

---

# Reference Architecture for Active System 1000 with VMware vSphere


---

*Release 1.1 for Dell PowerEdge Blade Servers, Dell Networking Switches, Dell Compellent Storage Center, and Dell Active System Manager*

Dell Virtualization Solutions Engineering

Revision: A00

July 2013



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

© 2013 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. *Dell*, the *Dell logo*, *OpenManage*, *Compellent*, *Kace*, *EqualLogic*, *PowerVault*, and *PowerEdge* are trademarks of Dell Inc. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries. *Microsoft*, *Windows*, *Hyper-V*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *VMware*, *vSphere*, *ESXi*, *vMotion*, *vCloud*, and *vCenter* are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. *Linux* is the registered trademark of Linus Torvalds in the U.S. and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

July 2013

### Revision History

Revision	Description	Date
A00	Initial version	July 2013

## Table of Contents

1	Introduction.....	3
2	Audience.....	3
3	Solution Overview .....	4
4	Design Principles .....	13
5	Prerequisites and Datacenter Planning .....	13
6	Architecture .....	15
6.1	Dell Blade Network Architecture .....	16
6.2	Network Connectivity .....	17
6.3	Network Configuration .....	18
6.4	Storage Architecture .....	21
7	Management Infrastructure .....	23
7.1	Dell Active System Manager .....	24
7.2	Dell OpenManage Essentials (OME).....	25
7.3	Dell Repository Manager (DRM) .....	25
7.4	Dell Management Plug-in for VMware vCenter (DMPVV) .....	25
7.5	Dell Compellent Enterprise Manager .....	26
7.6	VMware vCloud Connector.....	26
8	Connecting Active System 1000 to Datacenter Network.....	26
8.1	Connecting the S55 OOB switch to Datacenter Network .....	26
8.2	Connecting to a Dell Networking Datacenter Network .....	27
8.3	Connecting to a Cisco Nexus Datacenter Network .....	28
9	Scalability .....	29
10	Delivery Model .....	29
11	Additional Supported Configurations .....	33
11.1	Dell Networking S5000 as LAN Top of Rack Switch.....	33
11.2	Cisco Nexus 5548 as LAN Top of Rack Switch .....	34
12	Reference.....	36

## Figures

Figure 1: Active System 1000v Overview .....	4
Figure 2: Dell Blade Servers, Dell Networking Switches, Brocade Fibre Channel Switches, and Dell Compellent Storage .....	5
Figure 3: Active System 1000v Network Topology (Logical View) .....	15
Figure 4: I/O Connectivity for PowerEdge M620 Blade Server .....	16
Figure 5: vSwitch, NPAR and I/O Module Configuration .....	19
Figure 6: Fibre Channel SAN Logical Connectivity .....	21
Figure 7: Fibre Channel SAN Logical Connectivity with Additional Storage Controllers .....	22
Figure 8: Management Components .....	24
Figure 9: S55 Connectivity to Datacenter OOB Network (Default Option) .....	27
Figure 10: Alternative OOB Connectivity - S55 Switch to S4810 Switches .....	27
Figure 11: Active System 1000v connectivity to Dell Networking Z9000 Switch .....	28
Figure 12: Active System 1000v connectivity to Cisco Nexus 5548 .....	28
Figure 13: Active System 1000v Single Chassis: Rack Overview .....	30
Figure 14: Active System 1000v Two Chassis and Maximum Storage: Rack Overview .....	31
Figure 15: Active System 1000v Configuration for additional storage requirements: Rack Overview ....	32
Figure 16: Active System 1000v using S5000 as ToR Switch .....	34
Figure 17: Active System 1000v using Cisco Nexus 5548 as ToR Switch .....	35

## Tables

Table 1: Solution Components .....	6
Table 2: VLAN Overview .....	20

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

The Active System 1000 solution, a member of Dell Active Infrastructure family, is an infrastructure solution that has been designed and validated by Dell Engineering. It is available to be racked, cabled, and delivered to your site to speed deployment. Dell Services will deploy and configure the solution tailored for business needs, so that the solution is ready to be integrated into your datacenter. Active System 1000 is offered in configurations with either VMware® vSphere® (Active System 1000v) or Microsoft® Windows Server® 2012 with Hyper-V® role enabled (Active System 1000m) hypervisors. This paper defines the Reference Architecture for the VMware vSphere-based Active System 1000v solution.

Active System 1000v includes Dell PowerEdge™ M1000e blade chassis with Dell I/O modules (Dell PowerEdge™ M I/O Aggregator or Dell Networking MXL switch), Dell PowerEdge™ M620 blades, Dell Compellent™ Storage, Dell Networking network switches, Brocade Fibre Channel switches, and VMware vSphere 5.1 Update 1. The solution also includes Dell PowerEdge™ R620 servers as management servers. Dell Active System Manager, VMware vCenter Server, Compellent Enterprise Manager, and Dell OpenManage™ Essentials are included with the solution.

Dell Active System Manager is the Active Infrastructure management software, a key component of the Active System 1000v. Active System Manager streamlines physical and virtual workload provisioning through automation and standardization. Through capabilities, such as template-based provisioning, end-to-end automation, workflow orchestration and resource pooling, Active System Manager enables IT to respond rapidly to business needs, maximize data center efficiency, and strengthen quality of IT service delivery. In Active System 1000v, the Dell Active System Manager virtual appliance is deployed within the management cluster. Templates can be used to deploy the compute clusters.

## 2 Audience

IT administrators and IT managers, who have purchased or are planning to purchase an Active System configuration, can use this document to understand the component details of the solution.

### 3 Solution Overview

This section provides a high-level product overview of the VMware vSphere, Dell PowerEdge blade servers, Dell I/O modules, Dell Networking S4810, S55, Brocade 6510 Fibre Channel Switches, and Dell Compellent Storage, as illustrated in Figure 1 and Figure 2. Readers can skip the sections of products with which they are familiar.

Figure 1: Active System 1000v Overview



**VMware vSphere 5.1**

- vMotion, Storage vMotion
- VMware HA and DRS

**Dell PowerEdge Blade Servers**

- Energy efficient PowerEdge M1000e enclosure
- 12<sup>th</sup> generation M620 blade server
- Flex Address
- CMC and iKVM for enclosure management

**Dell PowerEdge M I/O Aggregator or Dell Networking MXL switch**

- Highest Performance in a Blade Switch
- Highest Density in a Single Blade Switch
- Scalability & Modular to Fit Your Business

**Dell PowerEdge Rack Servers for Management Cluster**

- 12<sup>th</sup> 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 Networking S4810 Switches**

- 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

**Dell Networking 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

**Brocade 6510 Fibre Channel Switches**

- 1U 48 port high density switch with 8Gbps Fibre Channel

**Dell Compellent Storage Center**

- Fluid Data Architecture with Thin Provisioning and Automated Tiered Storage
- Centralized management using Enterprise Manager
- Fast Track

**Integrated Management**

- Dell Active System Manager
- Dell Management plug-in for VMware vCenter
- Compellent vSphere plug-in
- Compellent Enterprise Manager
- OpenManage Essentials

**Cloud Enablement**

- VMware vCloud Connector for Dell vCloud connectivity

Figure 2: Dell Blade Servers, Dell Networking Switches, Brocade Fibre Channel Switches, and Dell Compellent Storage

