

Reference Architecture for Active System 50 with Microsoft Hyper-V

Release 1.1 for Dell 12th Generation PowerEdge Rack Servers, Dell Networking Switches, Dell EqualLogic iSCSI SAN, Microsoft Windows Server 2012 with Hyper-V role, and Dell Active System Manager

Dell Global Solutions Engineering

Revision: A00



THIS WHITE PAPER 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. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, *PowerEdge*, *EqualLogic*, and *OpenManage* are trademarks of Dell Inc. *Microsoft*, *Active Directory*, *Windows*, *Hyper-V* and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation. 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 Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Jun 2013 | Rev A00

Table of Contents

1	Introduction	1
2	Audience	1
3	Overview	2
4	Design Principles.....	8
5	Prerequisites and Datacenter Planning	9
6	Architecture	10
6.1	Server / Network Connectivity.....	11
6.2	Server / Storage Connectivity.....	13
6.3	Server / HA and Redundancy.....	13
6.4	Storage Architecture	13
7	Management Infrastructure	16
7.1	Active System Manager.....	16
7.2	Windows Server 2012 Built-In tools	17
7.3	Dell OpenManage Essentials	18
7.4	Out-of-band Management: iDRAC.....	19
7.5	Dell EqualLogic SAN HQ	20
7.6	System Center 2012 SP1	20
7.7	System Center 2012 Integration	20
7.8	Out-Of-Band Management Connectivity	21
8	Connecting Active System 50 to Datacenter Core Network.....	21
8.1	Connecting to Datacenter Network with Dell Networking Switches	22
8.2	Connecting to a Cisco Datacenter Network.....	23
9	Scalability.....	23
10	Delivery Model.....	24
11	References	25
11.1	Dell PowerEdge Server Documentation and Hardware/Software Updates.....	25
11.2	Dell EqualLogic	25
11.3	Dell Virtualization Documentation	25
11.4	Microsoft® Management Software.....	26

Tables

Table 1: Solution Components	3
Table 2. Active System 50 LAN and SAN Switch Overview	5
Table 3: Default Configurations	7
Table 4: Traffic Description	12
Table 5: Sample VLAN and subnet configuration.....	13

Figures

Figure 1: Active System 50 Component Overview	2
Figure 2: Network Topology (Logical View)	10
Figure 3: Network Configuration for PowerEdge Rack Servers.....	12
Figure 4: Manage Active System 50m with Server Manager	17
Figure 5: Resources Measuring in PowerShell 3.0	18
Figure 6: OME Dashboard	19
Figure 7: Connectivity of OOB management components.	21
Figure 8: Connectivity to Dell Networking Datacenter Network.....	22
Figure 9: Connectivity to a Cisco Network	23
Figure 10: Active System 50m Rack and Component Overview	24

Revision History

Revision	Description	Date
A00	Initial Version with Active Infrastructure 1.1 and Active System Manager 7.1	July 2013

1 Introduction

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

The Active System 50, a member of Dell Active Infrastructure family, is a converged infrastructure solution that 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.

The Active System 50 includes Dell™ PowerEdge™ R620 servers running Microsoft® Windows Server® 2012 Datacenter Edition with Hyper-V Role enabled, Dell EqualLogic™ PS6100 Series iSCSI storage and Dell Networking 7024 switches.

An optional Dell PowerEdge R420 server is available to host Dell and customer management tools. Additionally, this management server can be configured with the Dell Active System Manager to build private cloud solutions that address key needs for small to medium businesses and datacenters. The Active System 50 configuration varies in the number of EqualLogic PS6100 storage arrays to meet resource needs.

Active System 50 is offered in configurations with either VMware® vSphere™ or Microsoft® Windows Server® 2012 with Hyper-V® role enabled Hypervisors. The VMware vSphere solution is the Active System 50v and the Microsoft Hyper-V solution is the Active System 50m. This document defines the reference architecture for Active System 50m.

2 Audience

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

IT administrators and managers can use this document to understand the solution architecture.

3 Overview

The Active System discussed in this document is powered by Dell PowerEdge servers, Dell EqualLogic iSCSI storage, Dell Networking switches, and Microsoft Windows Server 2012 with Hyper-V role enabled. The solution implements Dell and Microsoft best practices. EqualLogic SAN HeadQuarters (SAN HQ) and Group Manager are included in the solution for storage array monitoring and management. The solution also includes the rack, four power distribution units (PDU), two uninterruptible power supplies (UPS), KMM (Keyboard, Monitor, Mouse), and management server.

Active System 50 includes two PowerEdge R620 servers and one EqualLogic PS6100 array. Storage expansion modules and other optional equipment are also offered in this release. Figure 1 below illustrates key features of major components in the solution.

Active System 50 includes two PowerEdge R620 servers and one EqualLogic PS6100 array. Storage expansion modules and other optional equipment are also offered in this release. Additional details for these are provided in the section below.

Figure 1: Active System 50 Component Overview




Table 1 below describes key solution components and the roles served in Active System 50m.

Table 1: Solution Components

Component	Description	Role
Compute Cluster	(2) Dell PowerEdge R620 servers running Windows Server 2012 Datacenter Edition and Hyper-V role enabled	Host virtual machines (VMs)
Storage	One to two Dell EqualLogic PS6100 Series controllers with 24 bay 2.5" SAS enclosures or 24 Bay 3.5" SAS enclosures	Provide iSCSI based shared storage for the Hyper-V cluster to host the VMs
Management	Optional (1) Dell PowerEdge R420 as a dedicated Management server	Host Dell Active System Manager, Dell EqualLogic SAN HQ, Dell OpenManage Essentials, System Center 2012 Virtual Machine Manager SP1 and other optional System Center 2012 components
LAN Switches	(2) Dell Networking 7024 switches	Support LAN traffic consisting of VM, Live Migration, Management, and Cluster traffic.
SAN Switches	(2) Dell Networking 7024 switches	Support iSCSI SAN traffic.

Microsoft Windows Server 2012: Microsoft® Windows Server® 2012 is Microsoft's flagship server operating system which provides the Hyper-V® virtualization platform. Hyper-V provides a virtualization platform that can consolidate Windows® and Linux® workloads enabling IT managers the ability to more fully utilize their available hardware resources.

Dell Active System Manager: Active System Manager sits at the center of Active System, and simplifies infrastructure configuration, collapses management tools, and drives automation and consistency. Through capabilities such as template-based provisioning, automated configuration, and infrastructure lifecycle management, Active System Manager enables IT to respond rapidly to business needs, maximize datacenter efficiency, and strengthen quality of IT service delivery.

The software enables a cloud like environment and supports pre-built provisioning templates and custom orchestrated workflows for a wide range of physical or virtual environments. Use and features of the Dell Active System Manager is described in more detail in subsequent sections of this document. Additional detail can also be found on the Dell website at: [Dell Active System Manager](#).

Dell OpenManage Essentials: The Dell OpenManage™ Essentials Console provides a single, easy-to-use, one-to-many interface through which to manage resources in multivendor operating system and hypervisor environments. It automates basic repetitive hardware management tasks – such as discovery, monitoring, and updates – for Dell servers, storage, and network systems. OME employs the embedded management of Dell 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. For more information on Dell Enterprise Management and OpenManage Essentials, see <http://content.dell.com/us/en/enterprise/data-center-system-management>.

