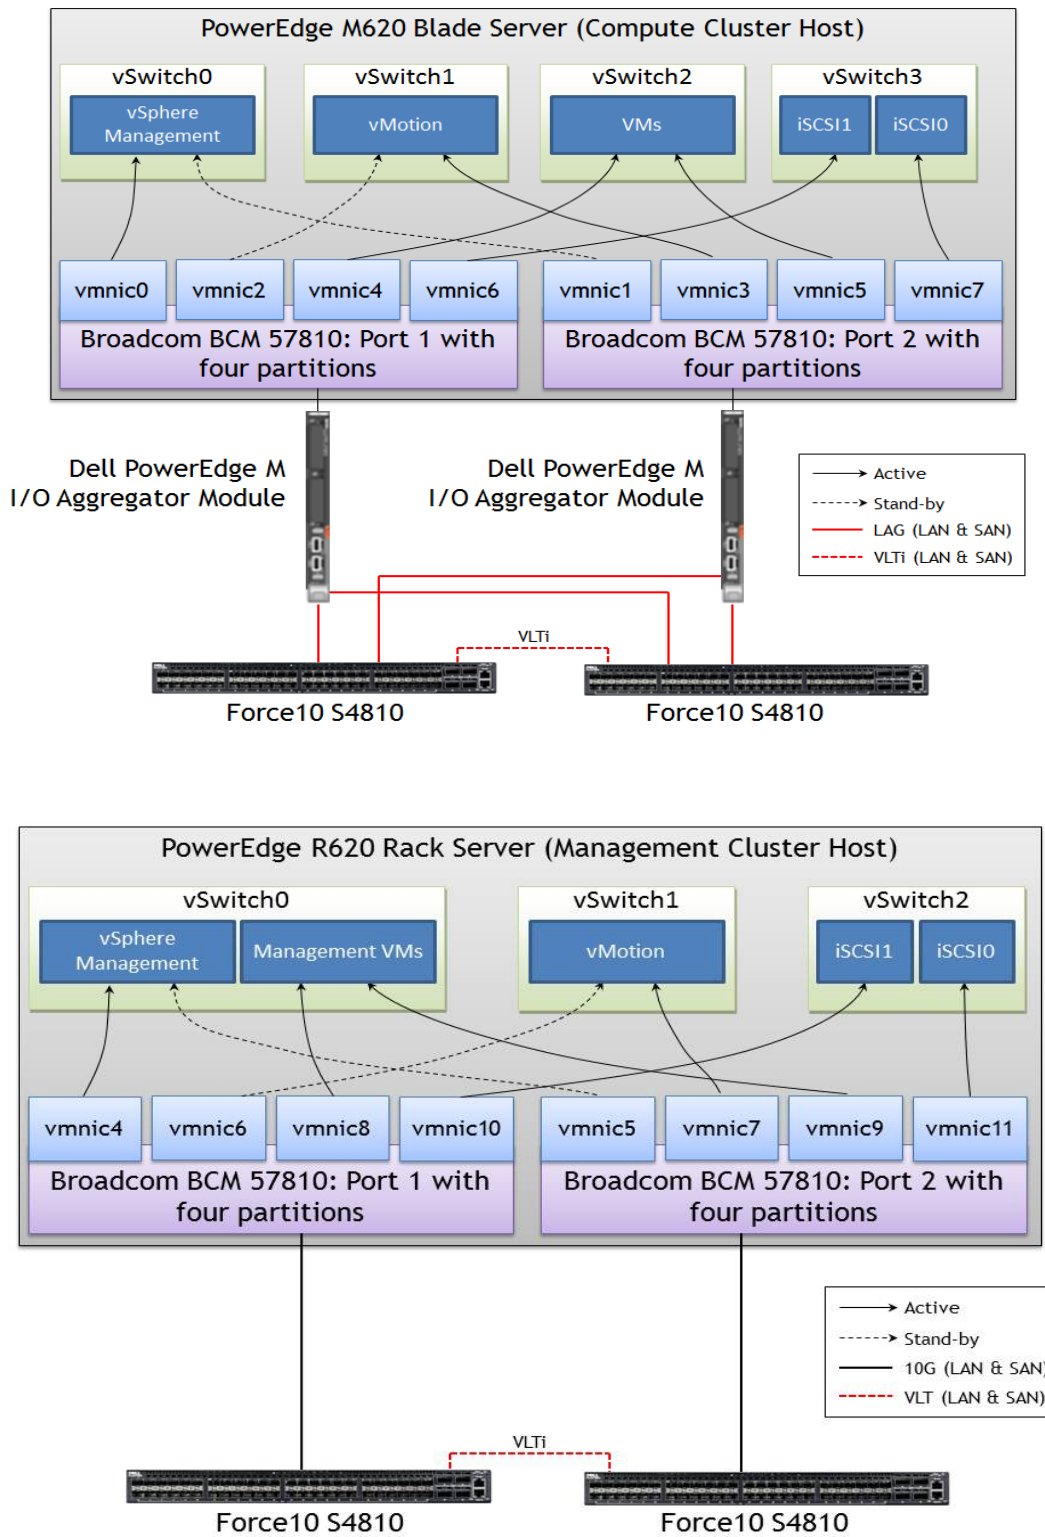
**Load Balancing and Failover:** This solution uses *Route based on the originating virtual switch port ID* configuration at the vSwitch for load balancing the LAN traffic. Any given virtual network adapter will use only one physical adapter port at any given time. In other words, if a VM has only one virtual NIC, it will use only one physical adapter port at any given time. The reason for choosing this option is that it is easy to configure and provides load balancing across VMs, especially in the case of a large number of VMs.

