# DVS Enterprise - Active System 800v with VMware Horizon View 5.2: Reference Architecture

*Release 1.0 for Dell PowerEdge 12th Generation Blade Servers, Dell Force10 Switches, and Dell EqualLogic iSCSI SAN*

**Dell Desktop Virtualization Solutions (DVS) Engineering**

**Revision: v1.0**

## Revision History

| Revision | Description | Date |
|----------|-------------|------|
| V 1.0 | Andrew Mc Daniel – Solution Development Manager DVS - Author | Feb. 2013 |

## Contents

## 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 DVS Enterprise - Active System 800v solution, a member of Dell Desktop Virtualization Solutions family, is a converged infrastructure solution that has been designed and validated by Dell™ Engineering for desktop virtualization. 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. DVS Enterprise - Active System 800v is offered with VMware® vSphere® (Active System 800v) hypervisor. This paper defines the Reference Architecture for DVS Enterprise - Active System 800v with VMware Horizon View 5.2.

DVS Enterprise - Active System 800v offers converged LAN & SAN fabric design to enable a converged infrastructure solution. The end-to-end converged network architecture in DVS Enterprise - 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 DVS Enterprise - Active System 800v reduces complexity and cost while bringing greater flexibility to the infrastructure solution.

DVS Enterprise - 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. VMware vCenter Server, EqualLogic Virtual Storage Manager for VMware, Dell OpenManage™ Essentials, and Active System Manager 7.0 are included with the solution.

# 2 Audience

IT administrators and IT managers - who have purchased, or are planning to purchase a DVS Enterprise - Active System 800v configuration – can use this document to understand the design elements, hardware and software components, and the overall architecture of the solution.

# 3   Overview

This section provides a high-level product overview of VMware vSphere, VMware Horizon View 5.2, 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.

**Figure 1.**   DVS Enterprise - Active System 800v Overview



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

**VMware Horizon View**

**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
- 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**
- 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
- Dell Active System Manager 7.0

**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 | 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><br>• 2x Dell PowerEdge M I/O Aggregator in each Dell PowerEdge M1000e chassis |
| **Storage** | • Up to 8 x Dell EqualLogic PS6110 series arrays |
| **Management Infrastructure** | • 2 x Dell PowerEdge R620 servers with embedded VMware vSphere 5.1 hosting management VMs.<br><br>• 1 x Dell Force10 S55 used as a 1Gb out-of-band management switch |
| **Management components hosted in the management infrastructure** | • VMware vCenter Server<br><br>• VMware Horizon View 5.2 Composer<br><br>• VMware Horizon View 5.2 Connection server<br><br>• Dell Management Plug-in for VMware vCenter<br><br>• Dell OpenManage Essentials<br><br>• Dell EqualLogic Virtual Storage Manager (VSM) for VMware<br><br>• Dell EqualLogic SAN HeadQuarters (HQ)<br><br>• VMware vCloud Connector<br><br>• Dell Repository Manager<br><br>• Dell Active System Manager 7.0 |

## 3.1   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.

- **vSphere Storage DRS™ and Profile-Driven Storage** – New integration with VMware vCloud® Director™ enables further storage efficiencies and automation in a private cloud environment.

- **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.

- **Space Efficient Sparse Virtual Disks -** SE Sparse Disks introduces an automated mechanism for reclaiming stranded space. SE Sparse disks also have a new configurable block allocation size which can be tuned to the recommendations of the storage arrays vendor, or indeed the applications running inside of the Guest OS. VMware Horizon View 5.2 is the only product that will use the new SE Sparse Disk in vSphere 5.1**.**

- **VMware vCenter Update Manager -** VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.

- **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.

- **vSphere Web Client** –The vSphere Web Client is now the core administrative interface for vSphere. This new flexible, robust interface simplifies vSphere control through shortcut navigation, custom tagging, enhanced scalability, and the ability to manage from anywhere with Internet Explorer or Firefox-enabled devices.

- **vCenter Single Sign-On** – Dramatically simplify vSphere administration by allowing users to log in once to access all instances or layers of vCenter without the need for further authentication.

For more information on VMware vSphere, see [www.vmware.com/products/vsphere](www.vmware.com/products/vsphere).

### 3.1.1    vSphere HA (Shared Tier 1)

Both Compute and Mgmt hosts are identically configured, within their respective tiers, and leverage shared storage so can make full use of vSphere HA. The Compute hosts should be configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware and View Composer (2000 VMs per vCenter). This will result in multiple HA clusters managed by multiple vCenter servers. With vSphere 5.1 a HA cluster can now accommodate 32 compute nodes.

Figure 2.    Compute Host Cluster
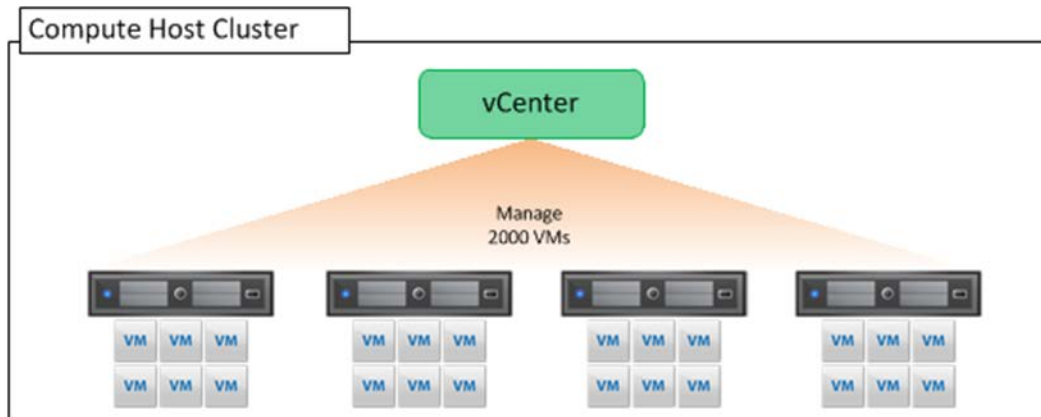


A single HA cluster will be sufficient to support the Management tier up to 10K users. An additional host can be used as a hot standby or to thin the load across all hosts in the cluster.

Figure 3.    Management Host Cluster
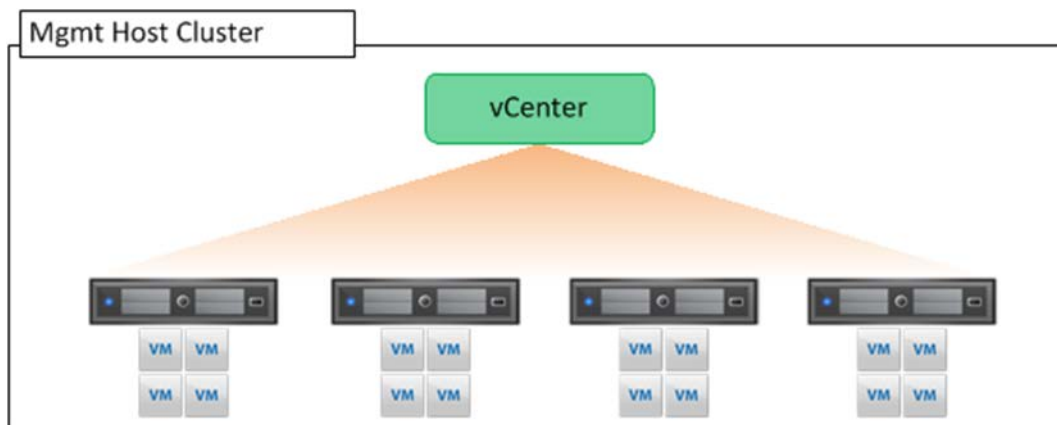


vSphere HA Admission control can be configured one of three ways to protect the cluster. This will vary largely by customer preference but the most manageable and predictable options are percentage reservations or a specified hot standby. Reserving by percentage will reduce the overall per host density capabilities but will make use of all hardware in the cluster. Specifying a failover host, on the other hand, will ensure maximum per host density numbers, is easier to maintain, but will result in hardware sitting idle.

**Figure 4.** HA Configuration



### 3.1.2 Distributed Resource Scheduler (DRS)

DRS provides the ability for automated VM placement within the cluster as they power on. DRS can be optionally used on the Compute tier set to Fully Automated with a Conservative migration threshold set.

**Figure 5.** DRS Configuration



Affinity rules should be used on the mgmt servers to ensure that View broker roles are properly distributed.

## 3.2   VMware Horizon View 5.2

**VMware Horizon View 5.2** is a desktop virtualization solution that delivers virtual desktops as an on-demand service to any user, anywhere. With VMware's desktop delivery technology, Horizon View 5.2 can quickly and securely deliver individual applications or complete desktops to the entire enterprise, whether they are task workers, knowledge workers or mobile workers. Users now have the flexibility to access their desktop on any device, anytime, with a high-definition user experience. With VMware Horizon View 5.2, IT can manage single instances of each OS, application and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management.

›   **VMware Horizon View 5.2 PCoIP protocol:** Was designed to deliver an uncompromised desktop experience. To deliver on this vision, PCoIP was architected to recognize different types of content and then use different compression algorithms based on the content type.

›   **VMware Horizon View 5.2 Administrator:** Provides a single management tool to provision new desktops or groups of desktops, and an easy interface for setting desktop policies. Using a template, you can customize virtual pools of desktops and easily set policies, such as how many virtual machines can be in a pool, or logoff parameters.

›   **VMware Horizon View 5.2 Composer**: Based on the mature Linked Clone technology, VMware Horizon View 5.2 Composer enables the rapid creation of desktop images from a golden image. Updates implemented on the parent image can be easily pushed out to any number of virtual desktops in minutes, greatly simplifying deployment, upgrades and patches while reducing desktop operational costs. With the core components of the desktop being managed separately the process does not affect user settings, data or applications, so the end-user remains productive on a working desktop, even while changes are being applied to the master image.

›   **VMware Horizon View 5.2 connection Server**: Acts as a broker between end users and virtual desktops they are allowed to access.

Features included in Horizon View support **usability, security, centralized control, and scalability**. This release of VMware Horizon View 5.2 adds the following new features and support.

•   On Microsoft Windows client devices, print from a virtual desktop to any local or networked printer that is defined on the Windows client device. This virtual printer feature solves compatibility issues and does not require you to install additional print drivers in a virtual machine.

•   On any client device, use the location-based printing feature to map to printers that are physically near the client system. Location-based printing does require that you install print drivers in the virtual machine.

•   Use multiple monitors. With PCoIP multiple-monitor support, you can adjust the display resolution and rotation separately for each monitor.