Dell EqualLogic SAN Headquarters (SAN HQ): SAN HQ provides in-depth reporting and analysis tools. With the information you get from SAN HQ, you can better customize, protect, and optimize your storage. Dell EqualLogic SAN HQ provides consolidated performance and robust event monitoring across multiple Dell EqualLogic groups. If action is required, the Dell EqualLogic Group Manager can be launched directly from SAN HQ giving you the ability to quickly adjust storage systems to meet business needs. For more information about Dell EqualLogic storage and SAN HQ, see <http://www.dell.com/us/enterprise/p/equallogic-products>.

Optional Microsoft System Center 2012 SP1: System Center 2012 is Microsoft's systems management platform which helps to monitor, manage, deploy, backup, and more. System Center 2012 SP1 introduces support for Windows Server 2012, enabling users to take advantage of Windows Server 2012's features in a managed environment. This solution utilizes either a customer provided or a trial version of System Center 2012 SP1 and focuses on Virtual Machine Manager 2012. Further information on all Microsoft System Center components can be found at www.microsoft.com/systemcenter.

System Center 2012 Virtual Machine Manager SP1 (VMM): VMM is Microsoft's centralized virtualization manager. It provides the fundamental services for creating and managing virtual machines, optimizing resources, and rapid deployment of virtualized infrastructure.

VMM 2012 SP1 supports the features included in Windows Server 2012: network virtualization, VHDX support, SMB 3.0 file share, and the enhanced storage management, host provisioning, and Live Migration capabilities.

Dell PowerEdge Rack Servers: The Dell PowerEdge R620 server uses two 8-core Intel® Xeon® E5-2660 series 2.2GHz processors, 64GB of memory, and Intel chipset architecture in a 1U rack mount form factor. The PowerEdge R620 is a powerful general purpose platform with highly expandable memory and I/O capabilities to match. Its extensive scalability and balanced design allows it to handle very demanding workloads. The server features two CPU sockets and 24 memory slots supporting 2, 4, 8, 16, or 32GB DIMMs to meet the memory demands of a virtualized infrastructure.

The PowerEdge R620 has rack Network Daughter Card (rNDC) 1Gb Ethernet ports and an additional quad port 1Gb Ethernet card has been added to each, providing a total of eight 1Gb ports. Four of these ports are utilized for LAN traffic and the remaining four are for SAN traffic. In addition, each PowerEdge R620 server is configured with a PERC H710 RAID controller that hosts Windows Server 2012 Datacenter.

Energy-efficient design features include power-supply units sized appropriately for system requirements, innovative system-level design efficiency, policy-driven power and thermal management, and highly efficient standards-based Energy Smart components. For more information, see the Dell PowerEdge R620 specifications at Dell.com/PowerEdge.

Dell PowerEdge R420 Management Server: The optional management server uses Intel® Xeon® E5-2400 series 2.2GHz processors in a 1U rack mount form factor. The server features two CPU sockets and 12 memory DIMM slots supporting 2, 4, 8, 16 or 32GB DIMMs. The PowerEdge R420 server has two onboard 1Gb ports and an add-in dual port 1Gb NIC which provide two ports for LAN traffic and two ports for SAN traffic.

In the Active System 50 Solution, the PowerEdge R420 Management Server is optional. This choice provides customers with the flexibility to manage the Active System 50 from a VM running on the Active System 50 cluster, from an existing physical server that runs virtualization management software and meets the vStart solution requirements, or from a VM that meets the requirements and runs in a separate virtualized environment that can communicate with the Active System 50 solution. More details about the Active System 50 Solution requirements are provided below

For more information, see the Dell PowerEdge R420 guides at Dell.com/PowerEdge.

Embedded Management with Dell’s Lifecycle Controller: The Lifecycle Controller is the engine for advanced embedded management and is delivered as part of iDRAC7 Enterprise in Dell PowerEdge 12th generation servers. For more information on Dell Lifecycle Controllers, see <http://content.dell.com/us/en/enterprise/dcsm-embedded-management>.

Dell Networking 7024 Switches: At the heart of the solutions network configuration are four Dell Networking 7024 switches. These managed Layer 3 Gigabit Ethernet switches offer the enterprise-class level of performance required for this configuration. The LAN switches use a link aggregation (LAG) configuration that enables connection redundancy and added bandwidth where required. The Active System 50m dedicates two switches for LAN traffic and two for SAN traffic. The switches are segregated to minimize latency for iSCSI traffic. In addition, this design decision allows for integration into environments that may have already implemented separate networks for LAN and SAN traffic. If the existing environment has a unified fabric (LAN and SAN on a single fabric), then the LAN and SAN switches provided can be uplinked into this unified environment. Each switch supports 24 x 1Gb connections and has two expansion bays that support either 10Gb Ethernet modules and/or stacking modules. VLANs are implemented to support solution management, security, and network traffic segmentation, and routing is leveraged to provide flexible connectivity.

Table 2 below provides the switch capabilities and hardware configuration details. For detailed information, refer to the 7024 user manuals.

Table 2. Active System 50 LAN and SAN Switch Overview

PC7024	Components
Switch Capability	24 10/100/1000 BASE-T auto-sensing Gigabit Ethernet Switching ports The 10Gbps Stacking Modules are used in Ethernet mode and configured as an Inter Switch Link (ISL) Link Aggregation Group (LAG). One 10 Gigabit SFP+ Ethernet module per LAN switch is used to uplink into a core network (if needed)
Performance	Switch Fabric Capacity up to 224 Gb/s per switch; forwarding rate up to 160 Mbps

Dell EqualLogic PS6100 Series Storage: EqualLogic PS6100 1GbE arrays include Dual quad-port 1GbE controllers, high-performance 300GB, 10K, 2.5” SAS HDDs Cache-to-flash memory de-staging and up to 7.2TB of capacity in a 2U chassis. Each storage controller has four 1Gb Ethernet network interfaces, and a 100Mb interface dedicated to management. Management features include:

- Group Manager, which integrates all management into one console for both file and block-storage protocols
- Dell EqualLogic Host Software that extends the functionality of the array-based software to enable cooperation with the host
- Host Integration Tools for Microsoft®, VMware® and Linux® products, as well as Dell EqualLogic SAN Headquarters (SAN HQ) management software

Contact your Dell sales representative for more information on Dell EqualLogic PS6100 Series storage configurations and sizing guidelines. For additional information, see <http://www.dell.com/us/enterprise/p/equallogic-ps6110-series?~ck=anav>.

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

- **The Dell Multi-UPS Management Console (MUMC):** Designed for multi-host server environments.
- **The Dell UPS Local Node Manager (ULNM):** Created for single-host environments, Dell’s node manager provides graceful automatic shutdown of any UPS, load segment, or connected device during prolonged power disruptions.
- **Free multi-platform software:** Manage your UPS systems with software that works intuitively across most popular virtualization platforms.
- **Large on-unit display:** Change control settings easily with a large, backlit LCD control panel offering the same color-coded system as Dell servers.
- **Virtual Environments:** Integrates with VMware, Microsoft, and Citrix. Manage virtual machine start/stop profiles through plug-ins like VMware vCenter. Use advanced virtualization functions like vMotion™ and LiveMigration to move VMs.

Power Distribution Unit (PDU): As the name suggests, PDUs distribute power from the main power source to the individual components within the 24U rack. Dell PDUs utilize a combination of worldwide standard IEC power outlet connections with regionalized input options allowing flexibility across a variety of global power infrastructures. The Active System 50 uses four 16A, 120-240V, half-height PDUs. Consult with your Dell Sales and Services representatives about your local requirements.