**Uplinks:** There are several options to uplink the Force10 switches to the core network. Selecting the uplink option depends on the customer core network and customer requirements. One simple option is to create multiple uplinks on each switch and connect them to the core network switches. Uplink LAGs can then be created from the Force10 S4810 switches to the core network.

## 7.2  Storage Architecture

EqualLogic PS6110 provides capabilities essential to the Active System 800v design, like 10Gb connectivity, 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 EqualLogic MEM and EqualLogic VSM for VMware.

### 7.2.1  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.

### 7.2.2  RAID Array Design

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, especially in random I/O situations. RAID 50 generally provides more usable storage, but has less performance than RAID 10. RAID 6 provides better data protection than RAID 50.

For more information on configuring RAID in EqualLogic, refer to the white paper, How to Select the Correct RAID for an EqualLogic SAN.

### 7.2.3  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 VMware configuration files when sizing these 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.

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 that can created.

For more information, refer to the current EqualLogic Firmware (FW) Release Notes available at the EqualLogic Support site.

## 7.2.4  Drive Types and Automated Tiered Storage

Dell EqualLogic PS6110 arrays, with the 10Gb dual-controller configuration, provide high bandwidth for data flows. This bandwidth is complemented with a large variety of drives in multiple speeds and sizes, including 10K RPM and 15K RPM SAS drives, 7.2K RPM NL-SAS drives and solid-state disks. The reference architecture presented in this document shows EqualLogic PS6110X arrays with 24 x 10K RPM SAS drives in each array. The disk and array type should be selected by carefully considering the workload requirements. Active System 800v supports a maximum of 8 x PS6110 arrays.

EqualLogic PS arrays provide IT organizations numerous techniques for storage tiering as a standard part of their all-inclusive feature set. These techniques extend the automation at the core of the PS Series design philosophy, while allowing broad customization of storage tiers to suit a wide range of business and organizational requirements.

## 7.2.5  Multipath Configuration

The Dell EqualLogic PS Series storage array supports multiple iSCSI SAN connections for performance and reliability. Multi-Path I/O (MPIO) provides multiple paths from servers to storage, delivering fault tolerance, high availability, and improved performance. Active System 800v uses EqualLogic Multipath Extension Module (MEM) for VMware vSphere to enable MPIO for the iSCSI storage.

EqualLogic MEM offers:

- Ease of installation and iSCSI configuration in ESXi servers

- Increased bandwidth

- Reduced network latency

- Automatic load balancing across multiple active paths

- Automatic connection management

- Automatic failure detection and failover

- Multiple connections to a single iSCSI target

Once installed, the EqualLogic MEM will automatically create iSCSI sessions to each member that a volume spans. As the storage environment changes, the MEM will respond by automatically adding or removing iSCSI sessions as needed.

Reference Architecture for Active System 800 with VMware vSphere

As storage I/O requests are generated on the ESXi hosts, the MEM plug-in will intelligently route these requests to the array member best suited to handle the request. This results in efficient load balancing of the iSCSI storage traffic, reduced network latency and increased bandwidth.

For more information on EqualLogic MEM, refer to white-paper, Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 5.1, 5.0 and 4.1 and PS Series SANs.

# 8 Management Infrastructure

Within the Active System 800v solution, two Dell PowerEdge R620 servers and one Dell Force10 S55 1/10GbE Ethernet switch are used for the management infrastructure. The Force10 S55 switch is used for out-of-band management connectivity for Dell CMC, Dell iDRAC, and the management ports on Dell EqualLogic arrays. The management cluster infrastructure imitates the compute cluster in using converged network infrastructure and configuration. The PowerEdge R620 servers are connected to the Force10 S4810 switches using Broadcom 57810 Dual Port 10Gb network adapters. The management servers are connected to the EqualLogic storage through the two Force10 S4810 switches.