- Access USB devices and other peripherals that are connected to the local device that displays your virtual desktop. You can specify which types of USB devices end users are allowed to connect to. For composite devices that contain multiple types of devices, such as a video input device and a storage device, you can split the device so that one device (for example, the video input device) is allowed but the other device (for example, the storage device) is not.

- Use View Persona Management to retain user settings and data between sessions even after the desktop has been refreshed or recomposed. View Persona Management has the ability to replicate user profiles to a remote profile store (CIFS share) at configurable intervals. You can also use a standalone version of View Persona Management on physical computers and virtual machines that are not managed by View.

Horizon View offers the following **security features**, among others:

- Use two-factor authentication, such as RSA SecurID or RADIUS (Remote Authentication Dial-In User Service), or smart cards to log in.

- Use pre-created Active Directory accounts when provisioning View desktops in environments that have read-only access policies for Active Directory.

- Use SSL tunneling to ensure that all connections are completely encrypted.

- Use VMware High Availability to host desktops and to ensure automatic failover.

**Scalability features** depend on the VMware virtualization platform to manage both desktops and servers:

- Integrate with VMware vSphere to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for your virtual desktops.

- Use the Horizon View storage accelerator feature to support end-user logins at larger scales with the same storage resources. This storage accelerator uses features in the vSphere 5 platform to create a host memory cache of common block reads.

- Configure View Connection Server to broker connections between end users and the virtual desktops that they are authorized to access.

- Use View Composer to quickly create desktop images that share virtual disks with a master image. Using linked clones in this way conserves disk space and simplifies the management of patches and updates to the operating system.

The following features provide **centralized administration and management**:

- Use Microsoft Active Directory to manage access to virtual desktops and to manage policies.

- Use View Persona Management to simplify and streamline migration from physical to virtual desktops.

- Use the Web-based administrative console to manage virtual desktops from any location.

- Use a template, or master image, to quickly create and provision pools of desktops.

- Send updates and patches to virtual desktops without affecting user settings, data, or preferences.

- Integrate with Horizon Workspace so that end users can access View desktops through the Horizon Workspace Web Client, as well as use the Horizon Workspace Web Client from inside a View desktop.

When implementing a VMware Horizon View 5.2 solution it's critical to design the vSphere cluster and View components with high availability and resource management according to best practices. The Dell DVS Engineering team leverages the strategic relationship with VMware to evaluate and implement the most current best practices for a given configuration.

## 3.3  Summary of Horizon View 5.2 Features

**End User Experience**

- Support for Windows 8 based desktops

- Horizon View 5.2 offers support for Windows 8 desktops.

- Facilitates a smooth transition in rolling out Windows 8.

- Enables IT to leverage the latest Windows capabilities in VDI

- Hardware Accelerated 3D Graphics

- Horizon View5.2 provides a rich workstation class user experience with high performance graphics

- Enables shared-access to physical GPU hardware for 3D and high performance graphical workloads.

- Very cost effective as multiple VMs share the same GPU resource.

- Offers full compatibility with hosts lacking physical GPUs


**Improved Video and VOIP communications with Microsoft Lync 2013 support**

- Horizon View 5.2 offers tighter integration with Microsoft Lync and Office applications

- Full collaboration capabilities with Microsoft Lync on View Desktops.

- Full support for Unified Communications VoIP and Video using Lync client on View desktops.

- Support for both RDP and PCoIP.

**Streamlined access to View Desktops from Horizon**

- View desktops can now be accessed via Horizon gateway.

- Horizon provides a single point of access for end users to desktops, data and apps.

- This provides a one-stop shop for all end-user access to their corporate workloads

**Easily connect to desktops from any device with HTML Access**

- View Desktops can now be accessed through a HTML5 capable web browser via Horizon.

- This provides install-free access to Desktops from ANY modern device.

**Ease of Management**

- Large Pool creation with elimination of 8 host limits and multiple vLAN support

- Horizon View 5.2 has support for large View pools with more than 8 hosts.

- This offers OPEX savings with less admin time spent on common operations.

- Even more OPEX savings with fewer pools to manage in multi-thousand user deployments.

- Tech Preview of a new Integrated Service Console in the VC Web Client

- Horizon View 5.2 offers a View plugin into vSphere Web client.

- Allows for easier desktop support and troubleshooting.

- Offers a simple interface for novice administrative users which allows for increased efficiency

- Support for VC Virtual Appliance based deployments

- Horizon View 5.2 is fully compatible with Virtual Appliance-based VirtualCenter deployments.

- This eliminates VirtualCenter dependencies on Windows

- Easier installation and upgrades.

**Total Cost of Ownership**

- Substantial storage capacity savings for persistent desktops with Space Efficient Disks

- View 5.2 leverages a vSphere capability to offer a new disk format for VMs on VMFS.

- Space Efficient disks provide reduced storage capacity requirements (lower CAPEX) for persistent desktops, even on lower-tier hardware

- Unused space is reclaimed and View composer desktops stay small.

## 3.4  Dell PowerEdge Blade Servers

**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.

**Blade Servers:** The PowerEdge M620 blade server is the Dell 12[th] 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 24 x 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 enclosure provides three redundant fabrics using six I/O modules. The modules can be populated with Ethernet switches, Fibre Channel (FC), and pass-through modules. InfiniBand™ switch modules are also supported.

**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.

The table below shows the configuration of the compute nodes used for the Desktop sessions.

Table 2. M620 Compute Host Configuration

| Shared Tier 1 Compute Host – PowerEdge M620 |
| --- |
| 2 x Intel Xeon E5-2680 Processor (2.7Ghz) |
| 192GB Memory (12 x 16GB DIMMs @ 1600Mhz) |
| VMware vSphere 5.1 on internal SD |
| Diskless |
| Broadcom 57810-k 10Gb DP KR NDC |
| iDRAC7 Enterprise w/ vFlash, 8GB SD |

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.

## 3.5   Dell PowerEdge M I/O Aggregator

The Dell PowerEdge M I/O Aggregator (IOA) is a flexible 1/10GbE aggregation device that is automated and pre-configured for easy "plug and play" deployment into converged iSCSI and FCoE networks. The key feature of the PowerEdge M I/O Aggregator 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/O Aggregator are automatically all part of a single link aggregation group (LAG), and thus there is no need for Spanning-tree.

The PowerEdge M I/O Aggregator provides connectivity to the CNA/Network adapters internally and externally to upstream network devices. Internally the PowerEdge M I/O Aggregator 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/O Aggregator 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/O Aggregator 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/O Aggregator 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/O Aggregator is reached by connection through the CMC's management port. This one management port on the CMC allows for management connections to all IO 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

## 3.6  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 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

OpenManage Server Administrator (OMSA) is installed on each ESXi server to provide comprehensive one-to-one management solution.

For more information on OpenManage Essentials, see http://content.dell.com/us/en/enterprise/data-center-system-management

## 3.7  Dell Force10 S4810 Switches

The Force10 S-Series S4810 is an ultra-low-latency 10/40 GbE Top-of-Rack (ToR) switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 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 via a future software enhancement, 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.

## 3.8   Dell Force10 S55

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 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.

## 3.9   Dell EqualLogic PS6110 Series iSCSI SAN Arrays

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 Data Center Bridging (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

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

- 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.

## 3.10 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.

Table 3.    R620 Management Host Configuration

| Shared Tier 1 Mgmt Host — PowerEdge R620 |
| --- |
| 2 x Intel Xeon E5-2680 Processor (2.7Ghz) |
| 96GB Memory (6 x 16GB DIMMs @ 1600Mhz) |
| VMware vSphere 5.1 on internal SD (Dual 2GB) |
| Diskless |
| Broadcom 5720 QP 1Gb NDC |
| Broadcom 57810  DP 10Gb |
| iDRAC7 Enterprise w/ vFlash, 8GB SD |
| 2 x 750W PSUs |

For more information, see the PowerEdge R620 guides at [Dell.com/PowerEdge](Dell.com/PowerEdge).

## 3.11 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 Management Plug-in for VMware vCenter).

## 3.12 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 DVS Enterprise - Active System 800v Solution.

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

2. **Redundancy with no single point-of-failure:** Redundancy is incorporated in every critical aspect[1] of the solution, including server high availability features, networking, and storage.

3. **Management:** Provide integrated management using VMware vCenter, Dell Management Plug-in VMware vCenter, Dell OpenManage Essentials, and Compellent plug-in for VMware vCenter.

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

5. **Integration into an existing data center:** This architecture assumes that there is an existing 10 Gb Ethernet infrastructure with which to integrate.

6. **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.

7. **Racked, Cabled, and Ready to be deployed:** DVS Enterprise - Active System 800v is available partially 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.

8. **Power, Cooling, and Weight Considerations:** DVS Enterprise - 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 for the solution.

9. **Flexible configurations:** DVS Enterprise - 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.

---

[1] Out of band management is not considered critical to user workload and does not have redundancy.

# 5 Reference Architecture

This solution consists of a PowerEdge M1000e chassis populated with PowerEdge M620 blade servers running VMware ESXi. Figure 6 provides the high-level reference architecture for the solution.