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

Keyboard, Monitor, Mouse (KMM): A 1U KMM console (touchpad, keyboard, and LCD) is cabled to the management server providing the ability to walk up to the rack and manage the entire solution. The KMM is optional for the Active System 50 solution.

PowerEdge 2420 Rack Enclosure: A single 24U rack is required to support either configuration. Blanking panels are included to ensure optimal airflow.

The solution is available in three default configurations, allowing users to better tailor the solution to their environment. Table 3 below describes these default configurations.

Table 3: Default Configurations

Component	Active System 50 with Active System Manager	Active System 50 with Management Server	Active System 50
Host Server	(2) Dell PowerEdge R620	(2) Dell PowerEdge R620	(2) Dell PowerEdge R620
Management Host(s)	(1) Dell PowerEdge R420 with Windows Server 2012 Datacenter Edition, Hyper-V Role Enabled	(1) Dell PowerEdge R420 with Windows Server 2012 Datacenter Edition, Hyper-V Role Enabled	None
Default Storage Device	(1) Dell EqualLogic PS6100X, each array with 24 10K 300GB 2.5" SAS drives	(1) Dell EqualLogic PS6100X, each array with 24 10K 300GB 2.5" SAS drives	(1) Dell EqualLogic PS6100X, each array with 24 10K 300GB 2.5" SAS drives
Management Components Licenses	Dell Active Systems Manager with ten license bundle	None	None

4 Design Principles

1. **Redundancy to minimize single point of failure:** The system is designed to mitigate failure points. NIC teaming and Multipath I/O (MPIO) are used to provide failover across the redundant network interfaces. iSCSI storage redundancy is achieved with multiple ports and storage controllers. For network traffic, NIC ports are teamed in such a way to distribute traffic across separate ports. Active System 50 also includes redundant power supplies connected to separate Power Distribution Units (PDUs). The solution is further optimized to reduce components such as NICs, cables, and IO Modules, then utilizes hypervisor-based high-availability to provide virtual machine failover. Out-of-Band (OOB) Management is not architected with this level of redundancy since mission critical workloads will continue to operate in the event of an OOB management failure.
2. **Management:** Integrated management using Active System Manager, Dell OpenManage Essentials, EqualLogic SAN HQ, and optional Microsoft System Center 2012 SP1 VMM.
3. **Dedicated iSCSI Network:** Dedicated NICs and switches are provided for iSCSI storage traffic to isolate the storage traffic from LAN traffic. This ensures minimal latency for storage I/O operations.
4. **Flexible configurations:** Dell Active System 50m is pre-configured to suit most customer needs for a virtualized infrastructure. Each rack server is configured with sufficient CPU, memory, and network adapters. Where desired, Active System 50m also supports processor, memory, and storage options. Additional options, such as wide or deep racks are also available.
5. **Racked, Cabled, and Ready to be deployed:** This Dell Active System is available to be racked, cabled, and delivered to the customer site ready for deployment. Components are configured and racked to optimize airflow and thermals.

5 Prerequisites and Datacenter Planning

Power, Cooling, and Weight Considerations: Dell Active System 50m solution is configured with four 16A, 120-240V, half-height Power Distribution Units (PDUs) to meet the power requirements of the components, as well as regional constraints. Power consumed, cooling required, and information regarding rack weight are provided to enable customers to plan appropriate power and cooling for the solution.

To support either of the configurations, the following components are required to be present in the customer environment:

- An existing Ethernet infrastructure with which to integrate is required.
- The addition of servers, switches, and iSCSI storage arrays to an existing or new datacenter requires planning for IP addresses and VLANs.
- Active Directory® Domain Services (AD DS) must be available on the network. The Hyper-V hosts are joined to an existing or new domain. Cluster Services also require AD DS. Consult with your Dell Sales and Services representatives for more details.
- Domain Name Server (DNS) must be available on the management network.
- Network Time Protocol (NTP) Server - NTP is recommended on the management network.
- Database to support System Center 2012 SP1 VMM - For a list of supported databases refer to: [Requirements for System Center 2012 - Virtual Machine Manager](#).
 - If IT Administrators wish to install VMM on the Dell PowerEdge R420 Management Server or as a VM, then a route must exist between the management server (physical or as a VM) and the database used.
- The database is presumed to have maintenance and backup configured as per the business needs of the customer.
- 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 50 with Hyper-V*.

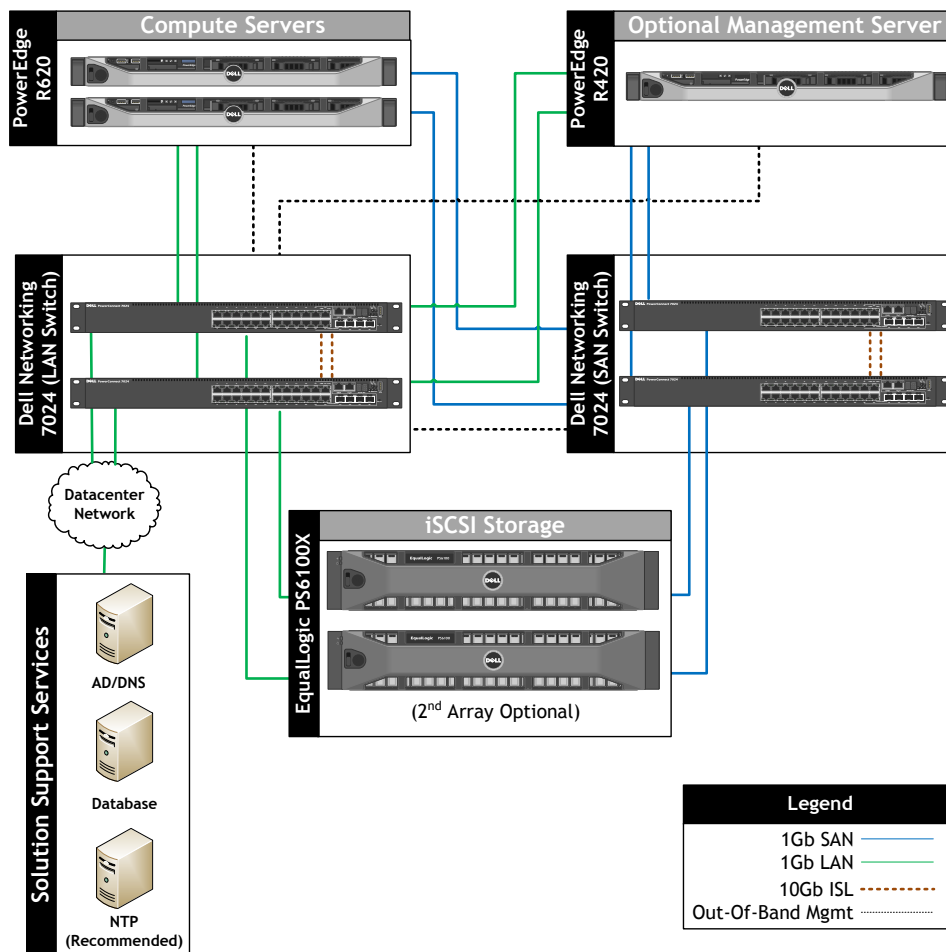
6 Architecture

This section describes the architecture and implementation detail for the Active System 50m. Hyper-V network traffic in this solution is comprised of five distinct types: Virtual Machine (VM), Management, Live Migration, Cluster Private, and iSCSI. In addition, support for Out-of-Band Management (OOB) is included. Two separate networks are created to support different traffic types:

- **LAN** - This network supports management, VM, Live Migration, Cluster Private and out-of-band management. In addition, uplinks to core infrastructure provide connectivity to the solution support services (AD DS, DNS, and NTP).
- **SAN** - This network supports iSCSI data. Uplinks are supported to connect into an existing iSCSI network; however, these uplinks are not required for full solution functionality. SAN switch out-of-band management also occurs on this network.

