

Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with EqualLogic Storage

A Solution Guide for Small and Medium Business

By Shabana M and Lance Boley

Dell Virtualization Solutions Engineering

www.dell.com/virtualization/businessready

Feedback: solutionfeedback@dell.com

July 2010



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.

© 2010 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, PowerConnect, and OpenManage are trademarks of Dell Inc.; Intel and Xeon are registered trademarks of Intel Corporation; Microsoft and Windows are registered trademarks, and Hyper-V is a trademark of Microsoft Corporation in the United States and/or other jurisdictions. 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.

1 Table of Contents

2	Introduction.....	2
3	Audience and Scope	2
4	Overview	2
4.1	Dell PowerEdge R510 Server	3
4.2	Dell EqualLogic PS4000	4
4.3	Dell PowerConnect 5424 series Switch	4
4.4	PowerVault DL2200 Backup Solution.....	5
4.5	Solution Capabilities.....	5
5	Specification.....	6
6	Reference Architecture Diagram.....	8
6.1	Design Principles	8
6.2	Consolidation Solution.....	9
6.3	High Availability Solution	9
6.4	High Availability + Backup Solution	10
7	Network Architecture	11
7.1	Different types of network traffic.....	13
7.1.1	Virtual Machine Network/Cluster Public Network.....	14
7.1.2	Private / CSV and Live Migration Networks.....	14
7.2	iSCSI SAN Architecture	15
8	Configuring EqualLogic PS4000 Array.....	15
8.1	Storage Connectivity.....	15
8.2	EqualLogic Network Requirements.....	15
9	Management	16
9.1	Systems Management	16
9.2	Storage Management	16
9.3	Virtualization Management using Dell OEM Microsoft Systems Center Essentials 2010 Solution .	17
10	References.....	21
	Appendix A: Reference Network Configuration	21
	PowerConnect 5424 Configuration File for Public, Virtual Network, Live Migration, Private and CSV:..	22
	PowerConnect 5424 Switch Configuration File for iSCSI network:.....	22

2 Introduction

Business-Ready Configurations for Virtualization are a family of reference architectures offered by Dell and designed to aid with the ordering, deployment, and maintenance of a virtualization infrastructure. These architectures are designed to meet specific customer needs through the use of various server, storage, and virtualization technologies available from Dell.

The reference architectures defined in this document are targeted at small and medium business virtualization needs, although other environments may benefit from the provided information. The proposed infrastructure includes Dell™ PowerEdge™ R510 servers, Dell EqualLogic PS4000 series storage, Dell PowerConnect™ switches and Microsoft® System Center Essentials. Customers can quickly and confidently deploy these proven engineering architectures based on extensive design and engineering work, into production environments, thereby helping to eliminate much of the costly and time-consuming trial and error work often encountered during complex deployments. The solution is optimally configured to run virtualized workloads and is designed to provide redundancy with no single point of failure, as well as scalability and manageability. This document covers network architectures, storage configuration and best practices necessary for deploying and configuring the solution.

3 Audience and Scope

The intended audience for this white paper includes IT administrators, IT managers, and channel partners who are planning to deploy or resell Microsoft® Virtualization themselves or for their customers. These reference architectures provide an overview of the recommended servers, storage, software, and services. They can be used to plan, scope, and procure the required components to set up a virtualization infrastructure. It is assumed that the reader has a basic understanding of server virtualization (Hyper-V preferred), iSCSI storage, and networking concepts.

The solutions discussed in these reference architectures are based on Dell PowerEdge R510 servers, Dell EqualLogic PS4000 Series iSCSI arrays, and Microsoft Hyper-V R2 server virtualization. The solution is focused on Small Medium Business (SMB) deployments. Based on customer requirements, further customization of the recommended architecture may be required.

4 Overview

The reference architectures discussed in this white paper are centered on Microsoft's latest virtualization platform, Microsoft Windows Server 2008 R2 with Hyper-V (referred to subsequently as Hyper-V R2), PowerEdge R510 servers, and Dell EqualLogic PS4000 storage. The architectures leverage the benefits of the PowerConnect 5424 for both iSCSI and traditional networking, Dell PowerVault DL2200 disk-to-disk backup solution with CommVault®, and Dell OEM Microsoft System Center Essentials 2010 (subsequently referred to as SCE 2010), which provides an integrated solution for managing physical and Hyper-V R2 virtual infrastructures. This section provides an overview of the hardware components used for the Business-Ready Configuration followed by the key capabilities it is offering.