Figure 6. DVS Enterprise - Active System 800v Network Topology (Logical View)
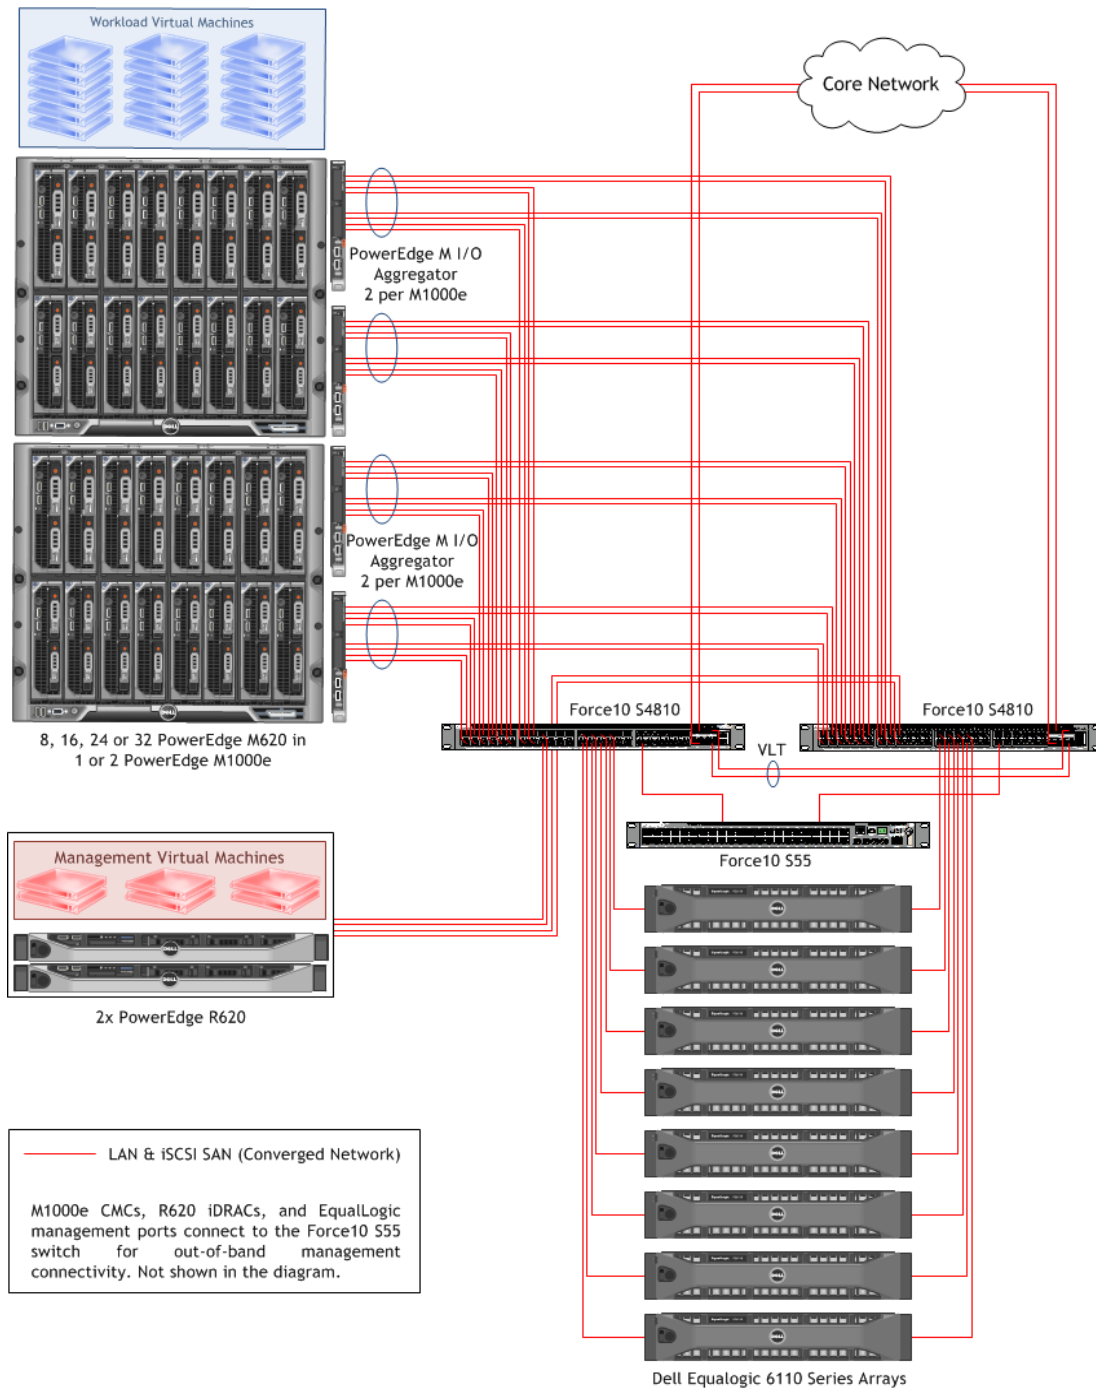


The figure above shows high-level logical connectivity between various components. Subsequent sections of this document provide more detailed connectivity information.

# 6   Dell Blade Network Architecture

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. 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.

In this solution, the Chassis Fabric A 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 7 below illustrates how the fabrics are populated in the PowerEdge M1000e blade server chassis and how the I/O modules are utilized.

**Figure 7.**   I/O Connectivity for PowerEdge M620 Blade Server
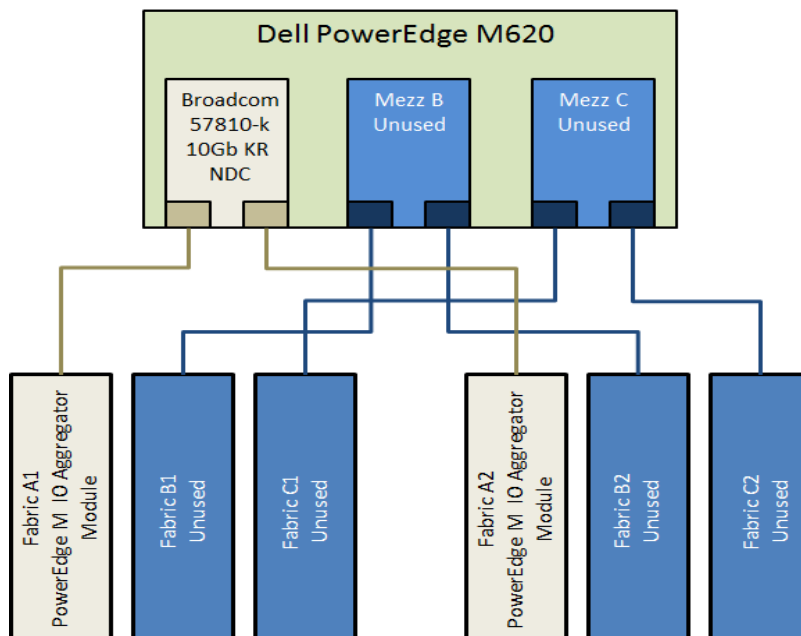


**Network Interface Card Partition (NPAR):** NPAR allows splitting the 10GbE pipe on the NDC with no specific configuration requirements in the switches. With NPAR, administrators can split each 10GbE port of an NDC into four separate partitions, or physical functions, and allocate the desired bandwidth 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, and hypervisor. DVS Enterprise - 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 Converged Network Architecture

One of the key attributes of the DVS Enterprise - 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 DVS Enterprise - Active System 800v to utilize these technologies, features, and capabilities to support converged network architecture.

## 7.1 Converged Network Connectivity

The DVS Enterprise - 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 DVS Enterprise - 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 DVS Enterprise - 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 8 below illustrates the resultant logical converged network connectivity within the DVS Enterprise - Active System 800v solution.

Figure 8.    Converged Network Logical Connectivity
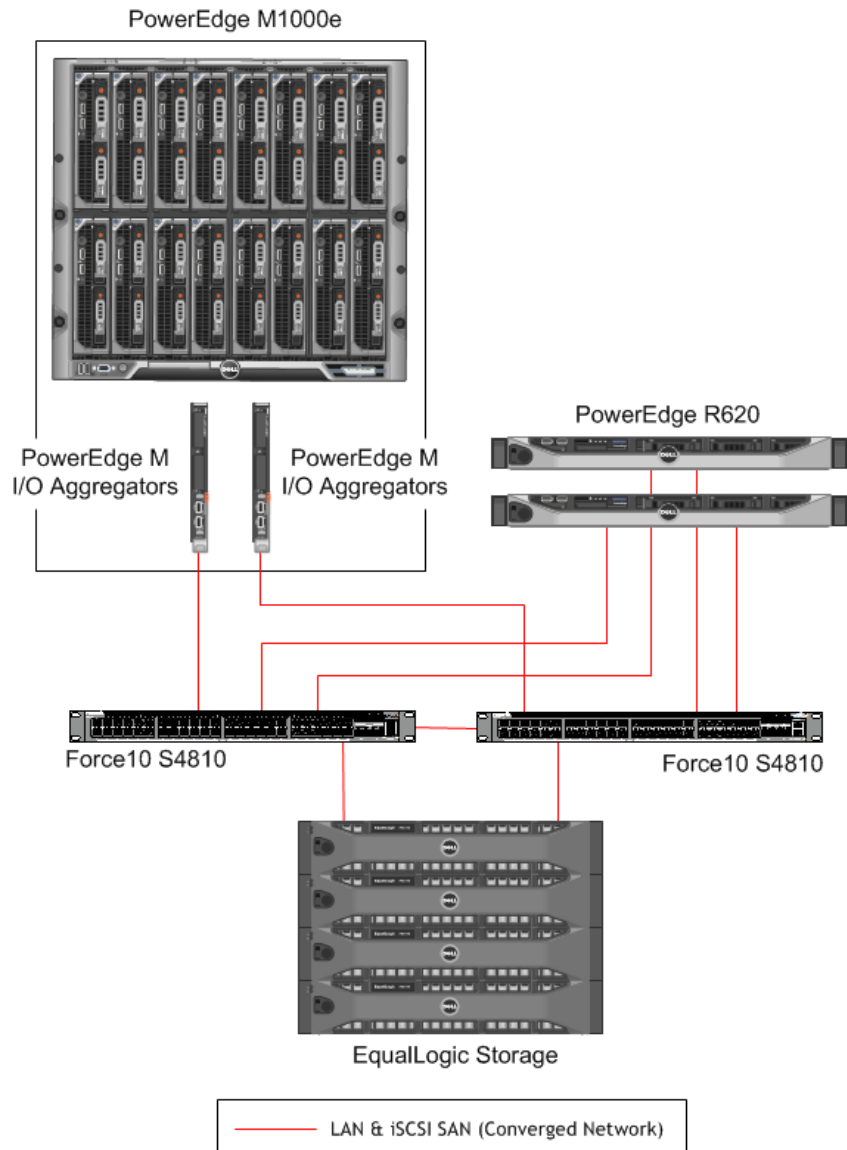
## 7.2   Converged Network Configuration

This section provides details of the different configurations in the DVS Enterprise - 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 DVS Enterprise - 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 DVS Enterprise - 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 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 DVS Enterprise - 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 DVS Enterprise - 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 DVS Enterprise - Active System 800v solution to include an end-to-end converged network design, without requiring support from the VMware vSphere hypervisor for DCB.

Figure 9 below provides a conceptual view of converged traffic with Data Center Bridging in DVS Enterprise - Active System 800v.