Figure 2 below depicts the high-level reference architecture for the solution including solution components and redundant connectivity for each I/O fabric.

Figure 2: Network Topology (Logical View)



6.1 Server / Network Connectivity

iSCSI storage redundancy is achieved with redundant NICs, switches, and storage controllers. For both network and iSCSI traffic, the redundant NICs are selected such that they are mapped across the rack Network Daughter Card (rNDC) and add-in controllers to avoid any single point of failure. Hyper-V High Availability (HA) is provided by Windows Server 2012 Failover Clustering.

The network traffic on each Server includes iSCSI, as well as traffic for the parent partition (hypervisor), Live Migration, cluster heartbeat and cluster shared volume, and child partitions (virtual machines).

Each Dell PowerEdge R620 rack server is configured with a Broadcom BCM5720 rNDC with four 1Gb ports, and a Broadcom BCM5719 add-in network adapter providing four 1Gb ports. Two ports on each network adapter (four total) are connected to the LAN 7024 switches. They are teamed using Windows Server 2012 native NIC teaming. Virtual network adapters are then created. In turn, the virtual adapters are shared with the parent partition, segmented with VLANs, and bandwidth weighted (BW).

The other four network connections are connected to the SAN 7024 switches. They are individually configured with IP addresses and Jumbo Frames to connect to the EqualLogic storage array. The EqualLogic Host Integration Tools (HIT) for Microsoft are installed on each compute node and manage the Multipath I/O (MPIO).

Each switch pair is connected together with a Stacking Module configured in Ethernet mode as an ISL Link Aggregation Group (LAG). The Stacking Module, in Ethernet mode, is configured into a 20Gbps LAG to provide a path for communication across the switches.

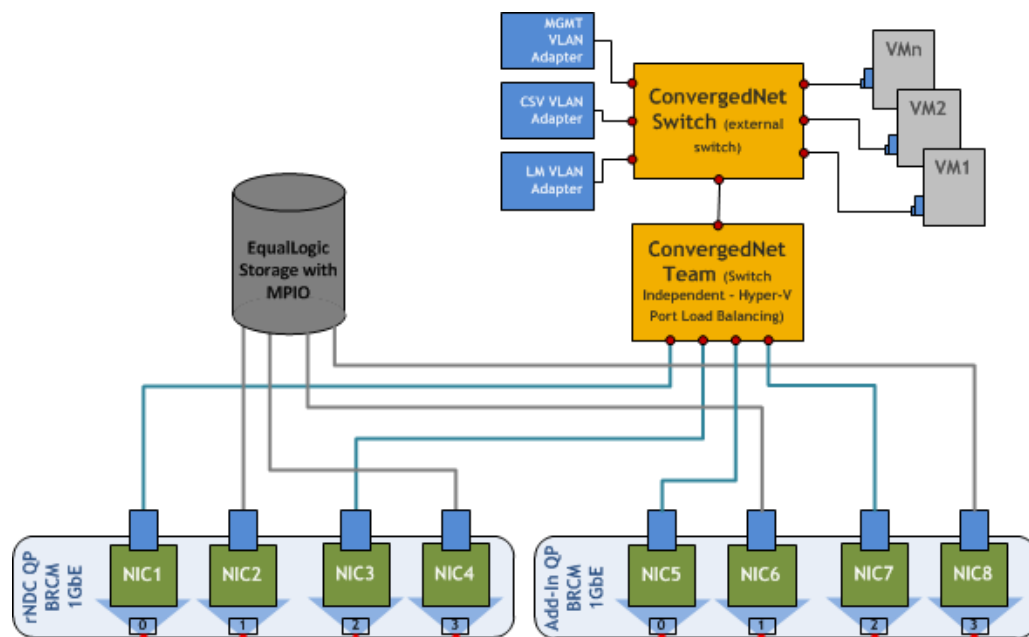
Spanning Tree Protocol (STP) PortFast is enabled on all the server and storage ports and allows the port to transition to forwarding traffic faster. Ports left for future expansion are disabled to prevent any unauthorized access or misconfiguration. The uplink ports on each switch are configured in a LAG. In addition, a separate LAG is created for the two Stacking Module Ethernet links.

The solution provides four 1Gb uplinks for LAN and eight 1GB uplinks for SAN from each switch to link into an existing core network infrastructure. If the core network infrastructure supports 10Gb Ethernet, then 10Gb uplink modules may be added to each switch; however, this option is beyond the scope of this document.

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.

On the Dell PowerEdge rack servers, each of the 1GbE ports are used for either LAN or SAN functionality. Four ports are used for LAN where one team is created and configured as a switch independent team using the Windows Server 2012 Load Balancing/Failover (LBFO) tools while four ports are individually configured and used for SAN connectivity to the EqualLogic storage. Figure 3 below illustrates the network configuration for the Dell PowerEdge servers.

Figure 3: Network Configuration for PowerEdge Rack Servers



The VLAN IDs are assigned to the virtual network adapters that in turn are assigned to the parent partition to segregate the traffic on the host between the cluster, Live Migration, parent management, and other types of traffic are described in Table 4 below. The VLAN configuration used in the Dell Active System configuration is listed in Table 5 below.

Table 4: Traffic Description

Traffic Type	Use
Hypervisor Management	Supports virtualization management traffic and communication between the host servers in the cluster.
Live Migration	Supports migration of VMs between the host servers in the cluster.
Tenant VM	Supports communication between the VMs hosted on the cluster and external systems.
Cluster	Supports internal cluster network communication between the servers in the cluster.
Out-of-Band Management	Supports configuration and monitoring of the servers through the iDRAC management interface, storage arrays, and network switches.
iSCSI Data	Supports iSCSI traffic between the servers and storage array(s). In addition, traffic between the arrays is supported.

Table 5: Sample VLAN and subnet configuration

Traffic Type	Sample VLAN	Sample Subnet
OOB	10	192.168.10.X
Management	20	192.168.20.X
Live Migration	30	192.168.30.X
Cluster Private	40	192.168.40.X
VM	100	192.168.100.X

6.2 Server / Storage Connectivity

In the Active System 50m configuration, each Dell PowerEdge server uses a Dell PERC H710 internal RAID controller connected to two SAS HDDs configured in a RAID 1. This RAID volume hosts Windows Server 2012 for the hypervisor OS.

Each server also includes Broadcom BCM 5720 rNDC and a BCM 5719 NIC for attaching to SAN volumes.

6.3 Server / HA and Redundancy

The Dell PowerEdge rack server design includes high availability and redundant features such as redundant fans and power supplies that are distributed to independent power sources. The servers also use PERC H710 controllers with two hard disks configured with RAID 1 to prevent server crashes in the event of single disk failures.

6.4 Storage Architecture

EqualLogic PS6100 storage provides capabilities essential to the Active System 50 design, like flexibility in configuring RAID arrays and creating volumes, thin provisioning, and storage tiering, while providing tight integration with Microsoft Windows Server 2012 for better performance and manageability through the use of The EqualLogic Host Integration Tools (HIT) Kit and SAN HQ.