Note that the EqualLogic storage is shared between the management cluster 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 R620 servers run VMware ESXi 5.1 hypervisor and are a part of the unique vSphere Cluster for management. VMware High Availability is enabled in that cluster to provide HA for virtual machines. Admission control is disabled in the VMware HA Cluster. If admission control is enabled, VMware HA would prevent putting one of the management servers in maintenance mode, since this would violate HA policy of having more than one active server in the cluster.

The Active System 800v solution includes the necessary management components required to manage the Active System 800v infrastructure, including the Converged Infrastructure management software, Dell Active System Manager. The following management components are included in the Active System 800v solution.

- Dell Active System Manager
- VMware vCenter Server
- Dell Management Plug-in for VMware vCenter
- Dell OpenManage Essentials (OME)
- Dell EqualLogic Virtual Storage Manager (VSM) for VMware
- Dell EqualLogic SAN HeadQuarters (HQ)
- Dell Repository Manager
- VMware vCloud Connector

These components are installed in virtual machines installed as virtual machines in the management infrastructure as illustrated in Figure 7.

Figure 7: Management Components
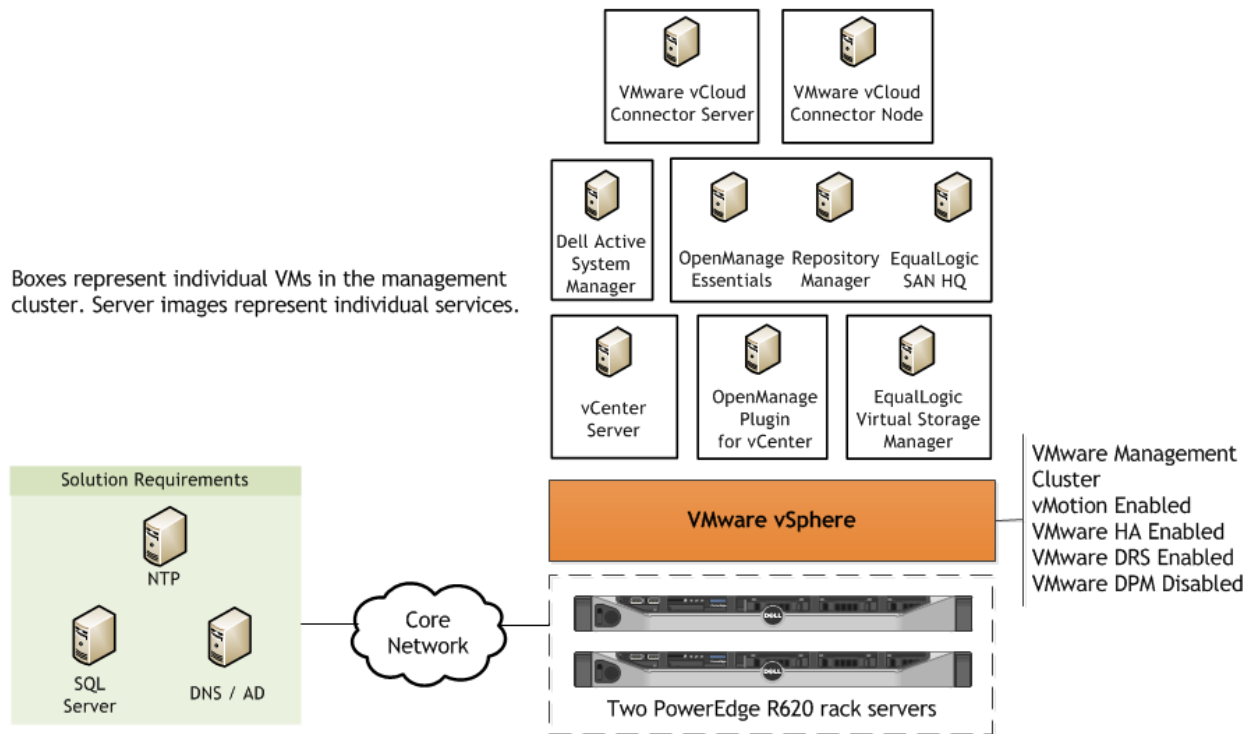
The remainder of this section will provide an introduction to each component and how they are integrated into the Active System 800v solution.

## 8.1 Dell Active System Manager

As described in section 3, the Dell Active System Manager is the Active Infrastructure management software that is part of the Active System 800v solution. 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.

For more information on Dell Active System Manager, see Dell Active System Manager.

## 8.2 Dell OpenManage Essentials (OME)

In the Active System 800v, Dell OpenManage Essentials (OME) is sized and configured to monitor the Active System 800v 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 utilized a local SQL Express database. For fullest functionality, direct internet access, or through a proxy, is recommended.

Within the Active System 800v, OME is utilized for discovery, inventory, and hardware level monitoring of blade & Rack servers, blade chassis, PowerEdge M I/O Aggregator modules, EqualLogic storage, and

Force10 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 component thought 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.