**Figure 9.** Conceptual View of Converged Traffic Using DCB
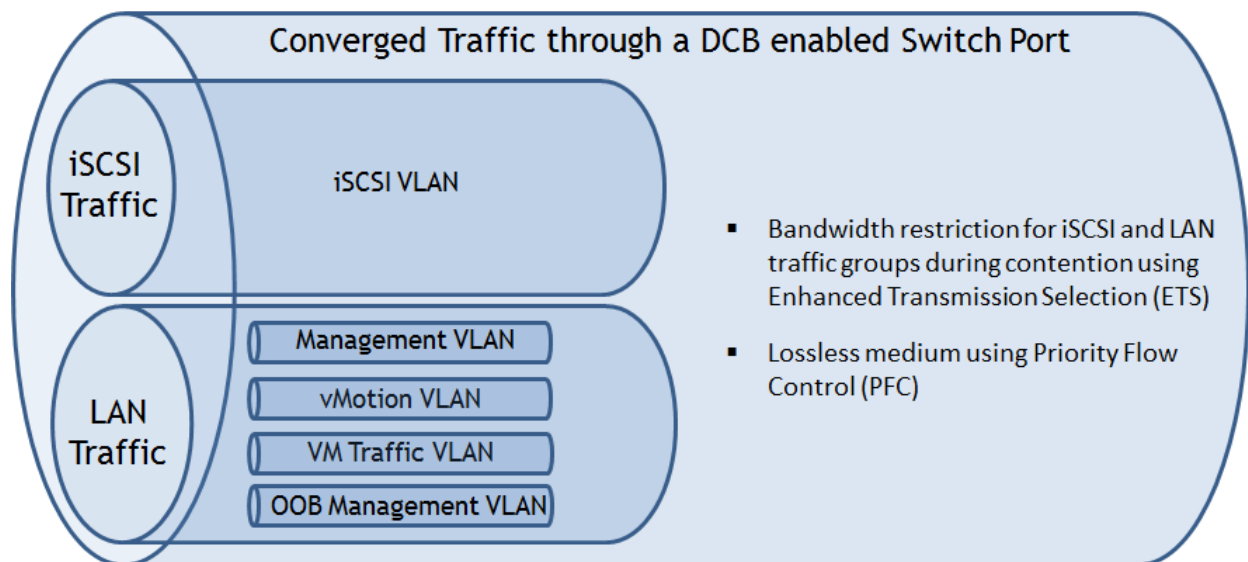


**Virtual Link Trunking (VLT) for S4810s:** Inside each DVS Enterprise - Active System 800v, a Virtual Link Trunking interconnect (VLTi) is configured between the two Force10 S4810 switches using the Virtual Link Trunking (VLT) technology. Aldo, VLT peer LAGs are configured between the PowerEdge M I/O 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 DVS Enterprise - 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 DVS Enterprise - Active System 800v.

**iSCSI hardware offload:** In DVS Enterprise - 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. Table 4, below, provides an overview of different traffic types segregated by VLANs in the DVS Enterprise - Active System 800v, and which edge devices they are associated with.

Table 4.    VLAN Overview

| Traffic Type (VLAN segregation) | Description | Associated Network Device |
|---|---|---|
| Management | vSphere management traffic and<br><br>DVS Enterprise - 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 DVS Enterprise - 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 10.

Figure 10.    vSwitch and NPAR Configuration for the Hypervisor Hosts
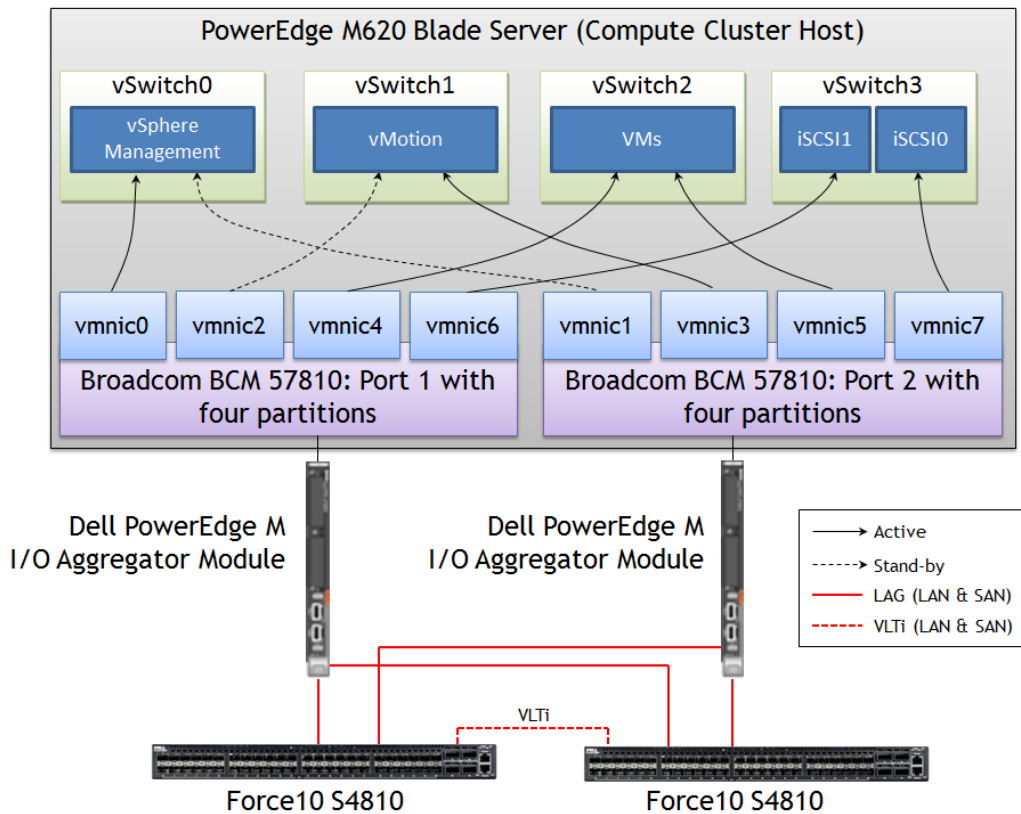
**Load Balancing and Failover:** This solution uses *Route based on the originating virtual switch port ID* configuration at the vSwitch for load balancing the LAN traffic. Any given virtual network adapter will use only one physical adapter port at any given time. In other words, if a VM has only one 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.

# 8  Storage Architecture

EqualLogic PS6110 provides capabilities essential to the DVS Enterprise - 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.

## 8.1  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.

Each tier of Equallogic storage should be managed as a separate pool or group to isolate specific workloads. Shared Tier 1 arrays used for hosting VDI sessions should be managed together, while shared Tier 2 arrays used for hosting Mgmt server role VMs and user data should be managed together.

Figure 11.    Tier 1 and Tier 2 Pool Configuration
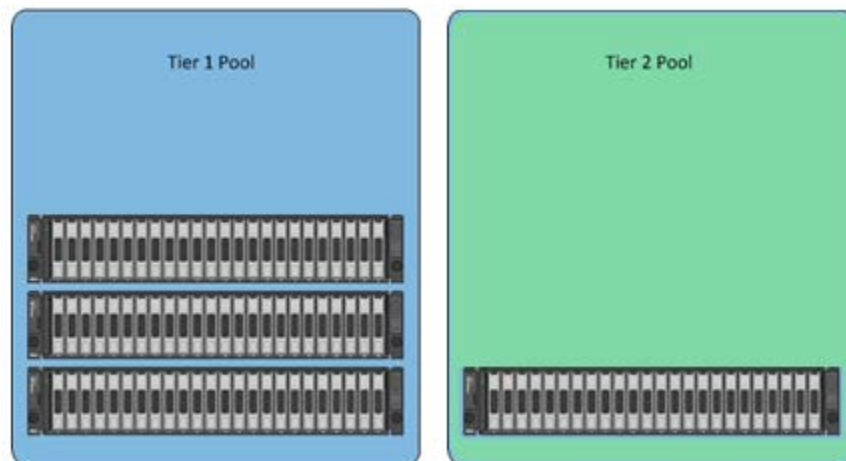
## 8.2   Shared Tier 1

Shared storage will be leveraged for Tier 1 used for VDI execution. Based on the heavy performance requirements of Tier 1 VDI execution, it is recommended to use separate arrays for Tier 1 and Tier 2. We recommend using 350GB LUNs for VDI and running 50 VMs per volume. Using the 6110XS as the shared Tier 1 array, sizing to 1000 basic users per array, we will require 20 x 350GB volumes per array.

Table 5.    Tier 1 Volume Sizing Recommendations

| Volumes | Size (GB) | Storage Array | Purpose | File System |
|---------|-----------|---------------|---------|-------------|
| VDI-1 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-2 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-3 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-4 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-5 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-6 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-7 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-8 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-9 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |
| VDI-10 | 350 | Tier 1 | 50 x desktop VMs + Replica | VMFS 5 |

## 8.3   PS6110XS

Table 6.    PS6110XS Configuration Details

| Model | Features | Options | Uses |
|-------|----------|---------|------|
| Equallogic PS6110XS | 24 drive hybrid array (SSD + 10K SAS), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 10Gb | 13TB – 7 x 400GB SSD + 17 x 600GB 10K SAS | Tier 1 array for shared Tier 1 solution model (10Gb) |

Figure 12.　PS6110XS
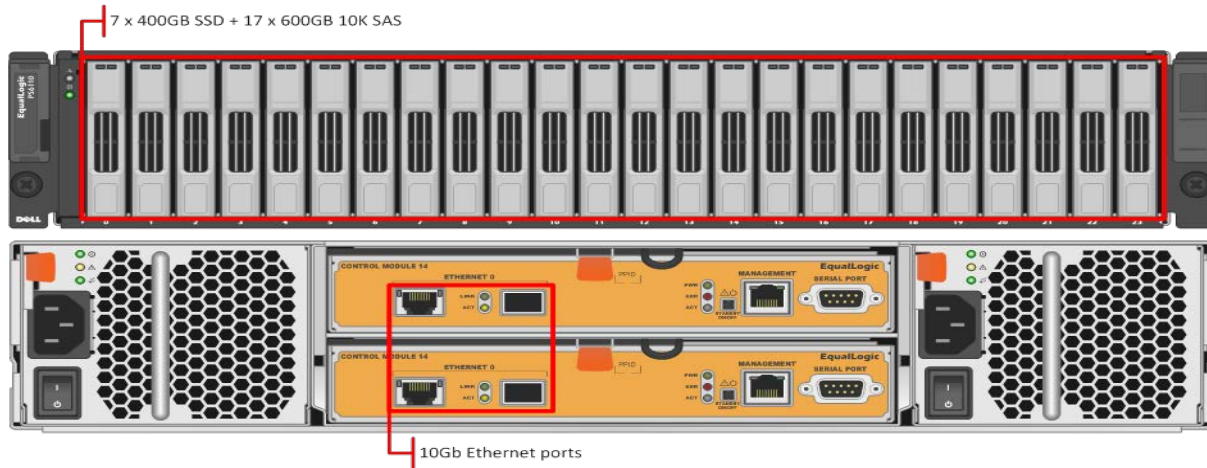


## 8.4  Shared Tier 2

Tier 2 is shared iSCSI storage used to host the mgmt server VMs and user data. The Equallogic 6110E series is utilized for this requirement. The table below outlines the volume requirements for Tier 2. Large disk sizes can be chosen to meet the capacity needs of the customer, up to 3TB disks. The user data can be presented either via a file server VM using RDM for small scale deployments. This is the only native NTFS volume on the array. All SQL disks should be presented as VMDKs. RAID 50 can be used in smaller deployments (>500 sessions). Larger scale deployments with higher performance requirements will need to use RAID 10 to maximize IOPS. The following depicts the component volumes required to support a 500 user environment. Additional Management volumes should be created as needed along with size adjustments as applicable for user data and profiles

Table 7.　Tier 2 Volume Sizing Recommendations

| Volumes | Size (GB) | Storage Array | Purpose | File System |
|---|---|---|---|---|
| Management | 300 | Tier 2 | vCenter, Connection Server, File Server and SQL | VMFS 5 |
| User Data | 2048 | Tier 2 | File server/ NAS | NTFS |
| User Profiles | 20 | Tier 2 | User profiles | VMFS 5 |
| SQL DATA | 100 | Tier 2 | SQL | VMFS 5 |
| SQL LOGS | 100 | Tier 2 | SQL | VMFS 5 |
| TempDB Data | 5 | Tier 2 | SQL | VMFS 5 |
| TempDB Logs | 5 | Tier 2 | SQL | VMFS 5 |
| SQL Witness | 1 | Tier 2 | SQL (optional) | VMFS 5 |
| Templates/ ISO | 200 | Tier 2 | ISO storage (optional) | VMFS 5 |

## 8.5   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. DVS Enterprise - 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.