6.4.1 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.

6.4.2 Storage Network

Each Dell PowerEdge server is configured with a Broadcom BCM5720 rNDC and a BCM5719 Add-in PCI-E Card. Ports from each adapter are connected to the Dell PowerEdge 7024 switches. The servers utilize iSCSI and EqualLogic MPIO for traffic flow from each node to the SAN.

6.4.3 Performance

Dell EqualLogic PS6100 Series uses redundant controller architecture with two active 1GbE front-end ports per controller. To allow the network port to dedicate the 4GbE bandwidth to SAN traffic, the 100MbE out-of-band port is used for configuration and performance monitoring.

6.4.4 Drive Types

The Dell EqualLogic PS6100 Series arrays use a high performance 10K RPM SAS drive in the default configuration. This provides a balance of size and performance. The Dell EqualLogic PS6100 is available with drives in multiple speeds and sizes, including 10K RPM and 15K RPM SAS drives, 7.2K RPM NL-SAS drives, and solid-state disks.

6.4.5 RAID Array Design

The storage array RAID configuration is highly dependent on the workload in your virtual environment. The EqualLogic PS series storage arrays support RAID types RAID 6, RAID 10, and RAID 50. The RAID configuration depends on workloads and customer requirements. In general, RAID 10 provides the best performance at the expense of storage capacity, especially in random I/O situations. RAID 50 generally provides more usable storage, but has less performance than RAID 10. RAID 6 provides better data protection than RAID 50.

6.4.6 Storage Tiering

Tiering storage is the practice of physically partitioning data into multiple distinct classes based on price, performance, or other attributes. Data may be dynamically moved among classes in a tiered storage implementation based on access activity or other considerations.

This is normally achieved through a combination of varying types of disks which are used for different data types. (i.e. Production, non-production, backups, etc.) Dell EqualLogic PS6100 Series dynamically moves data to the optimal storage tier based on actual use. The most active blocks reside on high-performance SSD and SAS drives, while infrequently accessed data migrates to lower-cost, high-capacity SAS drives. The result is network storage that remains in tune with application needs.

6.4.7 High Availability

In order to maintain continuous connectivity to stored data from the server, the controllers of the Dell EqualLogic PS6100 have port failover from active to standby controller. This port failover mechanism allows a port to fail over without having to fail over the controller. This results in a lower time when a controller is unavailable.

6.4.8 Multipath Configuration

For Windows Server 2012, the native MPIO provides functionality for Dell EqualLogic. The EqualLogic HIT kit provides additional load balance and tools. The multipathing solution uses the Round Robin load balancing algorithm to utilize all available paths. Dell recommends the EqualLogic HIT kit for its enhanced functionality and performance, but either can be use used.

6.4.9 iSCSI Authentication and Authorization

For authentication and authorization, the EqualLogic PS6100 Series Arrays provide IP address or IQN Name based Access Control Lists (ACL) or Challenge Handshake Authentication Protocol (CHAP. This feature allows hosts to be authenticated to the EqualLogic PS6100 iSCSI target prior to the target being available for use on the system.

6.4.10 Jumbo Frames

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

6.4.11 Thin Provisioning

Particularly in virtualization environments, thin provisioning is a common practice. This allows for efficient use of the available storage capacity. The volume and corresponding Cluster Shared Volume (CSV) may grow as needed, typically in an automated fashion to ensure availability of the volume. However, as storage becomes over-provisioned in this scenario, careful management and capacity planning is critical. Dell EqualLogic PS6100 Series offers thin provisioning and eliminates pre-allocated but unused capacity.

7 Management Infrastructure

The default management architecture consists of one PowerEdge R420 server running Microsoft Windows Server 2012 with Hyper-V role enabled. The management server is sized to run the Microsoft System Center 2012 suite of management components. At least one virtual machine needs to be deployed on the management Server when Active System Manager is to be installed. This VM hosts the Dell Active System Manager. Dell management tools such as OpenManage Essentials and EqualLogic SAN HQ can reside on the physical server or in a VM.

The management server network infrastructure duplicates the compute cluster in using a Hyper-V virtual network switch and virtual network adapters with weights and VLANs. The PowerEdge R420 server is connected to the PowerEdge 7024 switches using an onboard Broadcom BCM 5720 Dual Port 1Gb rNDC and a Broadcom BCM 5720 Dual Port 1Gb NIC. The management server is connected to the EqualLogic storage through the two 7024 switches.

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

7.1 Active System Manager

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

- 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

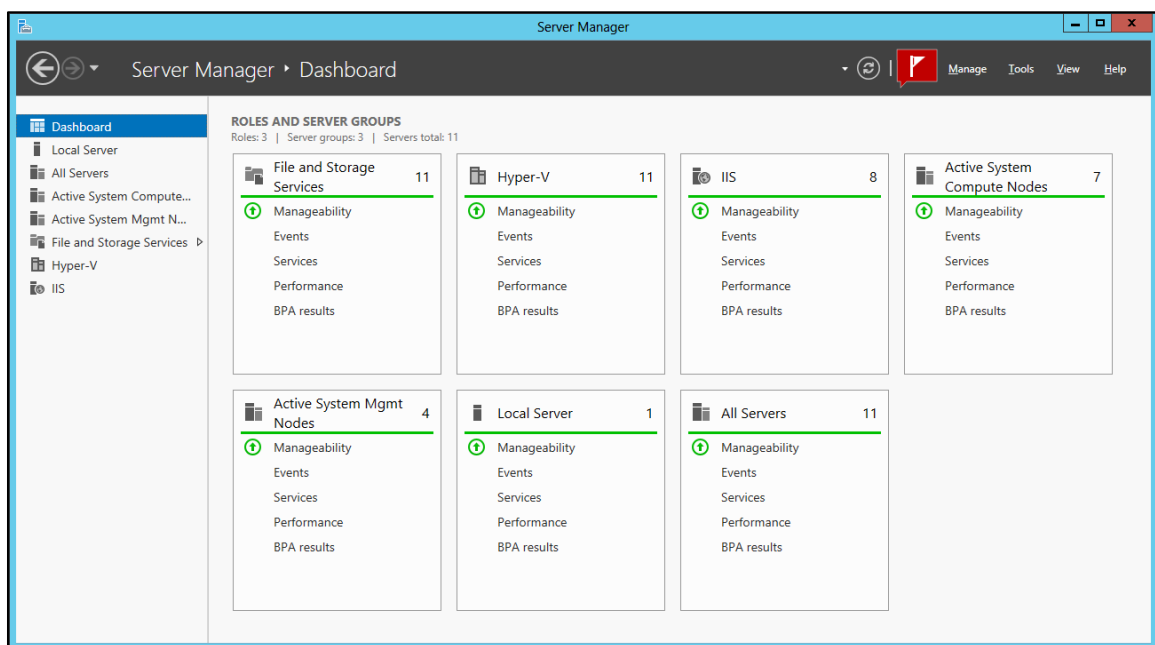
7.2 Windows Server 2012 Built-In tools

Windows Server 2012 provides the following built-in new or updated management features and tools for the virtualization solution:

Server Manager: Server Manager in Windows Server 2012 is re-designed as the primary management console for all the systems including the remote servers in your environment. The Server Manager can also manage the virtual machines and installed Windows roles on all the managed systems. It includes a role-based dashboard which can aggregate the data from managed servers. It reports system health state and surfaces issues which require your attention. The Server Manager also allows you to create customized groups for the managed systems.