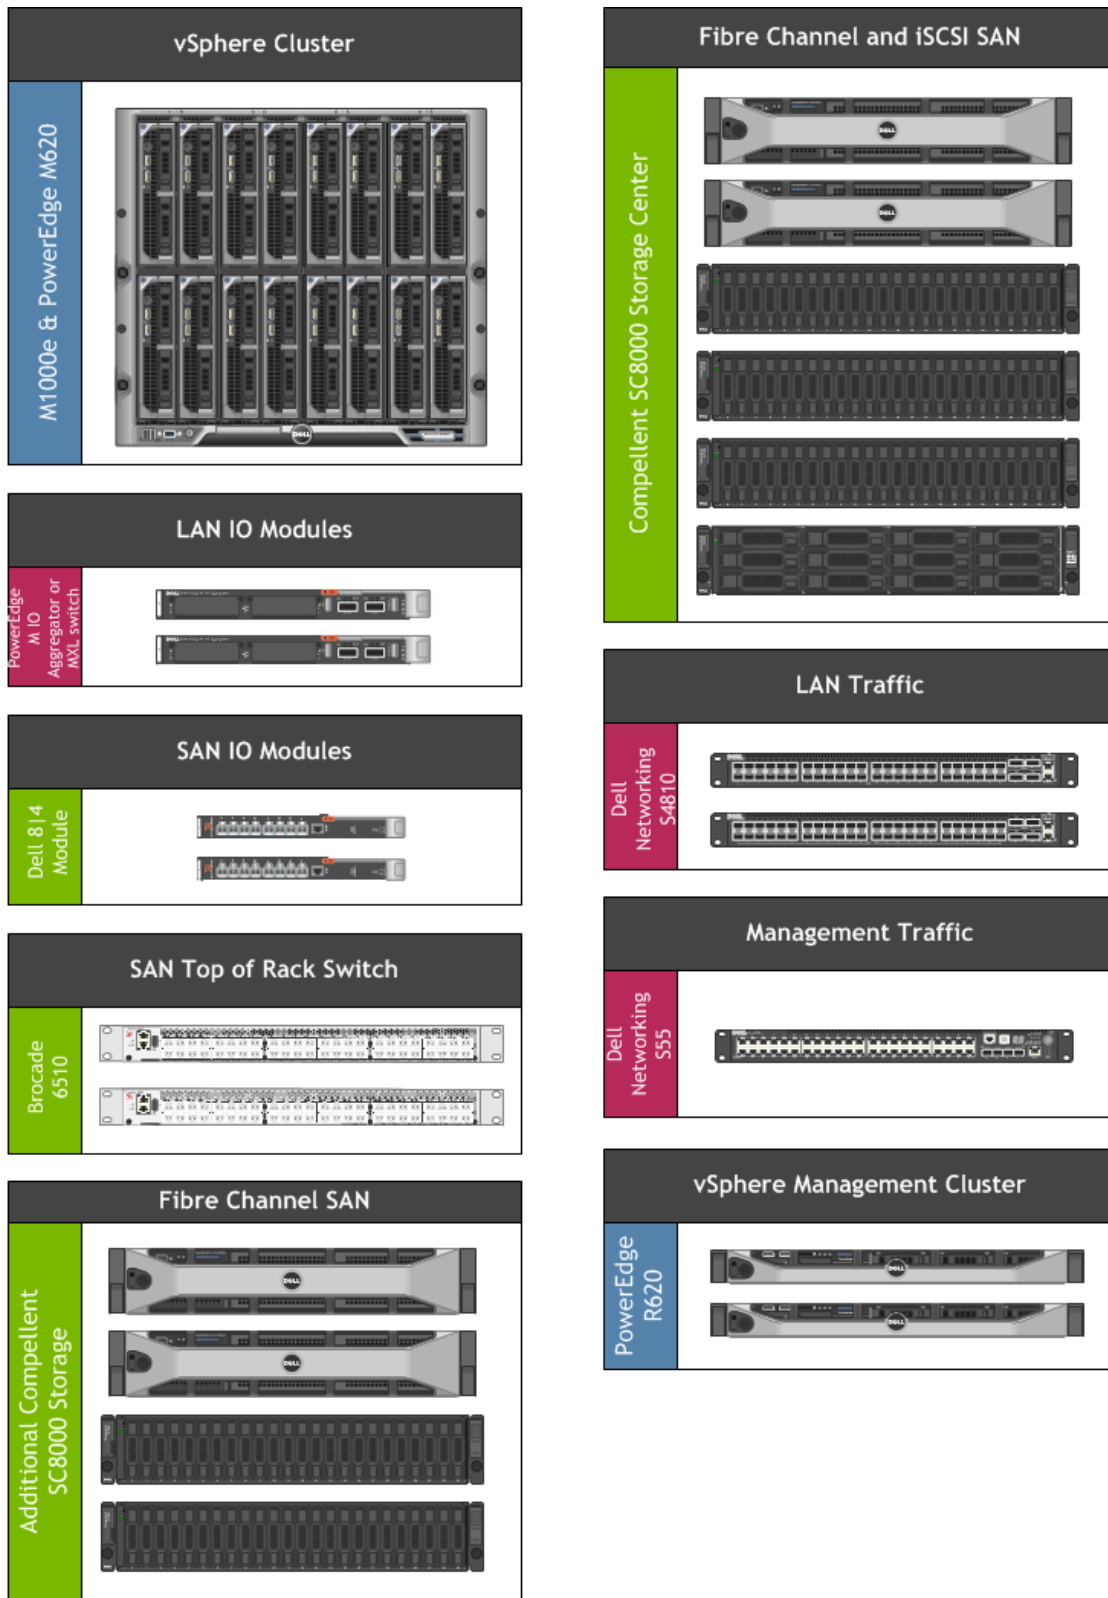




Table 1 below describes the key solution components and the roles served.

**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 Update 1
LAN Switch	Two Dell Networking S4810 switches 2x Dell PowerEdge M I/O Aggregator or Dell Networking MXL switches in each Dell PowerEdge M1000e chassis
SAN Switch	Two Brocade 6510 Fiber Switch 2x Dell 8   4 I/O modules per chassis
Storage	Dell Compellent SC8000 controllers with SC200/220 enclosures
Management Infrastructure	Two PowerEdge R620 servers hosting management VMs. One Dell Networking S55 used as a 1Gb management switch
Management components hosted in the management infrastructure	<ul style="list-style-type: none"> <li>• Dell Active System Manager</li> <li>• VMware vCenter Server</li> <li>• Dell Management Plug-in for VMware vCenter</li> <li>• Dell OpenManage Essentials</li> <li>• Compellent Enterprise Manager</li> <li>• Compellent Plugin for VMware vCenter</li> <li>• VMware vCloud Connector</li> <li>• Dell Repository Manager</li> </ul>

**VMware vSphere 5.1 Update 1:** VMware vSphere 5.1 Update 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 VM level. Upon host failure, VMware HA automatically re-starts VMs on other physical hosts running ESXi. VMware vSphere 5.1 Update 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.
- **VMware vCenter Update Manager:** VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.
- **VMware Storage vMotion™:** VMware Storage vMotion enables real-time migration of running VM disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.
- **Host Profiles:** Host Profiles standardize and simplify the deployment and management of VMware ESXi host configurations. They capture and store validated configuration information, including host compliance, networking, storage, and security settings.

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

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

The key capabilities of Dell Active System Manager are:

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

- Centralized Management – Intuitive centralized, role-based management and access through self-service web portal

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

**Dell Management Plugin 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 Plugin for VMware vCenter](#).

**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, EqualLogic™ and PowerVault™ storage, and 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 Windows®, Linux®, VMware® and Hyper-V® environments

For more information on OpenManage Essentials, see [Dell.com/openmanageessentials](http://Dell.com/openmanageessentials).

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

**Dell PowerEdge Blade Servers:** The Dell PowerEdge M620 blade server is the 12<sup>th</sup> generation 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 16GB DIMMs, or 768GB or RAM per blade - 12TB or RAM in a single PowerEdge 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 on the PowerEdge M620 offers three modular choices for embedded fabric capability. With 10Gb Converged Network Adapter (CNA) offerings from Broadcom, QLogic, and 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.

**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 Dell PowerEdge 12<sup>th</sup> generation servers. It includes 1GB of managed and persistent storage that embeds systems management features directly on the server, thus eliminating the media-based delivery of system management tools and utilities previously needed for systems management. Embedded management includes:

- Unified Server Configurator (USC) aims at local 1-to-1 deployment via a graphical user interface (GUI) for operating system install, updates, configuration, and for performing diagnostics on single, local servers. This eliminates the need for multiple option ROMs for hardware configuration.
- Remote Services are standards-based interfaces that enable consoles to integrate, for example, bare-metal provisioning and one-to-many OS deployments, for servers located remotely. Dell's Lifecycle Controller takes advantage of the capabilities of both USC and Remote Services to deliver significant advancement and simplification of server deployment.
- Lifecycle Controller Serviceability aims at simplifying server re-provisioning and/or replacing failed parts and thus reduces maintenance downtime.

For more information on Dell Lifecycle Controllers and blade servers, see

<http://content.dell.com/us/en/enterprise/dscm-embedded-management> and [Dell.com/blades](http://Dell.com/blades).

**Dell PowerEdge M I/O Aggregator:** The Dell PowerEdge M I/O Aggregator (M I/OA) is a flexible 1/10GbE aggregation device that is automated and pre-configured for easy deployment into converged iSCSI and FCoE (Fibre Channel over Ethernet) networks. The key feature of the PowerEdge M I/OA is that all VLANs are allowed as a default setting. This allows the top-of-rack (ToR) managed switch to perform all

VLAN management related tasks. The external ports of the PowerEdge M I/OA are automatically all part of a single link aggregation group (LAG), and thus there is no need for Spanning-tree. The PowerEdge M I/OA can use Data Center Bridging (DCB) and Data Center Bridging Exchange (DCBX) to support converged network architecture.

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

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

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

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

**Dell Networking MXL 10/40GbE Blade Switch:** The MXL switch provides 1/10/40GbE. The switch supports 32 internal 1/10GbE ports, as well as two fixed 40GbE QSFP+ ports and offers two bays for optional FlexIO modules. To ensure room to grow, uplinks via the FlexIO modules can be added or swapped as needed in the future. Choose from 2-port QSFP+, 4-port SFP+ or 4-port 10GBASE-T FlexIO modules to expand and aggregate (bi-directional) bandwidth up to 160 Gigabit per second. The MXL switch provides the flexibility to mix and match the FlexIO module types.

Like the M I/OA above, the MXL switch includes two integrated 40Gb Ethernet ports on the base module. These ports are 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 MXL Switch provides stacking capability for up to six interconnected blade switches allowing both stacking across chassis and local switching of traffic within the chassis. For more information, see <http://www.dell.com/us/business/p/force10-mxl-blade/pd>.

**Dell Networking S4810 Switches:** The Dell Networking 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 switch delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact S4810 switch 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 S4810 switch ideally suited for iSCSI storage environments. In addition, the S4810 switch incorporates multiple architectural features that optimize data center network flexibility, efficiency, and availability, including Dell Networking 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 Dell Networking switches, see <http://www.dell.com/networking>.

**Dell Networking S55:** The Dell Networking 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 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 Dell Networking switches, see <http://www.dell.com/networking>.

**Brocade 6510:** The Brocade 6510 switch is a high density FC switch providing 48 ports in a 1U form factor. The 6510 switch includes redundant power supplies and fans making it well suited to the high availability needs of virtualization infrastructures. It also includes the Ports-on-Demand capabilities for cost reduction when installing in smaller environments.

For more information on Brocade 6510 Fibre Channel Switches, see [Dell.com/brocade](http://Dell.com/brocade).

**Dell 8/4 Gbps FC SAN Module:** The Dell 8/4 Gbps FC SAN Module is a 24-port FC module with eight external ports and 16 internal ports that installs in a Dell PowerEdge M1000e Blade Enclosure. Built on industry-standard N\_Port ID Virtualization (NPIV) technology, the module eliminates the traditional challenges of heterogeneous switch-to-switch interoperability and can non-disruptively connect Dell blades to NPIV-enabled FC SANs, including Brocade, Cisco, McData, and others. The Dell 8/4 Gbps FC SAN Module eliminates incremental switch management and configuration by presenting FC connections as a logical device (rather than switch domains) to the SAN fabric. The module enables the benefits of port aggregation, failover, and redundancy without the complexities of additional SAN switches or additional switch domains.

For more information on Dell 8/4 Gbps FC SAN Module, see [Dell.com/us/enterprise/p/fc-san/pd](http://Dell.com/us/enterprise/p/fc-san/pd).