As storage I/O requests are generated on the ESXi hosts, the MEM 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.

# 9  Management Infrastructure

Within the DVS Enterprise - 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 converge 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 both the management cluster and the compute cluster reside on shared storage. Compute cluster runs the VDI sessions and these reside on Tier 1 storage which utilizes Equallogic PS6110XS's. The Management cluster runs the Management VM's, User Profile and Home Directories, etc and utilizes a Equallogic PS6110E array. 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 DVS Enterprise - Active System 800v solution includes the necessary management components required to manage the infrastructure, including the Converged Infrastructure management software, Dell Active System Manager. The following management components are included in the DVS Enterprise - Active System 800v solution.

- VMware vCenter Server
- VMware Horizon View 5.2 Composer
- VMware Horizon View 5.2 Connection 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
- Dell Active System Manager 7.0
- VMware vCloud Connector

These components are installed as virtual machines in the management infrastructure as illustrated in Figure 13.

Figure 13.    Management Components
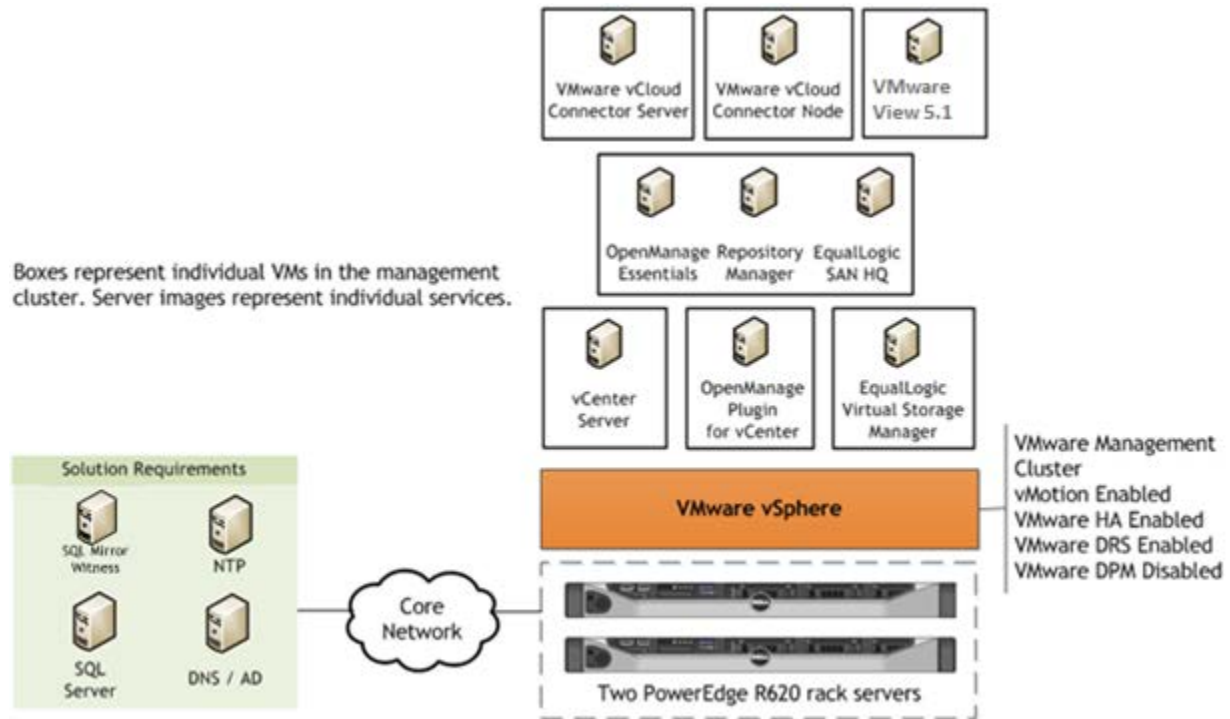


As discussed, all management servers will exist as VMs in a dedicated management server tier. Data disks should be used for role-specific application files/ data, logs, IIS web files, etc and should exist in the Management volume. Tier 2 volumes with a special purpose (called out above) should be presented in the format specified below:

Table 8.    Management VM Sizing Recommendations

| Role | vCPU | RAM (GB) | NIC | OS + Data vDisk (GB) | Tier 2 Volume (GB) |
|---|---|---|---|---|---|
| VMware vCenter | 2 | 8 | 1 | 40 + 5 | 100 (VMDK) |
| View Connection Server | 2 | 8 | 1 | 40 + 5 | - |
| SQL Server | 2 | 8 | 1 | 40 + 5 | 210 (VMDK ) |
| File Server | 1 | 4 | 1 | 40 + 5 | 2048 (RDM) |
| **TOTALS** | **7** | **28** | **4** | **180** | **2358** |

## 9.1   Dell OpenManage Essentials (OME)

In the DVS Enterprise - 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 DVS Enterprise - 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 Dell OpenManage Essentials, see Dell.com/openmanageessentials.

## 9.2   Dell Repository Manager (DRM)

Within the DVS Enterprise - 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.

## 9.3   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](#).

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

Within DVS Enterprise - 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 DVS Enterprise - 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.

## 9.5  Dell EqualLogic SAN HQ

Within the Active System 800v, 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.

## 9.6   Active System Manager 7.0

Active System Manager 7.0 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.

### 9.6.1   Features

- **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

## 9.8  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, '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.

# 10 VMware Horizon View 5.2 Infrastructure

Figure 14.    VMware Horizon View 5.2 Infrastructure
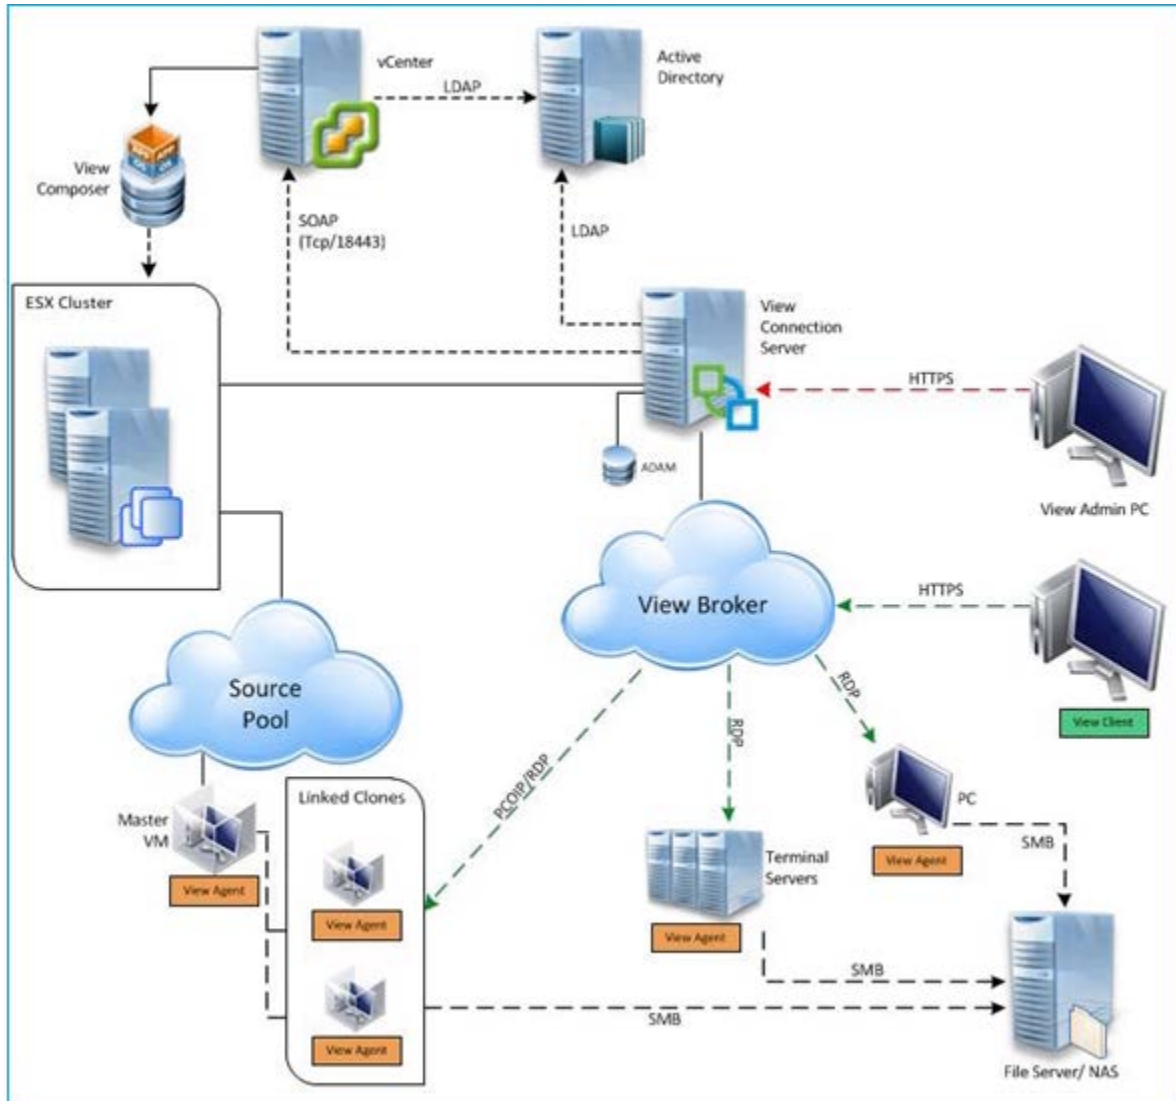


The DVS Enterprise - Active System 800v architecture with VMware Horizon View 5.2 is designed to provide maximum performance and scalability starting at very low user counts for SMBs and up to tens of thousands of users for large enterprise deployments. VMware Horizon View 5.2 brings with it many new features, as outlined in Section 3.2. This solution architecture follows a distributed model where solution components exist in tiers. The Compute tier is where VDI desktop VMs execute, the Mgmt tier being dedicated to the broker management server role VMs. Both tiers, while inextricably linked, scale independently.