Figure 4 below illustrates an example of managing with Server Manager.

Figure 4: Manage Active System 50m with Server Manager



Hyper-V Manager: In Windows Server 2012, more than 160 PowerShell cmdlets are added for Hyper-V. For host management, Hyper-V Manager allows the configuration of the physical GPU, Live Migrations, storage migration, and Hyper-V replication. For virtual machine management, Hyper-V Manager allows VM configuration options such as network hardware acceleration VM Queue (VMQ), Internet Protocol Security (IPSec), and Single Root I/O Virtualization (SR-IOV).

Failover Cluster Manager: With Failover Cluster Manager, the host configuration is validated for a Failover Cluster. Multiple failover clusters and high-availability virtual machines are created, configured and managed with the Failover Cluster Manager. In Windows Server 2012, the Failover Cluster Manager provides support for storage space, Hyper-V Replica, and new types of Live Migration. It also reports VM health states with more details, including Replication State, and allows monitoring of the services running inside VMs.

PowerShell 3.0: Several significant features are included in PowerShell 3.0 which ease and extend the management in the Windows-based environment. PowerShell 3.0 adds modules for Hyper-V,

storage, and network and more. Hyper-V cmdlets help automate the Hyper-V management tasks. With PowerShell 3.0, the resource utilization of VMs can be measured. This helps dynamic chargeback (vs. flat upfront cost) and capacity planning. Figure 5 below shows an example of how to measure a VM with PowerShell 3.0.

Figure 5: Resources Measuring in PowerShell 3.0

```

PS C:\Users\administrator.HER3> Enable-VMResourceMetering -VMName SQLSCVMDB
PS C:\Users\administrator.HER3> Add-VMNetworkAdapterACL -VMName SQLSCVMDB -Action Meter -Direction Inbound -RemoteIPAd
ress 10.10.0.0/16
PS C:\Users\administrator.HER3> Add-VMNetworkAdapterACL -VMName SQLSCVMDB -Action Meter -Direction Outbound -RemoteIPAd
dress 10.100.0.0/16
PS C:\Users\administrator.HER3> $report=Measure-VM -VMName SQLSCVMDB
PS C:\Users\administrator.HER3> $report | Format-List

```

ComputerName	: R610-2
VMId	: 280cdd17-e5aa-4a28-8706-caa1fa7ce05f
VMName	: SQLSCVMDB
MeteringDuration	: 00:03:44.6070000
AverageProcessorUsage	: 664
AverageMemoryUsage	: 4062
MaximumMemoryUsage	: 4096
MinimumMemoryUsage	: 4096
TotalDiskAllocation	: 61440
NetworkMeteredTrafficReport	: {Microsoft.HyperV.PowerShell.VMNetworkAdapterPortACLMeteringReport, Microsoft.HyperV.PowerShell.VMNetworkAdapterPortACLMeteringReport, Microsoft.HyperV.PowerShell.VMNetworkAdapterPortACLMeteringReport, Microsoft.HyperV.PowerShell.VMNetworkAdapterPortACLMeteringReport...}
AvgCPU	: 664
AvgRAM	: 4062
MinRAM	: 4096
MaxRAM	: 4096
TotalDisk	: 61440

```

PS C:\Users\administrator.HER3> $report.NetworkMeteredTrafficReport

```

LocalAddress	RemoteAddress	Direction	TotalTraffic(M)
::/0	Outbound	0	
10.10.0.0/16	Inbound	1529	
0.0.0.0/0	Inbound	1	
::/0	Inbound	1	
0.0.0.0/0	Outbound	7	
10.100.0.0/16	Outbound	0	
0.0.0.0/0	Inbound	126	
::/0	Outbound	0	
0.0.0.0/0	Outbound	1	
10.100.0.0/16	Outbound	8487	
10.10.0.0/16	Inbound	0	
::/0	Inbound	1	

CPU (MHz): Across hosts if LM.

Memory (MB): Average, Minimum and Maximum.

Disk(MB): VHD and the snapshots.

Network (MB): Use ACLs to separate traffics. Each ACL includes:

- Direction:** Inbound vs. Outbound
- Remote IP:** the source or destination of the network packet, depending on direction.

7.3 Dell OpenManage Essentials

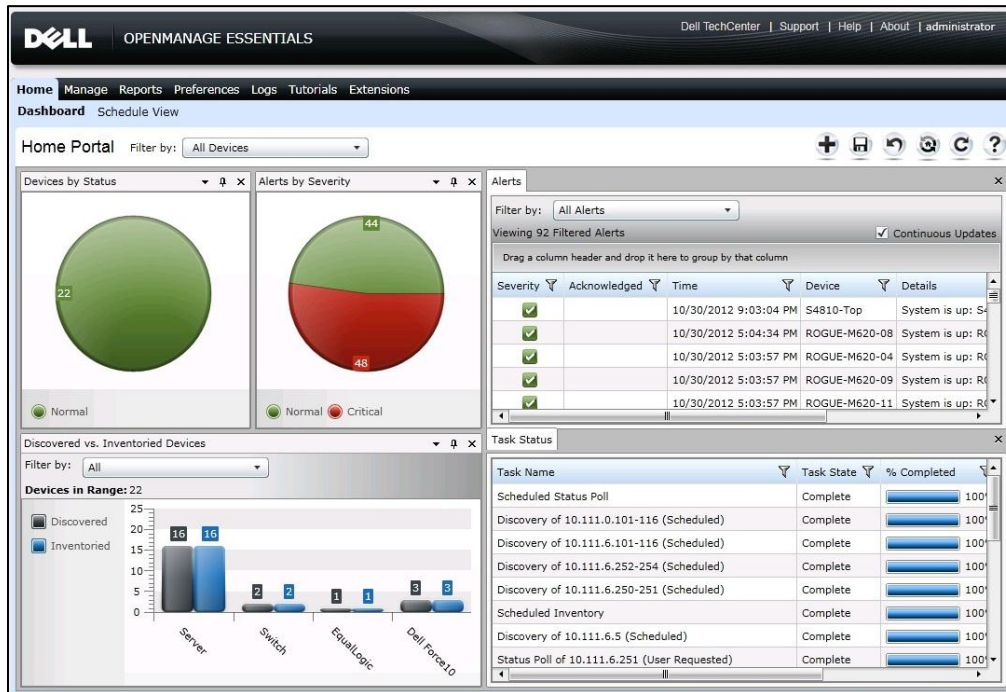
Dell OpenManage Essentials (OME) is a system management console that provides simple, basic Dell hardware management. Dell OME manages the host running Windows Server 2012. It supports agent-free hardware inventory and system update (e.g., BIOS, firmware) for the 11th and 12th generations of Dell PowerEdge servers. It also adds Dell SupportAssist extension for proactive “phone home” hardware issue resolution. Figure 6 shows the dashboard of Dell OME. Other features in Dell OME 1.1 are:

- Discover, inventory, monitor health, and view logs of Dell clients (i.e., desktops, portables)
- Discover, inventory, and monitor health of Dell Switches, Dell PDU and UPS devices
- Discover and collect limited inventory Linux servers via Secure Shell (SSH)
- Enhanced hardware inventory for ESXi servers with Dell OpenManage Server Administrator (OMSA)
- Alerts for hardware health status changes (not just when a hardware trap is sent)
- Ability to forward alerts in original format to support tiered event management
- A new command line interface (CLI) for automation of server discovery and grouping
- Calendar style task schedule view
- Group filtering on portal pages
- More configurable portal pages

- Grouped discovery ranges for common protocol settings
- Custom URL launch points to allow user-defined right-click actions on devices

For more information on OpenManage Essentials, see www.delltechcenter.com/ome.