**Dell Compellent SC8000 Storage Center:** The SC8000 is a 2U Storage Center controller built on the Dell 12th generation PowerEdge™ server platform with a custom configuration to support the needs of the enterprise storage controller. The SC8000 offers increased density, exceptional processing power, greater memory, faster PCIe Gen3 IO bus, improved diagnostics capability with the Integrated Dell Remote Access Controller (iDRAC), and exceptional power efficiency with Energy Star Platinum rated dual hot-swappable power supplies using Fresh Air™ technology. The SC8000 IO expansion consists of 7 PCIe Gen 3 (double the bandwidth of Gen 2) capable slots: 4 full-height and 3 low-profile slots. 1 full-height slot contains the controller cache card. The remaining 3 full-height and 3 low-profile slots are reserved for back-end and front-end IO expansion. The SC8000 controllers are connected to SC200/220 SAS enclosures in daisy chain loop. The SC200/220 are 2U 6Gb SAS enclosures. The SC200 supports up to twelve (12) 3.5” disk drives and SC220 supports up to twenty four (24) 2.5” disk drives. Supported drives can include a mix of drive speed and capacity in any slot within the enclosure.

For more information on Dell Compellent, see [Dell.com/Compellent](http://Dell.com/Compellent).

Features of the Dell Compellent SC8000 Storage Array include:

- **Fluid Data Architecture** - Storage is managed at the most granular level with built-in system intelligence to enable the dynamic flow of enterprise data.
- **Storage Virtualization** - Storage is virtualized at the disk level to create a flexible pool of storage resources shared by all servers all the time.
- **Thin Provisioning** - Allocation is completely separated from utilization so any size volume can be created at any time, yet capacity is only consumed when data is written.
- **Automated Tiered Storage** - Data dynamically cascades from tier to tier according to actual usage, freeing up high-performance drives for mission-critical applications.
- **Space-efficient Replays** - Continuous snapshots only capture changes in data for real-time protection with instant recovery to any point in time.
- **Thin Replication** - Data is replicated between local and remote sites using space-efficient snapshots and native IP or FC connectivity, eliminating the need for high-speed data links or identical system configurations.
- **Unified Storage Resource Management** - All storage resources are managed through a single point-and-click interface, providing a complete view of the entire storage environment.
- **Open, Agile Hardware Platform** - Storage is designed for persistence, not obsolescence, leveraging a single modular hardware platform coupled with technology independence.

**Compellent Enterprise Manager:** Compellent Enterprise Manager is included in the solution. It simplifies network storage management by providing a single, centralized console for the administration of multiple local and remote Compellent systems. Users can configure and verify remote replication processes, monitor storage capacity and disk utilization in real time, and generate comprehensive enterprise storage usage and performance reports.

For more information on Dell Compellent, see [Dell.com/Compellent](http://Dell.com/Compellent). Contact Dell sales representative for more information on Compellent storage configurations and sizing guidelines.

**PowerEdge R620 Management Server:** The Dell PowerEdge R620 uses Intel® Xeon® E5-2600 series processors and Intel chipset architecture in a 1U rack mount form factor. These servers support up to ten 2.5” drives and provide the option for an LCD located in the front of the server for system health monitoring, alerting, and basic management configuration. An AC power meter and ambient temperature thermometer are built into the server, both of which can be monitored on this display without any software tools. The server features two CPU sockets and 24 memory DIMM slots.

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

**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 the public cloud.

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



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

For more information, see [VMware vCloud Connector](#).

## 4 Design Principles

The following principles are central to the design and architecture of the Active System 1000v Solution.

1. **Redundancy with no single point-of-failure:** Redundancy is incorporated in the critical aspects<sup>1</sup> of the solution, including server high availability features, networking, and storage.
2. **Management:** Provide integrated management using VMware vCenter, Dell Management plug-in for VMware vCenter, Dell OpenManage Essentials, Compellent Enterprise Manager and Compellent plug-in for VMware vCenter.
3. **Cloud Enabled:** The solution also includes connectivity to public cloud using VMware vCloud Connector.
4. **Integration into an existing data center:** This architecture assumes that there is an existing 10 Gb Ethernet infrastructure with which to integrate.
5. **Hardware configuration for virtualization:** This solution is designed for virtualization for most general cases. Each blade server is configured with appropriate processor, memory, host bus, and network adapters as required for virtualization.
6. **Racked, Cabled, and Ready to be deployed:** Active System 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.
7. **Power, Cooling, and Weight Considerations:** Active System 1000v 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.
8. **Flexible configurations:** Active System 1000v is pre-configured to suit most customer needs for a virtualized infrastructure. The solution also supports additional options, such as configuring racks, server processors, server memory, and storage, based on customer needs.

## 5 Prerequisites and Datacenter Planning

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

1. An existing Ethernet infrastructure with which to integrate is required. 10Gb or 40Gb Ethernet infrastructure is recommended.
2. Active Directory® (AD) - VMware vCenter can be configured to use Active Directory.
3. Domain Name Server (DNS) must be available on the management network.
4. Network Time Protocol (NTP) Server must be available on the management network.

---

<sup>1</sup> Out of band management is not considered critical to user workload and does not have redundancy.

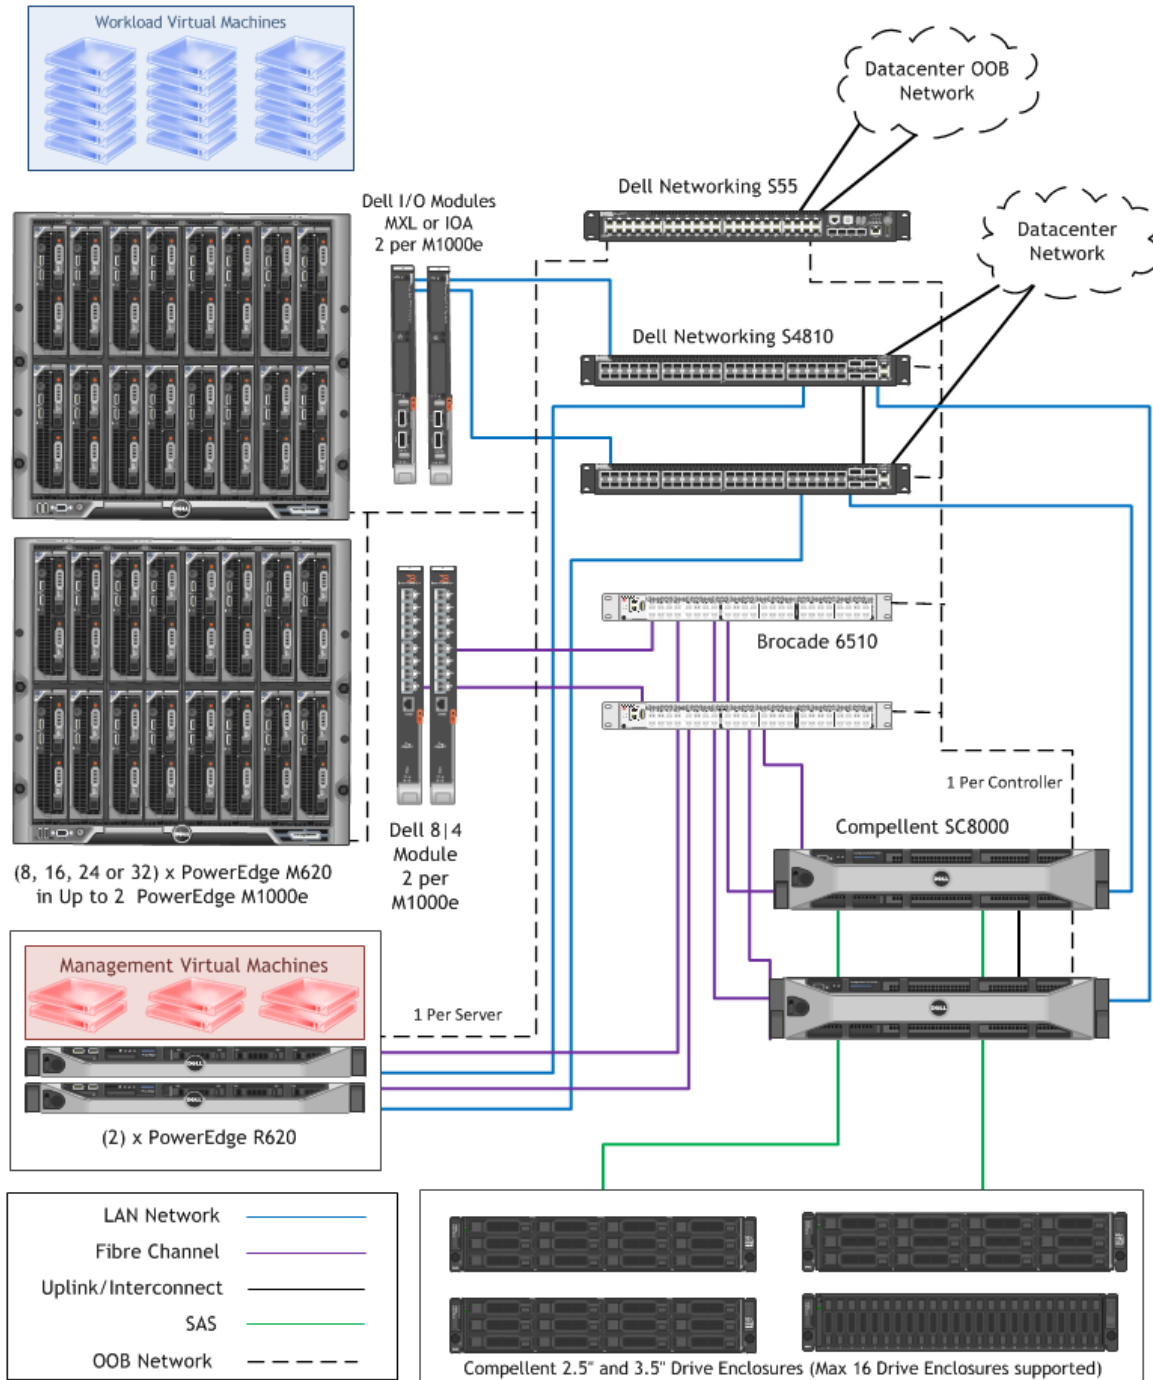


5. SQL Requirements - Required for vCenter, and Compellent Enterprise Manager.
6. SMTP is optionally utilized by other components of the management stack, including OME and vCenter, for notifications
7. Additional components, such as Dell Networking cables and transceivers, are needed to uplink the solution to the customer network. The necessary components depend upon customer networking and uplink requirements.
8. Sufficient power and cooling to support the solution must be present. Detailed power, weight, and cooling requirements for the datacenter are defined in the *Specification Guide for Active System 1000 with VMware vSphere*.