Table 9. Solution Components

| Solution Architecture Components | |
|---|---|
| Hypervisor | VMware vSphere 5.1 Desktop Edition |
| VDI Broker | VMware Horizon View 5.2 Premier |
| Non-persistent provisioning | View Composer 3.0 |
| Database software | Microsoft SQL Server 2008 R2 (x64) |
| Server OS | Microsoft Windows Server 2008 R2 SP1 (x64) |
| Desktop OS | Microsoft Windows 7 Enterprise (x86) |

Table 10. Virtual Machine Configuration

| Virtual Machine Configuration | |
|---|---|
| VM Version | 8 |
| Virtual CPU | vCPU x 1 |
| SCSI Controller | LSI Logic SAS |
| Network Adapter | VMXNET 3 |
| Disk Provisioning | Thin |
| Swap File | Store with VM |
| vRAM Reservations (Mgmt) | 50% of configured vRAM |
| vRAM Reservations (VDI) | 500MB – Basic, 1GB – Std, 1.5GB – Prem |

## 10.1 Load Balancing

The applicable core VDI roles will be load balanced via DNS in stacks <1000 users. Above that number, Citrix NetScaler VPX could be introduced to manage load-balancing efforts. This is called out in the scaling scenarios. All roles whose configurations are stored in SQL can be protected via an optional SQL mirror. With Storage configurations >1000 users, the file server VM could be replaced by a NAS such as the FS7600 for Equallogic.

# 11 Dell Wyse Cloud Clients

Dell Wyse Cloud Client devices and software provide superior security, reliability and energy efficiency when compared to a traditional PC. Dell Wyse desktop devices and software help streamline the delivery of VMware Horizon View 5.2 infrastructure to millions of users around the world. Thin Clients create a more secure environment that minimizes or eliminates exposure to data loss, viruses and malware. By utilizing thin clients as the access device for end user, deployments can benefit from centralized management and complete control of all endpoints. Since thin clients eliminate components with high failure rates, deployments can expect reduced costs and improved reliability over the life of a desktop virtualization deployment.

## 11.1 Dell Wyse P25

Experience uncompromised computing, with the benefits of secure, centralized management. The Dell Wyse P25 PCoIP zero client for VMware Horizon View 5.2 is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. Smaller than a typical notebook, this dedicated zero client is designed specifically for VMware Horizon View 5.2. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon, and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Dell Wyse P25 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

Figure 15.    Dell Wyse P25 and Display Recommendations



Click **HERE** for more information on the Dell Wyse P25.

Display recommendations for the P25 Zero Client are listed below



The P2412H shown above supports 1920x1080, VGA, DVI and USB. Other options include the E2213 with 1680x1050, VGA and DVI, and the E1913 with 1440x900, VGA and DVI.

## 11.2  Dell Wyse Z50D

Designed for power users, the new Dell Wyse Z50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the Z50D combines Dell Wyse-enhanced SUSE Linux Enterprise with a dual-core AMD 1.6 GHz processor and a revolutionary unified engine for an unprecedented user experience. The Z50D eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering unmatched processing speed and power, security and display performance, it's no wonder no other thin client can compare.

Figure 16.  Dell Wyse Z50D and Display Recommendations



Click **HERE** for more information on the Dell Wyse Z50D.

Display recommendations for the Z50D are listed below



The U3011 shown above supports 2560x1600, VGA, DVI, DP and USB. Other options include the P2212H with 1920x1080, VGA, DVI and USB, and the E2213 with 1680x1050, VGA and DVI.

## 11.3  Dell Wyse P45

Experience uncompromised computing with the benefits of secure, centralized management. The Dell Wyse P45 PCoIP zero client for VMware View is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. About the size of a notebook, this dedicated zero client designed specifically for VMware View. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of display performance available over 4 HD displays in a compact, energy-efficient form factor. The Dell Wyse P45 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

Figure 17.    Dell Wyse P45 and Display Recommendations



Click **HERE** for more information on the Dell Wyse P45.

Display recommendations for the P45 are listed below



The U3011 shown above supports 2560x1600, VGA, DVI, DP and USB. Other options include the U2212HM with 1920x1080, VGA, DVI and USB, and the U2711 with 2560 x 1440 (WQHD).

# 12 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. DVS Enterprise - 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.

Based on the best practices from VMware and DVS Engineering validation work the following scaling guidelines have been established for the virtual management infrastructure and the hosts.

Table 11. Management Host and VM Scaling and Sizing recommendations

| Number of View Sessions | Physical Management Hosts | View Connection Server | Virtual Center & Composer | Citrix Netscaler Option | MS Server 2008 R2 Standard License | MS SQL Server 2008 R2 Standard | vCPU | vRAM (GB) |
|---|---|---|---|---|---|---|---|---|
| 0-500 | Integrated into Compute Host vSphere Cluster | 2 | 1 | | 4 | 1 | 14 | 36 |
| 1000 | 2 x ESXi Management Hosts | | | 2 Optional | | | | |
| 2000 | No add | | 1 | | 1 | | 2 | 8 |
| 3000 | No add | 1 | | | 1 | | 4 | 10 |
| 4000 | Add 1 Management Host | | 1 | | 1 | | 2 | 8 |
| 5000 | No add | 1 | | | 1 | | 4 | 10 |
| 6000 | No add | | 1 | | 1 | | 2 | 8 |
| 7000 | Add 1 Management Host | 1 | | | 1 | | 4 | 10 |
| 8000 | No add | | 1 | | 1 | | 2 | 8 |
| 9000 | No add | 1 | | | 1 | | 4 | 10 |
| 10000 | No add | | 1 | | 1 | | 2 | 8 |
| Total Resourc | 4 ESXi Management hosts | 6 | 6 | 2 | 12 | 1 | 40 | 116 |
| 4 | 4 ESXi Management hosts Physical AVAILABLE resources at | | | | | | 51 | 307 |

Based on the DVS Engineering validation work, the per host density numbers below provide an idea of the number of users each host can accommodate. The Basic, Standard and Premium user workloads are generated by LoginVSI (from Login Consultants), the intensity of activity and applications increases as you go up through the workloads, giving real world relevance to the validation. There is of course always a ratio of basic, standard or premium users so actual densities will vary in production.

Table 12.    User Density per M620 based on User Type

| User Type | Horizon View 5.2 Per Host User (Non-Persistent) Density |
|---|---|
| Basic | 135 |
| Standard | 110 |
| Premium | 90 |

The following table simply identifies the number of blades which a PS6110XS. This sizing is approximate and number of session achieved will vary from customer to customer based on the workload generated by the end users.

Table 13.    Blade Density per PS6110XS for Tier 1 storage

| Number of Blades | Number of PS6110XS (Tier 1)  arrays |
|---|---|
| 8 | 1 |
| 16 | 2 |
| 24 | 3 |
| 32 | 4 |

Tier2 storage needs to be identified on a customer by customer basis, Section 8.4 identifies the minimum requirements, as user density increases Tier 2 will need to be sized appropriately to meet the individual customer's needs, DVS Engineering recommends a minimum of 2TB for every 500 users. A Windows File Server will support approximately 1,000 users.

# 13 End-User Workload Characterization

When designing a Desktop Virtualization Infrastructure it is critical to classify and assess the user workload requirements. There are three desktop workload classifications, each of which is bound by specific metrics and capabilities used by the DVS Engineering team to validate the end to end solutions.

Table 14.  User Workload Settings

| User Workload | Virtual CPU | VM Memory Allocation | VM Memory Reservation | OS Image |
|---|---|---|---|---|
| Basic | 1 | 1GB | 0.5GB | Shared |
| Standard | 1 | 1.5GB | 1GB | Shared |
| Premium | 1 | 2.5GB | 1.5GB | Shared |

## 13.1 Dell DVS vDesktop User Workloads

The tasks undertaken by the users in each of the workloads are outlined below;

### 13.1.1 Basic Workload

This workload emulates a typical task worker

- The light workload is very light in comparison to Standard

- Only 2 apps are open simultaneously

- Only apps used are IE, Word and Outlook

- Idle time total is about 1:45 minutes

### 13.1.2 Standard Workload

This workload emulates a medium knowledge working using Office, IE and PDF.  Once a session has been started the medium workload will repeat every 12 minutes.  During each loop the response time is measured every 2 minutes.  The medium workload opens up to 5 apps simultaneously.  The type rate is 160 ms for each character.  Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007, browse 10 messages.

- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsing Wired.com, Lonelyplanet.com and heavy flash app gettheglass.com.

- Word 2007, one instance to measure response time, one instance to review and edit document.

- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.

- Excel 2007, a very large randomized sheet is opened.

- PowerPoint 2007, a presentation is reviewed and edited.

- 7-zip: using the command line version the output of the session is zipped.

### 13.1.3  Premium Workload

The Premium workload is based on the standard workload; the differences in comparison to the Standard workload are:

- Type rate is 130 ms per character.

- Idle time total is only 40 seconds.

- The heavy workload opens up to 8 apps simultaneously

In order to ensure the validation environment is as representative of a production environment as possible Anti-Virus is deployed to all servers and virtual desktops.  The validation environment utilizes McAfee VirusScan Enterprise 8.7i, managed centrally by the ePolicy Orchestrator which controls agent deployment, definition updates and policies.

Since the Virtual desktops are based on a single read-only image streamed from the provisioning servers there are a number of exclusions that are configured to optimize performance of the On-Access scanning per McAfee and **VMware** best practices.