Figure 6: OME Dashboard



Dell Lifecycle Controller: Helps to reduce operating costs by simplifying server deployment and management. Key features include diagnostics, self-update (UEFI, Driver Pack update), firmware updates (BIOS, NIC FW, RAID Controllers), and hardware configuration. Also, the integrated iDRAC provides the mechanism for OME agentless management.

Dell OpenManage Server Administrator (OMSA): Dell OMSA provides a comprehensive, one-to-one systems management solution. OMSA provides the mechanism for the agent-based management for Dell OME. With agent-based management, OME can update the device drivers on the host in addition to other management capabilities provided by the agentless management. For more information, see <http://www.delltechcenter.com/page/OpenManage+Server+Administrator+-+OMSA>.

7.4 Out-of-band Management: iDRAC

The iDRAC on each server provides the flexibility to remotely manage the server through console redirection and virtual CD-ROM/DVD/floppy/flash capabilities.

7.5 Dell EqualLogic SAN HQ

Dell EqualLogic SAN HQ provides comprehensive monitoring of performance and health statistics for one or more Dell EqualLogic PS Series groups. The tool is a client/server application that runs on a Microsoft Windows server and uses SNMP to query the groups. Acting like a “flight data recorder” on an aircraft, SAN HQ collects data over time and stores it on the server for later retrieval and analysis. Client systems connect to the server and format and display the data in the SAN HQ GUI. SAN HQ enables the following as examples:

- Monitor one or more PS Series groups and store operational data for up to a year
- Obtain a centralized view of the health and status of multiple groups
- View IO rates, throughput, and latency for each volume, member, pool, or group
- Allow the same performance data to be viewed by multiple clients simultaneously
- Monitor and analyze capacity usage for groups

If action is required, the EqualLogic Group Manager can be launched directly from SAN HQ giving quick access to adjust storage systems to meet business needs.

7.6 System Center 2012 SP1

System Center 2012 SP1 Virtual Machine Manager (VMM) is Microsoft’s Virtualization management platform. VMM 2012 provides in-depth management of both hypervisor and VMs. It provides a system administrator the capability to create and deploy VM templates, manage library stores of VMs, hardware profiles, and image files, and even manage VMware environments. VMM 2012 also provides P2V (physical to virtual) functionality, thereby allowing a system administrator to convert physical servers to virtual machines. VMM 2012 integrates with the hypervisor, VMs, and System Center Operations Manager to provide a deep view of the system utilization.

7.7 System Center 2012 Integration

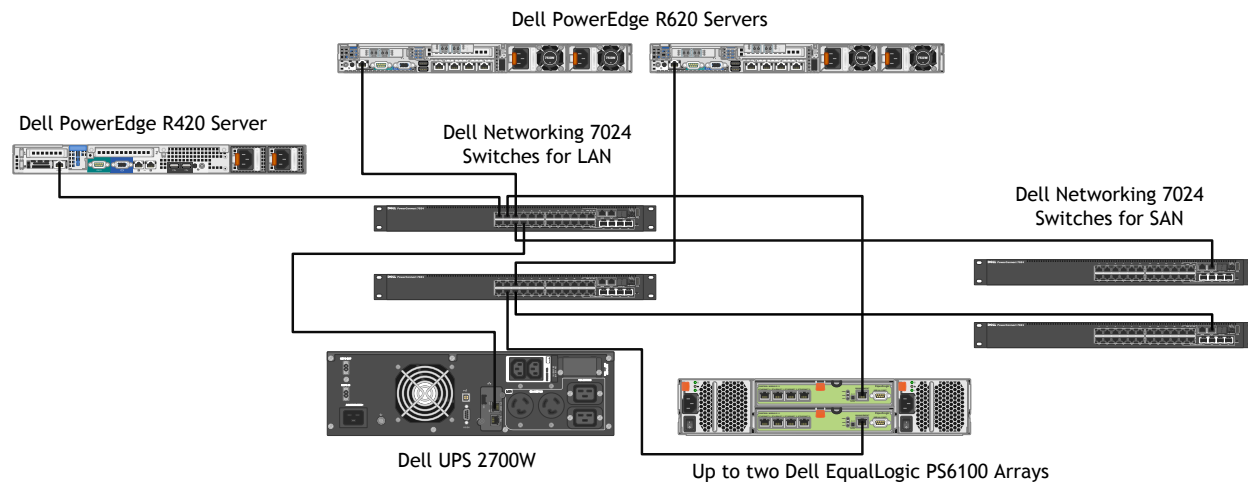
Dell provides management packs for the optional System Center components to monitor servers, storage, and networking components. These management packs allow System Center 2012 Operations Manager (SCOM) to monitor, report on, and take actions based upon alerts generated by the individual components. Dell Server PRO Management Packs can be integrated with VMM to monitor Dell servers and take remedial action when an inefficient system is identified.

Dell also provides integration tools for System Center Configuration Manager for the Lifecycle Controller, providing a framework for bare-metal deployment of servers and operating systems. Dell EqualLogic PS6100 Series also includes management tools for automating the deployment of virtual machines with VMM 2012 SP1. These optional components are beyond the scope of this document. Additional information can be obtained from a Dell Sales or Services representative.

7.8 Out-Of-Band Management Connectivity

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

Figure 7: Connectivity of OOB management components.



8 Connecting Active System 50 to Datacenter Core Network

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

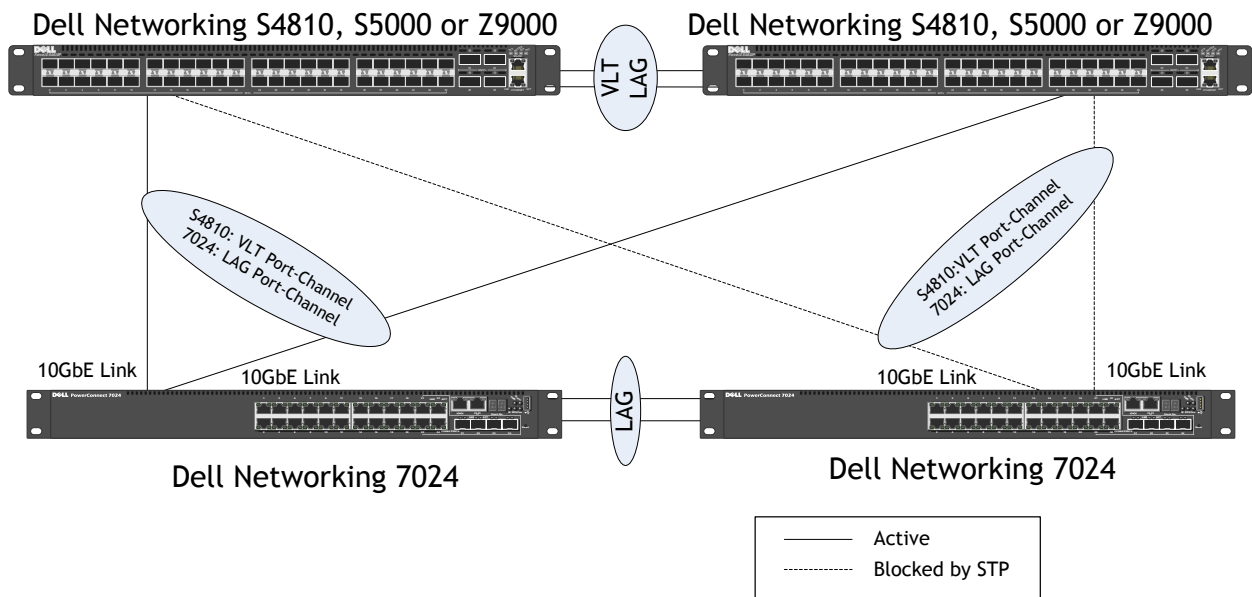
There are several ways to uplink the switches to the core network. Multiple uplinks can be created on each switch connecting to the core network switches. The solution provides four 1Gb uplinks from each switch to link into an existing core network infrastructure. If the core network infrastructure supports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

In this section, we provide examples to show how Active System 50 can be connected to a Datacenter network.