## 6 Architecture

This solution consists of a PowerEdge M1000e chassis populated with PowerEdge M620 blade servers running VMware ESXi. Figure 3 provides high-level reference architecture for the solution. The Figure 3 shows high-level logical connectivity between various components. Subsequent sections provide more detailed connectivity information.

Figure 3: Active System 1000v Network Topology (Logical View)



Note- Compellent Storage configuration is customizable depending upon customer requirements.

## 6.1 Dell Blade Network Architecture

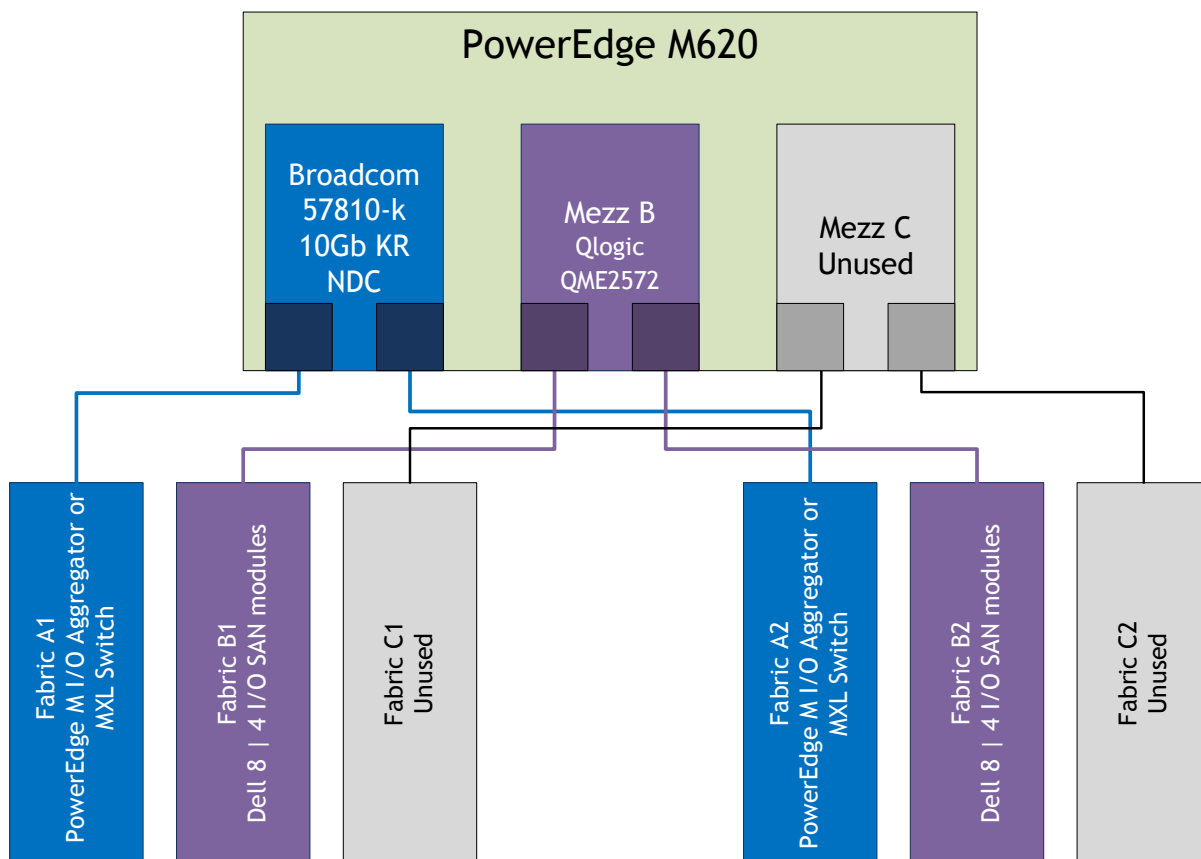
The Dell blade chassis has three separate fabrics referred to as A, B, and C. Each fabric has two I/O modules, for a total of six I/O modules slots in the chassis. The I/O modules are A1, A2, B1, B2, C1, and C2. Each I/O module can be an Ethernet physical switch, an Ethernet pass-through module, FC switch, 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 Dell I/O modules and is used for LAN. Fabric B contains Dell 8|4 Gbps SAN modules and is used for SAN. The Fabric C is unused.

PowerEdge M620 blade servers use a Broadcom 57810-k Dual port 10GbE KR bNDC (blade Network Daughter Card) to connect to the fabric A. Dell I/O modules uplink to Dell Networking S4810 network switches providing LAN connectivity. QLogic QME2572 8 Gbps Fibre Channel I/O mezzanine cards are used to connect to Dell 8|4 Gbps SAN modules. The uplinks of Dell 8|4 Gbps SAN modules connect to Brocade 6510 switches providing SAN connectivity.

Figure 4 below illustrates how the fabrics are populated in a Dell blade server chassis and how the I/O modules are utilized.

Figure 4: 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. The Active System 1000v solution takes advantage of NPAR. Partitions are created for various traffic types and bandwidth is allocated, as described in the following section.

## 6.2 Network Connectivity

This section describes the network architecture of Active System 1000v.

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

- **Connectivity between the Dell PowerEdge M620 blade servers and Dell I/O modules:** 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 Dell I/O Modules. The Dell I/O Modules 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 Dell I/O Modules. Internal ports 17-32 of each Dell I/O Module are disabled and not used.
- **Connectivity between the Dell I/O Module and Dell Networking S4810 switches:** The two Dell I/O modules are configured either to operate as a port aggregator for aggregating 16 internal ports to eight external ports or as a MXL switch.

The two fixed 40GbE QSFP+ ports on each Dell I/O Module are used for network connectivity to the two Dell Networking S4810 switches. These two 40GbE ports on each Dell I/O Module 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 Dell I/O Module, two links connect to one of the Dell Networking S4810 switches and the other two links connect to the other Dell Networking S4810 switch. Due to this design, each PowerEdge M1000e chassis with two Dell I/O modules will have total of 16 x 10Gb links to the two Dell Networking S4810 switches. This design ensures load balancing while maintaining redundancy.

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

**Connectivity between the two network switches:** The two S4810 switches are connected using Inter Switch Links (ISLs) using two 40 Gbps QSFP+ links. Virtual Link Trunking (VLT) is configured between the two 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.

## 6.3 Network Configuration

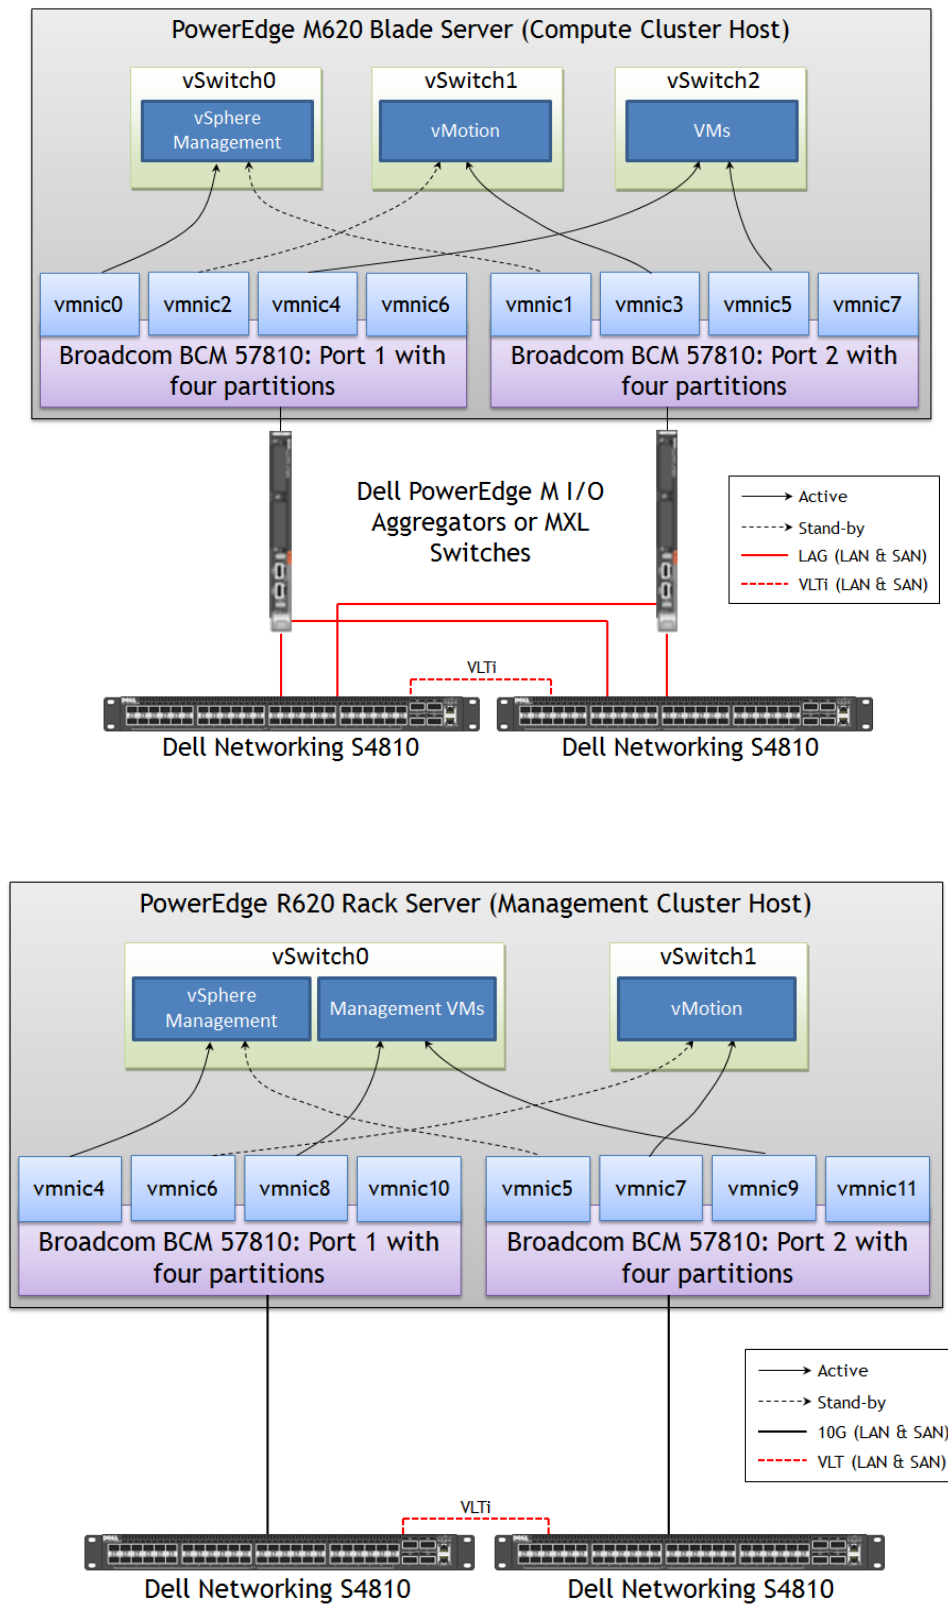
This section provides the network best practices for implementing VMware vSphere 5.1 Update 1 on Dell blade servers and Dell Networking S4810 switches.