# 14 Delivery Model

This Reference Architecture can be purchased as a complete solution, the Dell DVS Enterprise - Active System 800v.  This solution is available to be partially 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 17 below shows the DVS Enterprise - Active System 800v solution with a single chassis. Figure 18 shows DVS Enterprise - Active System 800v with two chassis and maximum storage enclosures. Note that all EqualLogic arrays shown in the figures are PS6110X/XS. 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 18.** *DVS Enterprise – Active System 800v Single Rack Overview*



Force10 S55
2x Force10 S4810

2x PowerEdge R620
KMM

4x EqualLogic PS6110 Arrays

PowerEdge M1000e

16x PowerEdge M620

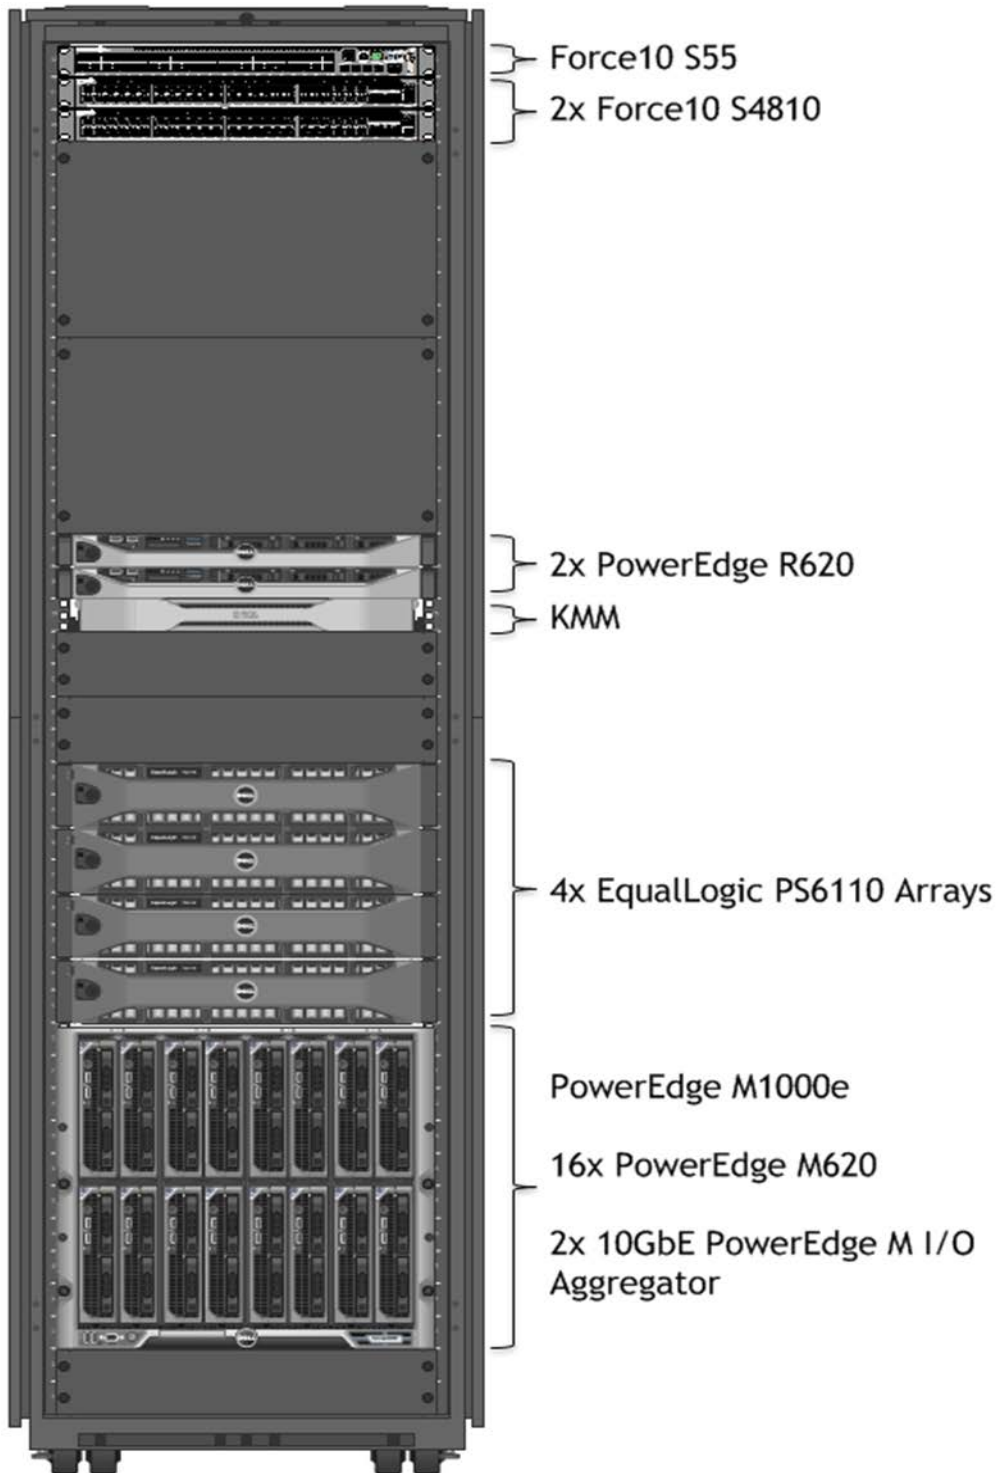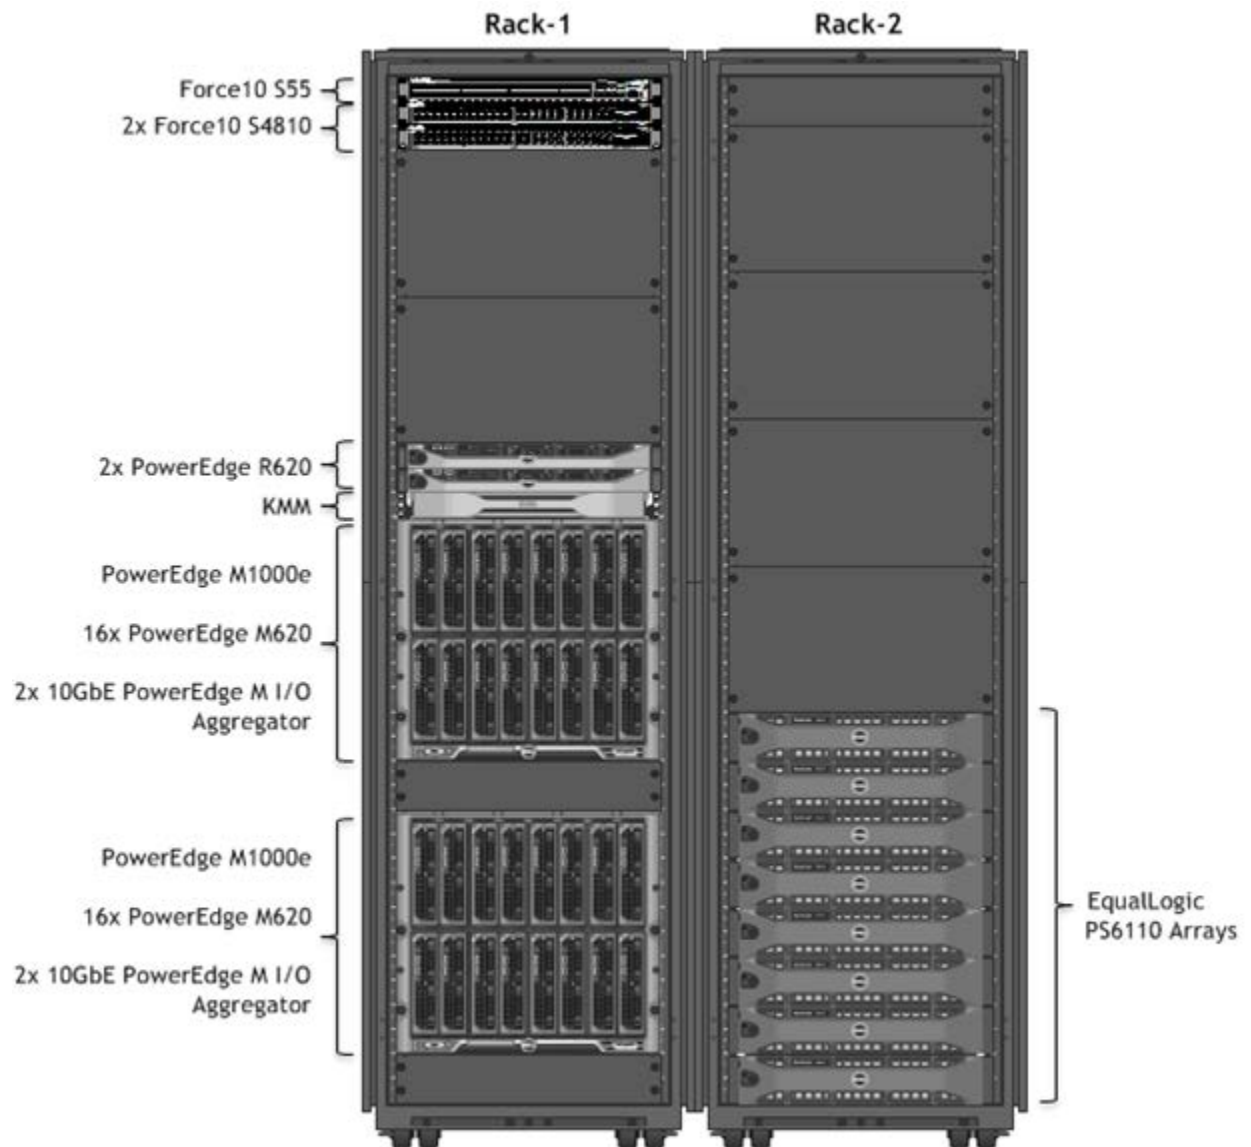2x 10GbE PowerEdge M I/O Aggregator

**Figure 19.** DVS Enterprise - Active System 800v Two Chassis and Maximum Storage: Rack Overview

# 15 Appendix A – Dell Mobile Clinical Computing - AlwaysOn Point of Care

Dell and VMware Solution Architecture - Secure, Reliable and Continuous Access to EMR and Patient Care Applications

Digitized medical records can reduce medical errors, improve patient safety and produce better clinical outcomes—but if the caregiver cannot access Electronic Medical Record (EMR) and clinical information systems efficiently or are required to change workstations to access particular applications, productivity suffers and user frustration can increase. Furthermore, security and HIPAA compliance requirements may stipulate that users provide login credentials for each application that is accessed, then exit from each application and log off the system when finished. These processes are cumbersome and time consuming, and the log off process may be overlooked altogether creating potential security and data protection issues.