8.1 Connecting to Datacenter Network with Dell Networking Switches

Figure 8 below illustrates an example with Dell Networking S or Z series switches. The two datacenter network switches are connected together using VLTi. Each 7024 switch in the Active System uses a two 10GbE port LAG connecting to the two Dell Networking switches. In this example, one LAG is active, while the other is blocked by STP. Multiple 1GbE ports can also be used, which may require 1Gb Dell Networking transceivers.

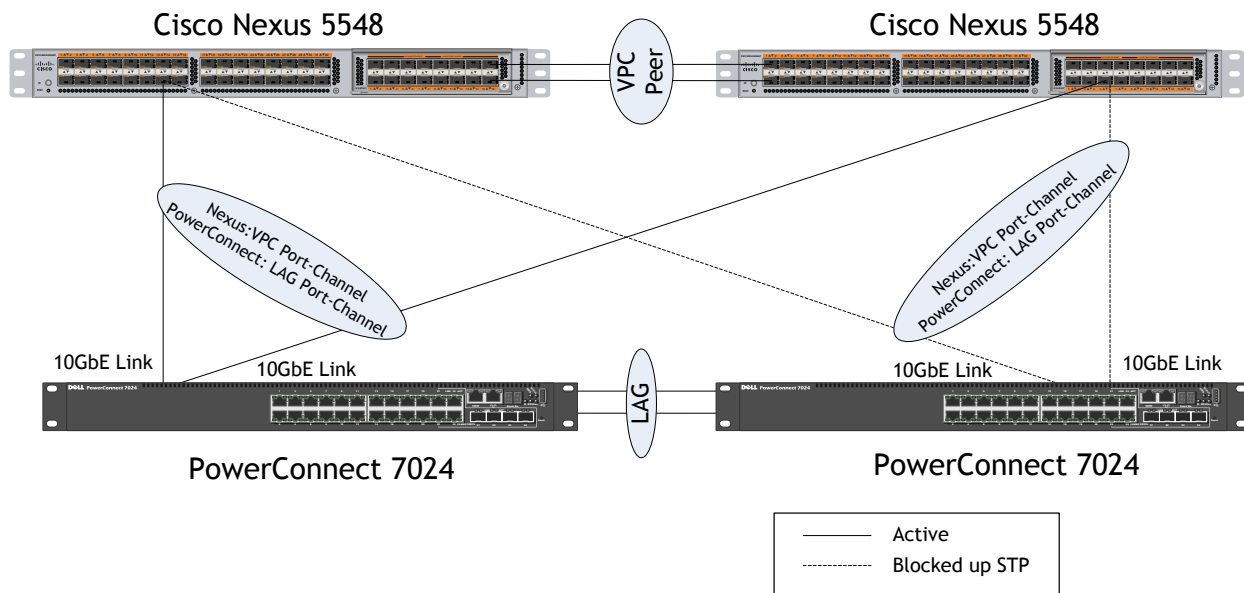
Figure 8: Connectivity to Dell Networking Datacenter Network



8.2 Connecting to a Cisco Datacenter Network

Figure 9 below illustrates an example with Cisco switches. The two Cisco switches can be connected together using VPC. Each 7024 switch in the Active System use a two 10GbE port LAG connecting to the two Cisco switches. In this example, one LAG is active, while the other is blocked by STP. Multiple 1GbE ports can also be used, which may require 1Gb Cisco transceivers.

Figure 9: Connectivity to a Cisco Network



9 Scalability

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

Scaling Compute and Network Resources: This solution is configured with four Dell Networking 7024 network switches. Only two Dell PowerEdge R620 Servers and one Dell PowerEdge R420 Server are available in the Active System 50 solution.

Scaling Storage Resources: Dell EqualLogic PS6100 Series storage can be scaled seamlessly and independent of the compute and network architectures. This design is currently limited to two enclosures.


10 Delivery Model

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

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

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

Figure 10: Active System 50m Rack and Component Overview



11 References

11.1 Dell PowerEdge Server Documentation and Hardware/Software Updates

Drivers and other downloads: Visit <http://support.dell.com>

Dell Networking Switch Documentation and Firmware Updates

<http://www.dell.com/support/drivers/us/en/04/Product/powerconnect-7024>

11.2 Dell EqualLogic

Dell EqualLogic PS Series Architecture Whitepaper

http://www.dell.com/downloads/global/products/pvaul/en/dell_equallogic_architecture.pdf

Host Integration Tools for Windows

<http://www.dell.com/downloads/global/products/pvaul/en/equallogic-host-software.pdf>

PS Series Storage Arrays: Choosing a Member RAID Policy

<http://en.community.dell.com/techcenter/b/techcenter/archive/2012/08/14/new-equallogic-raid-tech-report-considerations-and-best-practices-released.aspx>

http://en.community.dell.com/cfs-file.ashx/___key/communityserver-components-postattachments/00-19-86-14-80/TR1020_5F00_v44_5F00_RAID_5F00_Final.pdf

Using Tiered Storage in a PS Series SAN

http://en.community.dell.com/cfs-file.ashx/___key/communityserver-components-postattachments/00-19-86-14-79/TR1027_2D00_Tiered_2D00_Storage_2D00_v2.1_2D00_Final.pdf

Monitoring your PS Series SAN with SAN HQ

http://en.community.dell.com/cfs-file.ashx/___key/communityserver-components-postattachments/00-19-86-14-50/TR1050_2D00_Monitoring-Your-PS-Series-SAN-with-SAN-HeadQuarters.pdf

11.3 Dell Virtualization Documentation

Microsoft® Windows Server® 2012 With Hyper-V™ for Dell™ PowerEdge™ Systems Important Information Guide

ftp://ftp.dell.com/Manuals/all-products/esuprt_electronics/esuprt_software/esuprt_operating_system/ms-win-srvr-2012_User's%20Guide_en-us.pdf

Microsoft® Hyper-V Documentation

Windows Server 2012

<http://technet.microsoft.com/en-us/library/hh801901>

Windows Server 2012 Hyper-V Whitepaper

http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS%202012%20White%20Paper_Hyper-V.pdf

Reference Architecture for Active System 50 with Hyper-V

Failover Clusters in Windows Server 2012

<http://technet.microsoft.com/en-us/library/hh831579>

11.4 Microsoft® Management Software

Microsoft System Center 2012 SP1

<http://www.microsoft.com/en-us/server-cloud/system-center/datacenter-management.aspx>

System Center 2012 Virtual Machine Manager Deployment

<http://technet.microsoft.com/en-us/library/gg610669>