**Hypervisor network configuration for LAN traffic:** VMware ESXi hypervisor is configured for the LAN traffic associated with the blade servers. LAN traffic in Active System 1000v 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, Brocade management, and Compellent 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 cluster.

On the compute cluster hosts (the PowerEdge M620 blade servers), one vSwitch each is created for VM traffic, vSphere management traffic and vMotion 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 and vMotion 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 5.

Figure 5: vSwitch, NPAR and I/O Module Configuration



**Traffic isolation using VLANs:** LAN traffic is separated into four unique VLANs; one VLAN each for management, vMotion, VM traffic, and out-of-band management. 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 Dell Networking S4810 switches. Additionally, the Dell Networking S4810 switch ports that connect to the blade servers are configured in VLAN trunk mode to pass traffic with different VLANs on a given physical port. The table 2 below provides an overview of different traffic types segregated by VLANs in the Active System 1000v, and which network devices with which they are associated.

Table 2: VLAN Overview

Traffic Type (VLAN segregation)	Description	Associated Network Device
Management	vSphere management traffic and Active System 1000v 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
Out-of-Band Management	Out-of-Band Management traffic	iDRAC, CMC, Brocade, and Compellent Management Ports

**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. With this option, 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 good load balancing across VMs, especially in the case of a large number of VMs.

**Virtual Link Trunking (VLT) for S4810s:** Inside each Active System 1000v, a Virtual Link Trunking interconnect (VLTi) is configured between the two Dell Networking S4810 switches using the Virtual Link Trunking (VLT) technology. VLT peer LAGs are configured between the Dell I/O modules and Dell Networking S4810 switches, and also between the Dell Networking S55 switch and the Dell Networking S4810 switches.

Virtual Link Trunking technology allows a server or bridge to uplink a single trunk into more than one Dell Networking 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.

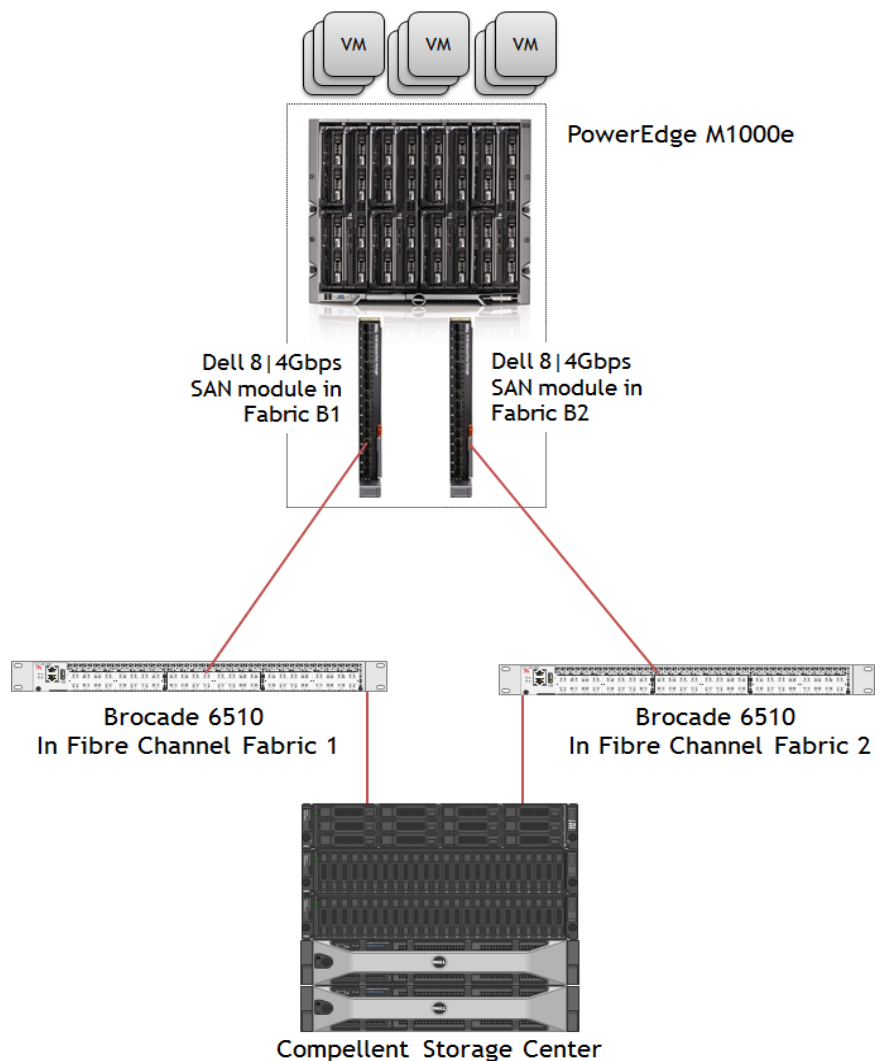
Note that the two switches can also be stacked together. However, this is not recommended, as this configuration may incur downtime during firmware updates of the switch or failure of stack links.

## 6.4 Storage Architecture

In this solution, Compellent Storage Center is connected to the Dell blade servers using Brocade 6510 FC switches.

**Fibre Channel Fabric Architecture:** The solution is configured with two FC fabrics as shown in Figure 6. The two fabric design ensures that changes to one fabric do not impact the other fabric.

Figure 6: Fibre Channel SAN Logical Connectivity



**Connectivity between the Dell FC SAN Module and Brocade 6510:** Each blade is populated with a QLogic QME2572 8 Gbps Fibre Channel I/O mezzanine card, which is used to connect to Fabric B. Fabric



B is populated with Dell 8|4Gbps SAN modules. The Dell FC SAN Module is configured to operate as a port aggregator for aggregating internal ports to external ports. Dell FC SAN module port aggregator operates in access gateway mode for providing N\_Port ID virtualization (NPIV) functionality. The following are the FC ports that the FC SAN Module uses:

- F\_Port - internal fabric port that connects a blade server (HBA)
- N\_Port - external node port that connects to a switch

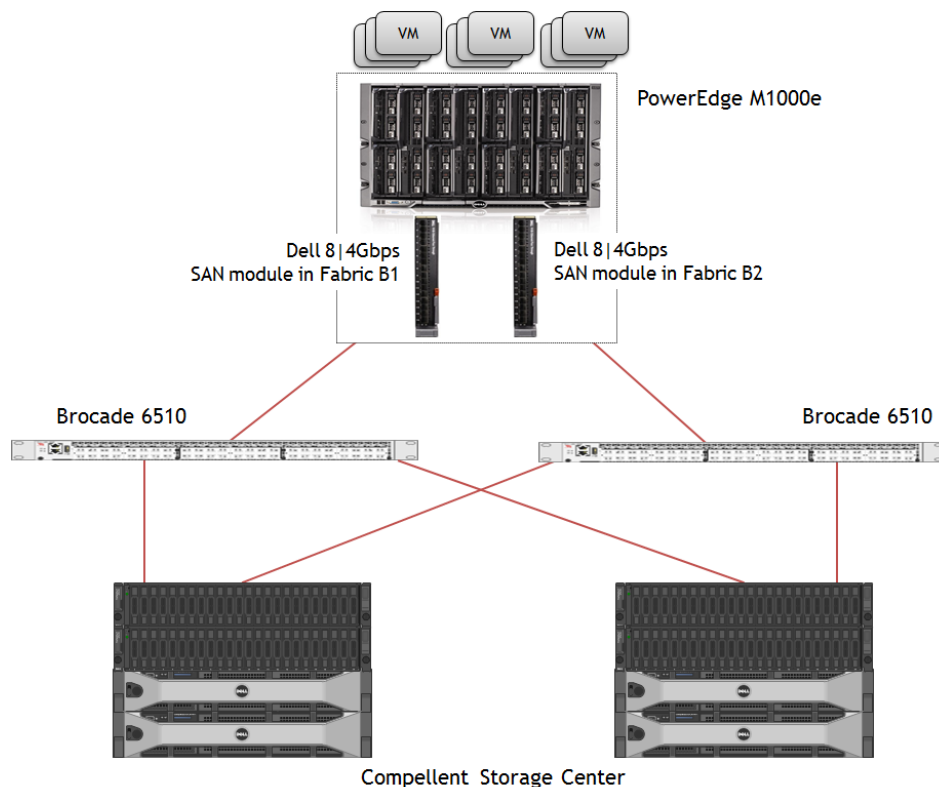
Brocade Access Gateway (AG) is a Fabric OS feature used to configure Enterprise fabric. Switches in AG mode are transparent to the host and the fabric. Thus number of hosts having access to the fabric can be increased without increasing the number of switch domains.