When electronic devices replace paper charts and physician prescription pads having quick secure and reliable access to the various applications from any device in clinical workspaces becomes critical. This is why modernization of the point-of-care desktop has become an urgent priority for hospital IT professionals as well as hospital caregivers. Desktops and patient care applications must be immediately accessible and available to caregivers regardless of workstation, device or location.

Dell Mobile Clinical Computing (MCC) and VMware have a solution. The Dell Mobile Clinical Computing – VMware AlwaysOn Point of Care™ solution provides an innovative virtualization architecture solution for the specific application access and stringent security requirements of medical organizations. By virtualizing desktops with VMware Horizon View 5.2™ and hosting them on VMware vSphere™, a key component of Dell's Mobile Clinical Computing architecture, and using this validated architectural design, healthcare organizations can now have unparalleled desktop and application reliability and availability – and secure rapid access from virtually any device. The caregiver can rapidly and securely access their desktop and application workspace by tapping a proximity card, swiping their finger on a biometric device or entering their username and password and they will be immediately presented with their clinical desktop to resume caring for patients.

The Dell MCC - AlwaysOn Point of Care offering allows clinicians and staff to achieve the rapid, secure access and continuous level of availability they demand. In a typical IT environment only the production environment is backed up. But by having an active-active desktop environment running identical desktop images, even if there is a failure at the primary site, end users can promptly access their desktops and applications. If a healthcare provider's infrastructure is compromised through a natural disaster or other unexpected outage, caregivers—who are many times among the first responders—can be assured they can reach their clinical desktops and applications where and when they are needed the most.

This new Dell and VMware architectural design features continuous monitoring capabilities, as well as load balancing with constant data replication across sites to ensure that if the primary site is down VMware Horizon View 5.2 will seamlessly failover to the secondary site so the caregiver experiences minimal disruption. As a result, IT can now deliver non-stop point of care desktops with all applications and data readily available where and when they are needed most.

## 15.1  Benefits

- •Validated Dell and VMware solution architecture

- •Rapid secure non-stop access to clinical desktops and Electronic Medical Record applications

- •Provides constant monitoring, load balancing and data replication features for optimizing performance and uptime

- •Single sign on access, session mobility, highly reliable

- • End-to-end Dell and VMware Horizon View 5.2 implementation with Dell Mobile Clinical computing Infrastructure

- •High performance, secure, reliable Dell Wyse thin clients

## 15.2  Solution Elements

Dell and VMware have jointly validated the Dell Mobile Clinical Computing - AlwaysOn Point of Care offering. This purpose built architecture integrates Dell's Mobile Clinical Computing solution with VMware and additional Dell technology partner solutions to meet the uniquely stringent need for security, data protection, flexible scaling, and multiple single sign on mechanisms for health care organizations.

### 15.2.1  Imprivata OneSign

Imprivata OneSign® provides Single Sign On and strong authentication, permitting users to access all workstations and applications they are authorized to use. By configuring and linking multiple instances of virtual appliances at both sites with fault tolerance, during the site failover, desktop agents can continuously look up the next working instance without disrupting the workflow. Users easily connect to virtual desktops and applications via Imprivata OneSign single sign on from the access point all the way through to the EMR application.

### 15.2.2  Dell and VMware Solutions for secure, reliable and continuous access to EMR and patient care applications

Dell and VMware have created the Dell Mobile Clinical Computing - AlwaysOn Point of Care of Care offering - an innovative virtualization architecture solution for the specific application access and stringent security requirements of medical organizations. It integrates technology and solutions from Dell, Dell Wyse, VMware and our thriving ecosystem of technology partners.

The solution leverages and integrates Dell's Mobile Clinical Computing and a highly available VMware virtualization infrastructure, with single sign on and secure high performance thin clients validated to health care industry requirements. Together, Dell, VMware, and our partners can deliver rapid and secure access to high performance continuous desktops, critical medical applications and records. More than that, we deliver true peace of mind for medical professionals—and the highest levels of patient care.

For additional information about how the Dell Mobile Clinical Computing with VMware AlwaysOn Point of Care offer is built and validated, please consult the Reference Architecture at:

www.dell.com/virtualdesktops    or    http://www.vmware.com/solutions/industry/healthcare/point-of-care.html or contact you Dell Partner Direct Reseller www.dell.com/partnerdirect

# 16 Appendix B - Dell DVS Enterprise - Mobile Secure Desktop

Dell and VMware Solution Architecture - Mobile, Secure Access to Applications and Data

Today's workforce is no longer tethered to traditional stationary desktops. New devices have proliferated in organizations of all sizes. Employees are increasingly mobile, and more than 60% of enterprise firms and 85% of SMB organizations are looking to initiate Bring Your Own Device (BYOD) programs. But while end users are eager to embrace "bring your own device" trends, IT departments—faced with tight budgets—are challenged with how to best support and manage these new devices while protecting corporate data as it is accessed across networks and locations.

Dell and VMware have a solution. The Dell Desktop Virtualization Solutions (DVS) Enterprise - Mobile Secure Desktop architecture provides an innovative way for IT to support device diversity and BYOD initiatives by improving user access and mobility, streamlining application updates, enhancing data security, and delivering the highest-fidelity user experience. This solution virtualizes desktops and hosts them on VMware® vSphere™, a key component of VMware Horizon View 5.2™, and uses Dell's DVS Enterprise infrastructure so organizations can now have unparalleled desktop and application access across devices and locations.

And because this solution ties desktop environments to user identities instead of devices, end users are free to access their data and applications from any qualified device, whether in the office or halfway around the world. This solution addresses three key requirements including:

## 16.1 Mobility

The Mobile Secure Desktop solution is built on Dell DVS Enterprise and VMware Horizon View 5.2. It places desktops in the datacenter and provides secure access to applications and data from a multitude of client devices including PC workstations, thin clients and mobile device. This enables BYOD support with persistence for true session mobility across devices—so users can access the same desktop from different devices. With persona management and optional support for user-installed applications, the Mobile Secure Desktop solution provides a personalized user experience across devices and sessions. In addition to providing session persistence across devices, VMware Horizon View 5.2 uses PCoIP protocol to deliver the best desktop user experience from any device.

## 16.2 Security

By integrating support for two-factor authentication (RSA SecurID, RADIUS authentication), the DVS Enterprise - Mobile Secure Desktop solution emphasizes data and application security. In addition to providing the right level of access to the right resources, it also simplifies patch management and update management. Since all the desktops are hosted in a centralized infrastructure, this solution streamlines deployment of updates and patches to desktops helping ensure that no vulnerabilities exist in the environment due to outdated, unpatched or orphaned systems. This solution also integrates VMware vShield™ to provide superior security for the environment.

## 16.3 Management

One of the key challenges facing organizations today is the ability to monitor and manage the entire end user compute environment including; desktops, access policies and service levels. The DVs Enterprise - Mobile Secure Desktop solution with optionally integrated vCenter Operations Manager (VCOps), provides an integrated dashboard with intelligent response on all desktop-related events. This helps IT admins to provide the right amount of intervention and guidance when the demands on the virtual infrastructure appear to exceed an expected range of behavior. The solution can also include vCenter Configuration Manager (vCM) for importing suggested configurations and to meet required regulatory compliance standards.

A critical aspect to mobile secure desktops and true BYOD support is securely managing and enabling corporate access. Dell Wyse project Stratus offers simple, secure, cloud-based management for today's dynamic IT environment. The integration of Dell Wyse Stratus into the Mobile Secure Desktop solution provides IT administrators with an intelligent and dynamic cloud-based console to securely manage and enable corporate access to any device, including smartphones, tablets, thin clients, zero clients, and PC's – regardless of whether the device is owned by the company or by an individual employee.

## 16.4 Solution Elements

Dell and VMware have jointly validated the Dell DVS Enterprise - Mobile Secure Desktop offering. This purpose built architecture integrates Dell's DVS Enterprise solution with VMware and additional Dell technology partner solutions to meet the need for flexible, mobile and secure end user computing in today's dynamic business environment.

### 16.4.1  Compliance - VMware vCenter Configuration Manager

A key requirements for many organizations is managing compliance to various government and industry regulations. vCenter Configuration Manager (vCM) automates critical configuration and compliance management tasks including configuration data collection, configuration change execution, configuration reporting, change auditing, and compliance assessment.

### 16.4.2  Cortado ThinPrint

Most of the use cases supported by this solution have a location-aware printing requirement. Cortado ThinPrint software, OEM'd by VMware, provides the ability to take advantage of location-aware printing from a wide range of devices.

### 16.4.3  Imprivata OneSign

Imprivata OneSign® provides Single Sign On and strong authentication, permitting users to access all workstations and applications they are authorized to use. By configuring and linking multiple instances of virtual appliances at both sites with fault tolerance, during the site failover, desktop agents can continuously look up the next working instance without disrupting the workflow. Users easily connect to virtual desktops and applications via Imprivata OneSign single sign on from the access point all the way through to the EMR application.

**Dell and VMware solutions for mobile, secure access to applications and data**

Dell and VMware have created the Dell DVS Enterprise - Mobile Secure Desktop solution – a validated architecture that integrates technology from Dell, VMware and the Dell technology partner ecosystem. The solution leverages Dell server, storage and networking infrastructure, mobile, wireless and wired networks, VMware Horizon View 5.2 and vSphere, vShield security services, management and monitoring components to protect data, monitor the infrastructure and secure access for virtual any end point device. This solution is optimized for organizations looking to drive higher levels of productivity by improving end-user access across devices and locations, reduce costs by streamlining desktop and application management, and enhance security and compliance.

The solution leverages and integrates Dell's DVS Enterprise end to end desktop virtualization infrastructure and highly available VMware virtualization platforms, with single sign on, personal management, and additional high performance industry recognized solutions. Together, Dell, VMware, and our partners deliver rapid and secure access to high performance desktops, business applications and corporate data with the Dell DVS Enterprise - Mobile Secure Desktop solution.

For additional information about how the Dell DVS Enterprise - Mobile Secure Desktop solution is built and validated, please consult the Dell DVS Enterprise - Mobile Secure Desktop Reference Architecture at www.dell.com/virtualdesktops or http://www.vmware.com/solutions/desktop/mobile-secure-desktop/overview.html or contact your Dell Partner Direct Reseller www.dell.com/partnerdirect

# 17 Reference

Dell Active Infrastructure reference:

- [Dell Active System](#)

- [Dell Active Infrastructure Wiki](#)

VMware references:

- [VMware vSphere Edition Comparisons](#)

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

- [VMware vSphere 5.1 documentation](#)

Dell PowerEdge References:

- [Dell PowerEdge M1000e Technical Guide](#)

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

- [NIC Partitioning (NPAR)](#)

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](#)