## 8.3   Dell Repository Manager (DRM)

Within the Active System 800v 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 Admins 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 M620. DRM can be used in conjunction with other OpenManage tools helps to 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](http://content.dell.com/us/en/enterprise/d/solutions/repository-manager).

## 8.4   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 800v 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 recommend that the DMPVV appliance has access to an internet connect either directly, or thought 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](#).

## 8.5   Dell EqualLogic Virtual Storage Manager (VSM) for VMware

Within Active System 800v, the Dell EqualLogic Virtual Storage Manager (VSM) for VMware is deployed as a virtual appliance within the management cluster and is attached to the VMware vCenter Server within the Active System 800v stack.  VSM communicates with the dedicated management interfaces of the EqualLogic storage enclosures over the out-of-band network. VSM enables customers to preform 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 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.

## 8.6   Dell EqualLogic SAN HQ

Within the Active System 800, 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 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.

## 8.7  VMware vCloud Connector

VMware vCloud Connector is an optional component of the Active System 800v 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 responsibility for the physical transfer of VM workloads.  Within the Active System 800v, 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 800v 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.

As described in the section 3.11 of this document, "Dell Cloud Connectivity using 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 Dell vCloud public cloud. The key capabilities provided by VMware vCloud Connector are:

- Expand your view across hybrid clouds. Use a "single pane of glass" management interface that seamlessly spans your private vSphere and public Dell vCloud environment.

- Extend your datacenter.  Move VMs, vApps, and templates from private vSphere to a Dell vCloud to free up your on-premise datacenter resources as needed.

- Consume cloud resources with confidence.  Run Development, QA, and production workloads using Dell vCloud, a VMware technology-based public cloud.

The Dell Cloud with VMware vCloud™ Datacenter is an enterprise-class, multi-tenant infrastructure-as-a-service (IaaS) public cloud solution that is hosted in secured Dell data centers. Utilizing VMware vCloud Connector, Dell Cloud provides you with unique hybrid cloud capabilities to extend your internal data center with Dell and VMware by transitioning your VMware virtualized workloads into our vCloud data center. vCloud hosting provides you with a secure, manageable, and flexible public cloud application.

# 9  Scalability

As workloads increase, the solution can be scaled to provide additional compute and storage resources independently.

**Scaling Compute and Network Resources:** This solution is configured with two Force10 S4810 network switches. Up to two PowerEdge M1000e chassis can be added to the two Force10 S4810 switches. In order to scale the compute nodes beyond two chassis, new Force10 S4810 switches need to be added. Additional switches can either be stacked together and/or connected to this distribution switch based on customer needs.

**Scaling Storage Resources:** EqualLogic storage can be scaled seamlessly and independent of the compute and network architectures. Additional EqualLogic PS6110 arrays of same or different

configuration can be added to the existing PS 6110 arrays. New volumes can be created or existing volumes can be expanded to utilize the capacity in the added enclosures. Active System 800v solution can scale up to maximum of 8 arrays. To scale beyond this, additional racks can be added which may require additional switches and networking.

# 10 Delivery Model

This Reference Architecture can be purchased as a complete solution, the Dell Active System 800v. This solution is available to be racked, cabled, and delivered to the customer site, to speed deployment. Dell Services will deploy and configure the solution tailored to the business needs of the customer and based on the architecture developed and validated by Dell Engineering. For more details or questions about the delivery model, please consult with your Dell Sales representative.

Figure 8 below shows the Active System 800v solution with a single chassis. Figure 9 shows Active System 800v with two chassis and maximum storage enclosures. Note that all EqualLogic arrays shown in the figures are PS6110X. If a different PS6110 array type is ordered, the actual rack configuration may be different from the one shown below. Also note that the switches shown in figures are shown mounted forward for representation. In actual use, ports face the back of the rack. The PDUs are not shown in the illustration because they will vary by region or customer power requirements.

**Figure 8: Active System 800v Single Chassis: Rack Overview**

Figure 9: Active System 800v Two Chassis and Maximum Storage: Rack Overview

# 11 References

## 11.1 Dell Active Infrastructure reference:

- [Dell Active System Manager](#)

- [Dell Active Infrastructure Wiki](#)

## 11.2 VMware references:

- [VMware vSphere Edition Comparisons](#)

- [VMware vSphere Compatibility Matrixes](#)

- [VMware High Availability (HA): Deployment Best Practices](#)

- [VMware Virtual Networking Concepts](#)

## 11.3 Dell PowerEdge References:

- [Dell PowerEdge M1000e Technical Guide](#)

- [Dell PowerEdge M I/O Aggregator Configuration Quick Reference](#)

- [NIC Partitioning (NPAR)](#)

## 11.4 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](#)

## 11.5 Dell Management reference:

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