**Connectivity between Brocade 6510 and Compellent Storage Controller:** In the solution, each Compellent SC8000 storage controller is configured with two dual port FC HBAs. Two ports per HBA are used to connect each storage controller to the two Brocade 6510 switches.

**Compellent Storage Connectivity:** In this solution, Compellent SC8000 storage controllers are configured with two quad port 6Gb/s SAS HBAs. Using the two quad port SAS HBAs, multiple Compellent SC200/220 enclosures are connected. The enclosures are grouped into two redundant daisy chained connections for optimal performance. Each daisy chain loop can have a maximum of 168 drives.

**Additional Storage Controllers for Increased Storage Requirements:** The solution also supports additional pair of storage controllers for increased storage requirements. The additional controller provides increased IOPS which may be needed for some workloads like VDI. The additional storage controller is connected to the Brocade 6510 switches as shown in the Figure 7.

Figure 7: Fibre Channel SAN Logical Connectivity with Additional Storage Controllers



**Performance:** Dell Compellent SC8000, with the dual-controller configuration, 8 Gb Fibre Channel interconnects provides high bandwidth for data flows. This bandwidth is complemented with a large variety of drives in multiple speeds and sizes. The SC8000 controllers also use virtual port IQNs and WWNs, thereby enabling higher throughput and fault tolerance.

**Drive Types and Automated Tiered Storage:** In the Active System 1000v solution, the number of storage enclosures and the drives in the enclosures can be customized based on customer requirements. Administrators can mix SSD and SAS drives in the same system, as well as SAS drives with the same form factor (but different speeds and capacities) in the same storage enclosure. A maximum of 16 enclosures is supported in Active System 1000v.

High speed drives are assigned to higher tiers and low speed drives to lower tiers. Compellent Storage Center automatically configures RAID levels for these tiers and automatically moves the data between the tiers based on access patterns. Compellent Fluid Data storage dynamically moves data to the optimal tier based on actual usage. The most active blocks reside on high-performance SSD, or SAS drives, while infrequently accessed data migrates to lower-cost, high-capacity SAS or SATA drives. For more details, refer to the [Automated Tiered Storage](#) web page. Automated Tiered Storage requires Data Progress licenses for Compellent.

**RAID Array Design:** Dell Compellent SC8000 supports RAID 5, 6, and 10. The Compellent Storage Center will dynamically set up RAID based upon the demands of applications accessing data on the storage tier(s).

**Multipath Configuration:** In the solution, VMware Native Multipath Plug-In (NMP) is used to provide multi-pathing. Path Selection Plug-Ins (PSPs) run with the VMware NMP and is responsible for choosing a physical path for I/O requests. Round Robin (VMW\_PSP\_RR) path selection algorithm is the recommended configuration for Compellent Storage Center. Round Robin uses a path selection algorithm that rotates through all available active optimal paths enabling load balancing across the paths. This ensures all the paths are used to provide the maximum bandwidth and balance I/O across the paths/fabrics.

## 7 Management Infrastructure

Two PowerEdge R620 servers and one Dell Networking S55 1Gb Ethernet switch are used for management infrastructure. The PowerEdge R620 servers are connected to the Dell Networking S4810 switches using Broadcom 57810 Dual Port 10Gb Network Adapters. The servers are connected to the Compellent storage through the Brocade 6510 switches using a QLogic QLA2562 8Gbps Fibre Channel Card.

Note that the Compellent storage is shared between management cluster and compute cluster. The Compellent storage must be sized so that sufficient bandwidth is allocated for both the management VMs and compute VMs.

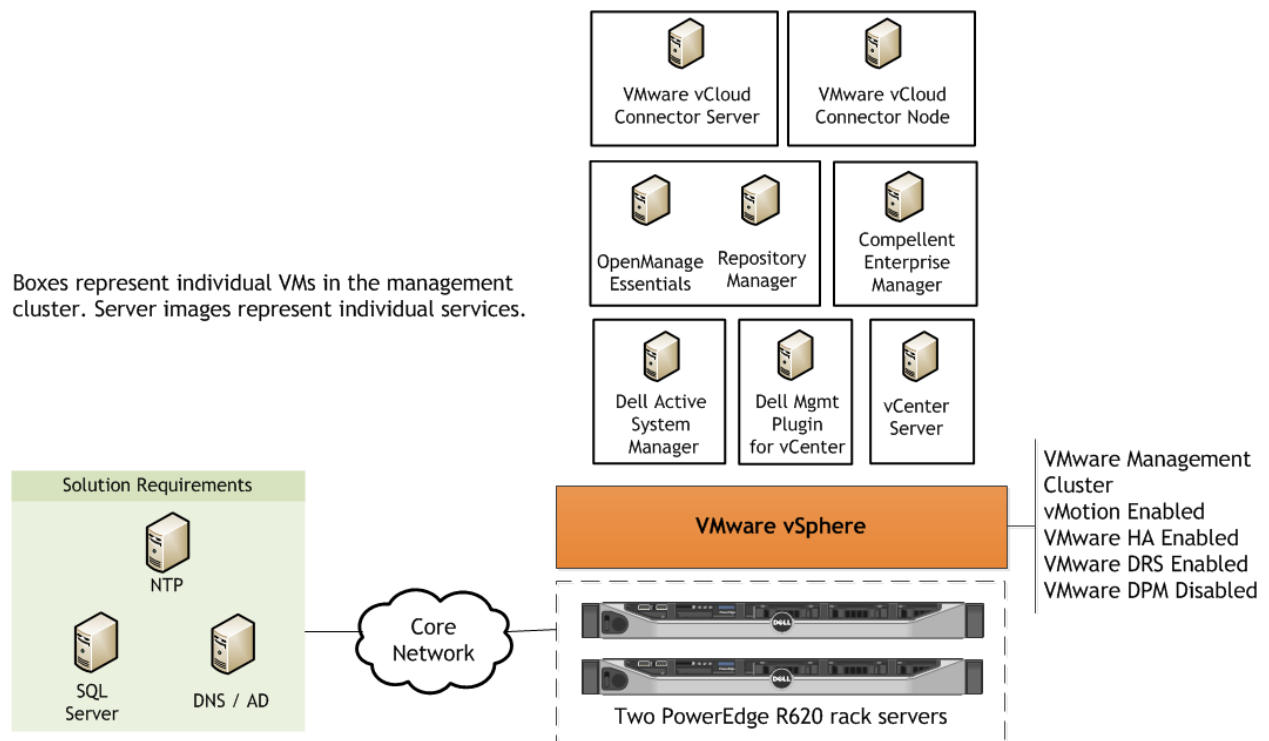
The PowerEdge R620 servers run VMware ESXi 5.1 Update 1 hypervisor and are a part of the unique vSphere Cluster. VMware High Availability is enabled in that cluster to provide HA for virtual machines. Admission control is disabled in the VMware HA Cluster. If admission control is enabled, VMware HA would 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 following management components are included in Active System 1000v solution as illustrated in Figure 8:

- Dell Active System Manager
- VMware vCenter Server
- Dell Management Plugin for vCenter
- Dell Compellent Enterprise Manager
- Dell OpenManage Essentials
- Dell Repository Manager
- VMware vCloud Connector Server
- VMware vCloud Connector Node

Compellent Plugin for vCenter is installed along with VMware vCenter client in VMware vCenter Server VM.

Figure 8: Management Components



The remainder of this section will provide an introduction to each component and how they are integrated into the Active System 1000v solution.

### 7.1 Dell Active System Manager

As described in section 3, the Dell Active System Manager is the Active Infrastructure management software that is part of the Active System 1000v solution. The Dell Active System Manager virtual

appliance is deployed on the management cluster. For fullest functionality, direct internet access, or access through a proxy, is recommended.

Active System Manager addresses key factors that impact service levels, namely infrastructure configuration errors, incorrect problem troubleshooting, and slow recovery from failures. Active System Manager dramatically improves the accuracy of infrastructure configuration by reducing manual touch points.

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

### 7.2 Dell OpenManage Essentials (OME)

In the Active System 1000v, Dell OpenManage Essentials (OME) is sized and configured to monitor the Active System 1000v solution components. It is deployed on a Windows Server 2012 virtual machine within the management cluster. High availability of the OME virtual machine is provided by VMware High Availability service. OME utilized a local SQL Express database. For fullest functionality, direct internet access, or through a proxy, is recommended.

Within the Active System 1000v, OME is utilized for discovery, inventory, and hardware level monitoring of blade and rack servers, blade chassis, Dell I/O modules, and Dell Networking 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 through WS-MAN and SNMP inventory calls. For instance, reporting is available to provide blade and rack server firmware versions or solution warranty status. OME can be used as the single point of monitoring for all hardware components within an enterprise.

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

### 7.3 Dell Repository Manager (DRM)

Within the Active System 1000v solution, Dell Repository Manager (DRM) is installed on the same Windows Server 2012 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>.

### 7.4 Dell Management Plug-in for VMware vCenter (DMPVV)

Dell Management Plug-in for VMware vCenter is deployed as a virtual appliance within the management cluster, and is attached to the VMware vCenter Server within the Active System 1000v stack. DMPVV communicates with the VMware vCenter Server, the hypervisor management interfaces, and server out-of-band management interfaces (iDRAC). For ease of appliance firmware updates and warranty

information, it is recommended that the DMPVV appliance has access to an internet connection either directly, or through a proxy. Dell Management Plug-in for VMware vCenter enables customers to:

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

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

## 7.5 Dell Compellent Enterprise Manager

Compellent Enterprise Manager simplifies storage management by providing a single, centralized console for the administration of multiple local and remote Compellent systems. Users can configure and verify remote replication processes, monitor storage capacity and disk utilization in real time, and generate comprehensive enterprise storage usage and performance reports.

## 7.6 VMware vCloud Connector

VMware vCloud Connector is an optional component of the Active System 1000v 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 have responsibility for the physical transfer of VM workloads. Within the Active System 1000v, two of these components, the server and the local node, are installed. The third component, 'remote' node VM, should be installed outside of the Active System 1000v 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 VMs to be transferred and the number of concurrent transfers anticipated.


# 8 Connecting Active System 1000 to Datacenter Network

It's likely that an Active System 1000 is connected into a data center infrastructure that consists of Dell Networking switches, Cisco switches, or those of some other vendor. Active System uses Dell Networking S4810 as the ToR switch and Dell Networking S55 for the OOB switch. In this section, we provide examples to show how Active System is connected to the Datacenter network.

## 8.1 Connecting the S55 OOB switch to Datacenter Network

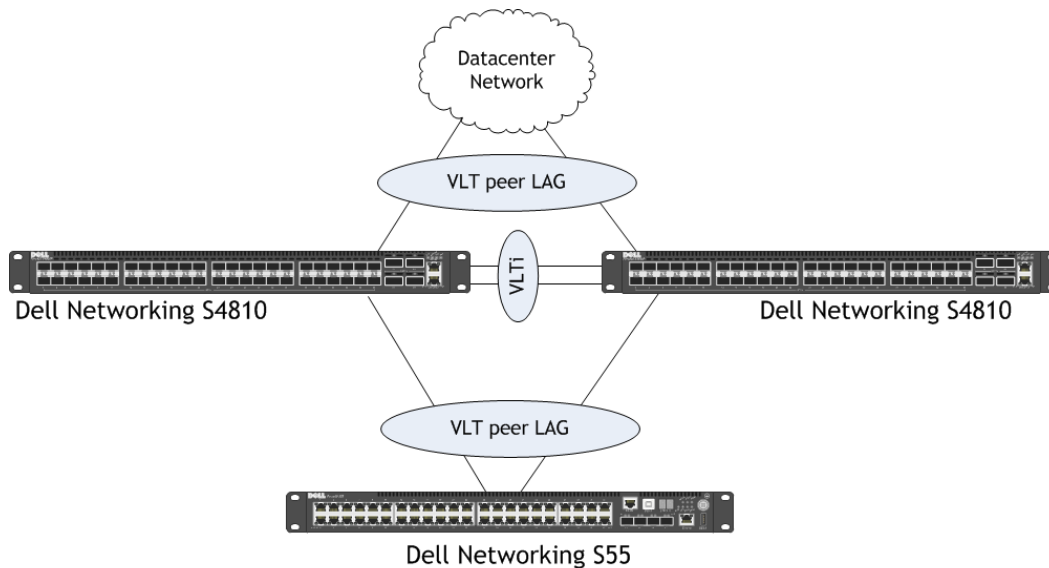
A Dell Networking S55 switch needs to be connected to the datacenter out of band network, if available, using two 1Gb uplinks from the S55 switch. This is shown in Figure 9.

Figure 9: S55 Connectivity to Datacenter OOB Network (Default Option)



If the datacenter OOB network is not available, the S55 switch can alternatively be connected to the S4810 switches for OOB connectivity to the datacenter, as show in Figure 10.


Figure 10: Alternative OOB Connectivity - S55 Switch to S4810 Switches



## 8.2 Connecting to a Dell Networking Datacenter Network

Figure 11 below demonstrates an example with Dell Networking Z9000 switches. The two Z9000 switches can be connected together using VLTi. The S4810 switches in the Active System use a 4-port VLT Peer LAG connecting to two Z9000 switches. The number of VLT Peer LAG links is flexible and can be changed according to use cases. The VLT Peer LAG will be used for VLAN traffic from the VLT uplinks, so they should be planned in an appropriate manner to avoid oversubscription.

Figure 11: Active System 1000v connectivity to Dell Networking Z9000 Switch

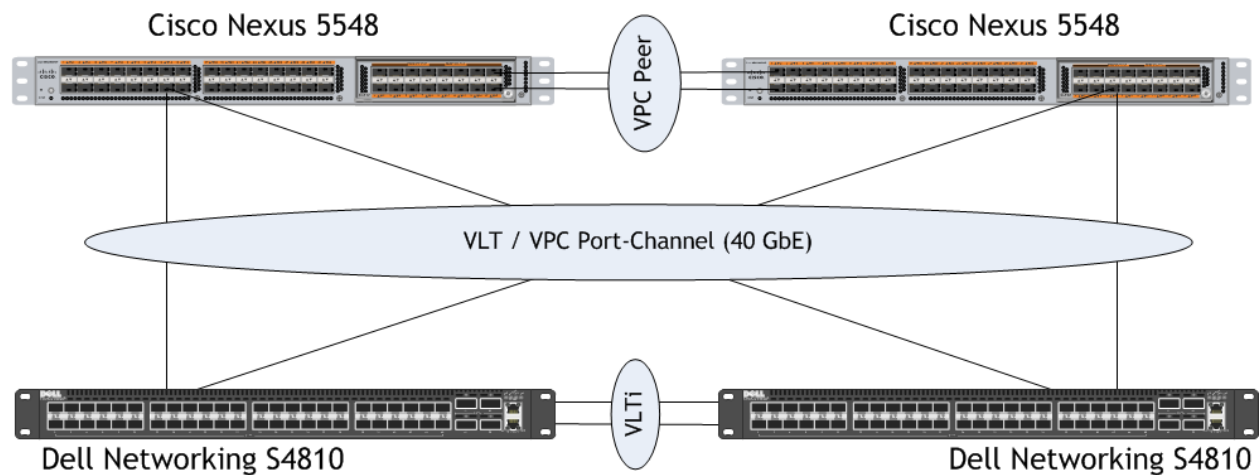


### 8.3 Connecting to a Cisco Nexus Datacenter Network

Figure 12 below demonstrates an example with Cisco Nexus 5548 switches. The Dell Networking S4810 switches in the Active System 1000 have a 4-port LAG/Port Channel linking into the two Cisco Nexus switches. The number of links is flexible and can be changed according to use cases. Also in this example, four ports were used for the vPC “Peer Link”. The vPC Peer Link will be sharing VLAN traffic from the vPC uplinks, so they should be planned in an appropriate manner to avoid oversubscription. The “Peer Keep-Alive” link is suggested to be configured to utilize the management port (this is the default) and will send little layer 3 traffic between the switches.

The two Cisco Nexus 5000 Series switches are configured with vPC and a high availability aggregation pair.

Figure 12: Active System 1000v connectivity to Cisco Nexus 5548





## 9 Scalability

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

**Scaling Compute and Network Resources:** This solution is configured with two Dell Networking S4810 switches. Up to two PowerEdge M1000e chassis can be added to the two Dell Networking switches. In order to scale the compute nodes beyond two chassis, new Dell Networking S4810 switches need to be added.

**Scaling Storage Resources:** Compellent storage can be scaled seamlessly and independent of the compute and network architectures. Additional drives and enclosures can be added to the existing controllers. New volumes can be created or existing volumes can be expanded to utilize the capacity in the added enclosures. The Active System 1000v solution can scale up to maximum of 16 array enclosures. To scale beyond this, additional racks or controllers can be added. Compellent SC8000 controller can scale up to a maximum of 960 drives.

## 10 Delivery Model

This Reference Architecture can be purchased as a complete solution, the Active System 1000v. 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 13 below shows the Active System 1000v solution with a single chassis and up to 16 compute hosts in a one rack configuration that also contains supporting Compellent storage. Figure 14 shows Active System 1000v with two chassis with up to 32 compute hosts and maximum of 16 storage enclosures in a two rack configuration. Figure 14 shows the configuration with one chassis and 16 compute hosts and four storage controllers available with the Active System 1000v solution. Note that switches shown in figures are shown mounted forward for representation. In actual use, ports face the back of the rack. PDUs shown are for illustration and will vary by region or customer power requirements. Additional PDUs are utilized within the rack.



Figure 13: Active System 1000v Single Chassis: Rack Overview

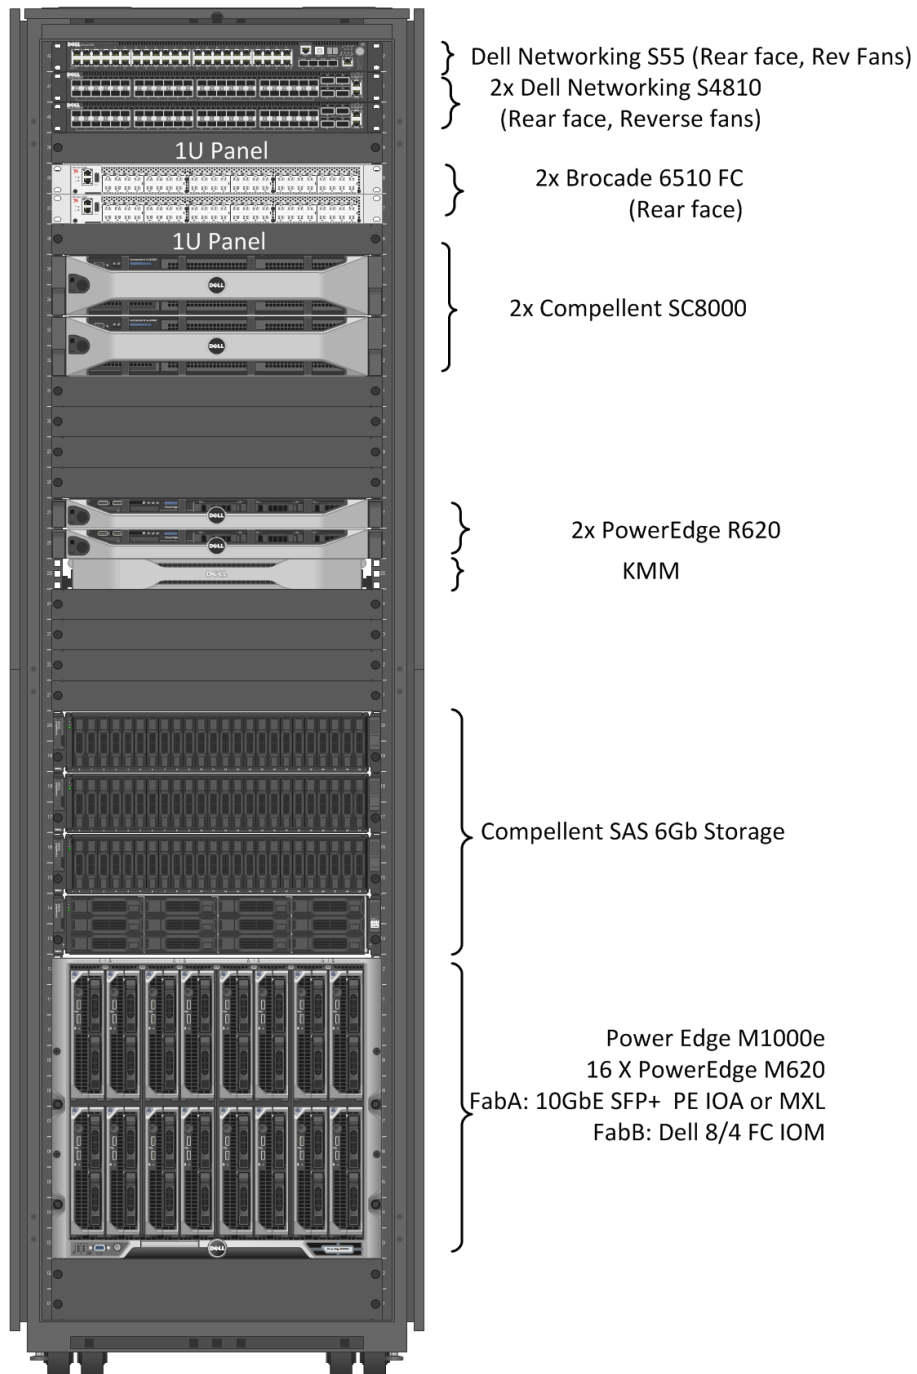


Figure 14: Active System 1000v Two Chassis and Maximum Storage: Rack Overview





Figure 15: Active System 1000v Configuration for additional storage requirements: Rack Overview



## 11 Additional Supported Configurations

Dell Active System 1000 supports two additional network configurations. The customer can choose to have a Dell Networking S5000 (1U 10/40GbE LAN/SAN switch equipped with native FC and FCoE capabilities) switch as a LAN Top of Rack switch or a Cisco Nexus 5548 (1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE) switch as LAN Top of Rack switch. These switches are not pre-racked in Active System 1000 solution.

### 11.1 Dell Networking S5000 as LAN Top of Rack Switch


The Dell Networking S5000 is a first-of-its-kind modular, 1 rack unit (RU) 10/40GbE top-of-rack (ToR) LAN/SAN switch equipped with native FC and FCoE capabilities. The S5000 switch's innovative system design is powered by an industry-hardened, and feature-rich operating system for maximum dependability and uptime. Dell Open Automation framework provides integrated automation, scripting, and programmable management for enhanced network flexibility in virtualized environments.

Key features include:

- **Pay-as-you grow modularity, designed to scale** - provides for greater deployment flexibility and IT budget allocation compared to fixed-port switches. The S5000 switch accommodates four modules allowing customers to populate a single module and add as necessary instead of buying all four modules at once.
- **High-density LAN/SAN convergence** - saves on the number of switches and rack space required, the S5000 has up to 1.3 to 2.6 times the port density per rack unit compared to industry alternatives. The S5000 has a maximum of 64 x 10GbE ports, or 48 x Ethernet/FC ports with 16 x 10GbE ports.
- **Feature-rich storage networking** - complete support for iSCSI, RDMA over Converged Ethernet (RoCE), Network Attached Storage (NAS), FCoE, and FC fabric services, all on the same platform.
- **Future-proof design for maximum investment protection** - with the modularity and system design, the S5000 switch hardware is future-proofed to support newer features and options when released without needing to sacrifice existing infrastructure investment.
- Easy integration, proven interoperability with leading adapter, switch, and storage vendors including Broadcom, Brocade, Emulex, Intel, and Qlogic.

The connectivity between the Dell Networking S5000 switch and the blade IO modules is similar to the connectivity between the S4810 switches and the blade IO modules. The blade IO Modules connect to the S5000 ToR switch through the 40Gb ports on the IO Modules to the 40Gb ports on the S5000 switch. The two Dell Networking S5000 switches are configured with Virtual Line Trunking (VLTi) using two 40 Gbps QSFP+ links. VLT Interconnects are created between the two 40 Gbps QSFP+ ports, providing a path for communication across the switches. Figure 16 below shows the connectivity between S5000 switches and blade I/O modules.

Figure 16: Active System 1000v using S5000 as ToR Switch



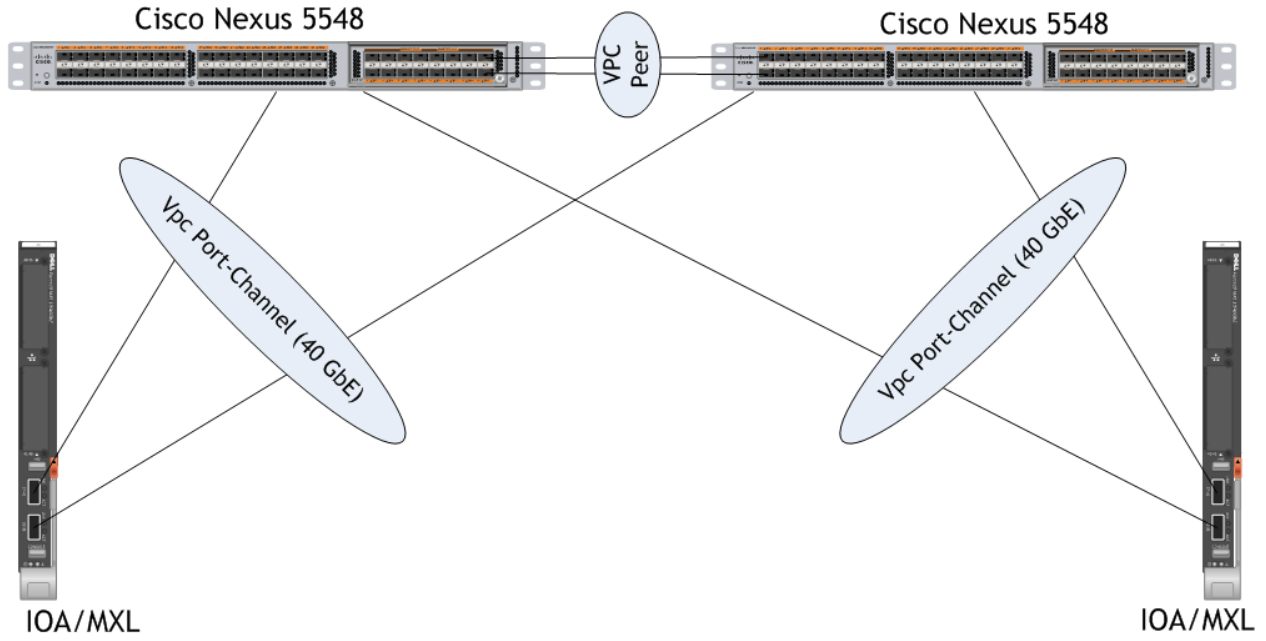
## 11.2 Cisco Nexus 5548 as LAN Top of Rack Switch

The Cisco Nexus 5548 is a 1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 fixed ports and one expansion slot. The Cisco Nexus 55000 Series switch enables consistent low-latency Ethernet solutions, with front-to-back or back-to-front cooling, and with data ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data center-class Cisco® NX-OS Software for high reliability and ease of management. The Cisco Nexus 5500 platform is well suited for enterprise-class data center server access-layer deployments across a diverse set of physical, virtual, storage-access, and high-performance computing (HPC) data center environments.

For more information on Cisco Nexus switches, see <http://www.cisco.com/en/US/products/ps9670/index.html>.

The connectivity between Cisco Nexus 5548 switches and the blade IO modules is also similar to the connectivity between the S4810 switches and the blade IO modules. The only difference is that the blade IO modules connect to the Cisco Nexus 5548 virtual PortChannel (vPC) using QSFP+ optical transceivers and QSFP+ optical breakout cables. The 40Gb ports on the IO modules have a Dell 40Gb QSFP+ optical transceiver. The 40Gb QSFP+ optical breakout cables connect to the Cisco Nexus 5548 switches. The two Cisco Nexus 5548 switches are configured as vPC peers. The Dell IO modules connecting to the Cisco Nexus 5548 are vPC port-channels with a throughput of 40Gbps. Figure 17 below shows the connectivity between Cisco Nexus 5548 switches and blade I/O modules.

Figure 17: Active System 1000v using Cisco Nexus 5548 as ToR Switch



## 12 Reference

- [Dell Active System Manager](#)
- [Dell Active Infrastructure Wiki](#)
- [Dell Networking Switch Details](#)
- [Dell PowerEdge M1000e Technical Guide](#)
- [Dell PowerEdge M I/O Aggregator Configuration Quick Reference](#)
- [Brocade 6510 Product Details](#)
- [Dell 8/4 Gbps Fibre Channel SAN module](#)
- [Dell Compellent Storage](#)
  - [Automated Tiered Storage](#)
  - [Fast Track](#)
  - [Dell Compellent Plug-In for VMware vCenter](#)
- VMware vSphere links:
  - [VMware vSphere Edition Comparisons](#)
  - [VMware vSphere Compatibility Matrixes](#)
  - [VMware High Availability \(HA\): Deployment Best Practices](#)
  - [VMware Virtual Networking Concepts](#)
- [Dell Management Plug-In for VMware vCenter references - Solution Brief](#)
- NPAR: [Enhancing scalability through network interface card partitioning](#)