There are three reference architectures discussed in this reference architecture: Consolidation an entry level configuration, a High Availability configuration and a High Availability + Backup configuration. A high-level overview of the hardware included in each configuration is presented in table 1.

	Consolidation	High Availability	High Availability + Backup
Virtualization Hosts (Hyper-V R2)	1 x PowerEdge R510	2 x PowerEdge R510	4 x PowerEdge R510
External Storage	None	1 x PS4000X (16 x 600 GB 10k disk)	Up to 2 x PS4000X (16 x 600GB 10k disk)
Network Infrastructure	None	2 x PowerConnect 5424	4 x PowerConnect 5424
Backup Device	None	None	DL2200 (CommVault)
Virtualization Management	Hyper-V Manager	1 x PowerEdge R410 with SCE 2010	1 x PowerEdge R410 with SCE 2010
Key features	<ul style="list-style-type: none"> Ideal for consolidation or deployment of 9 virtual machines 	<ul style="list-style-type: none"> Ideal for consolidation or deployment of 9 virtual machines High Availability of virtual machines using Failover Clustering iSCSI based shared storage array (up to 9.6 TB of storage capacity) Single console management using SCE 2010 	<ul style="list-style-type: none"> Ideal for consolidation or deployment of 27 virtual machines High Availability of virtual machines using Failover Clustering iSCSI based shared storage array (up to 9.6 TB of storage capacity). Can scale up to 2 EqualLogic PS 4000 arrays Single console management using SCE 2010 Disk-to-Disk backup with optional disk-to-tape backup using DL2200

Table 1: Three pre-configured Solution Hardware Overview

Note: Virtual machine count mentioned in the above table is based on virtual machines with an average of 3GB of RAM, 40GB of shared storage space. Actual results may vary.

4.1 Dell PowerEdge R510 Server

The PowerEdge R510 is a two socket 2U rack server designed with energy-optimized options supporting Intel® Xeon® 5600 series processors, DDR3 memory and advanced embedded management capabilities.

The R510 is available in three different chassis configurations, a four hard drive chassis, an eight hard drive chassis, and a twelve hard drive chassis. In the eight hard drive chassis option, the PowerEdge R510 provides up to 8TB of internal storage. The R510 server has one dual-port Broadcom 5716 Gigabit network adapter. This server provides high availability capabilities with hot plug drives, redundant power supplies and redundant cooling.

The Lifecycle Controller is the engine for advanced embedded management and is delivered as part of the optional iDRAC Express or iDRAC Enterprise in the PowerEdge R510 server. The Lifecycle Controller helps to simplify administration tasks by performing a complete set of provisioning functions in a pre-OS environment, from a single intuitive interface called the Unified Server Configurator (USC). Tasks performed include system deployment, system updates, hardware configuration, and diagnostics.

4.2 Dell EqualLogic PS4000

The Dell EqualLogic PS4000 is the iSCSI storage device and is designed for deployment in a small to medium business environment. It consists of two 1Gb Ethernet network ports per controller for data and one 10/100-Mb port dedicated for the management network, fast processors, 2 GB cache per controller, support for RAID 6, and a monitoring application, SAN HQ, at no additional cost. In addition to the features described above, Dell EqualLogic storage arrays provide the following capabilities:

Reliability: Dell EqualLogic PS4000 Series arrays have hot-swappable redundant components, a choice of RAID types, and hot-spare disks. They also include the Auto-Stat Disk Monitoring System (ADMS) which proactively scans disk drives in the background to help detect media anomalies and correct them.

Scalability: As each array is added to the storage group, the storage capacity and performance, in terms of both bandwidth and operations per second, are increased. This increased capacity can be utilized without downtime. Thin-provisioning permits predefinition of a group of volumes which are larger than the physical space, allowing for physical space to be added and utilized without downtime when necessary. Note that a PS4000 array group supports up to two PS4000 arrays.

Self-Managing Arrays: The arrays offer many self-managing features such as automatic load balancing and storage tiering. A single storage pool can have different models that offer a range of capacity and performance parameters. In addition, different arrays in a storage pool can be configured with different RAID levels, and volumes will automatically be migrated between the RAID levels based on performance data and usage patterns. All data and volume movement can be performed online with zero downtime.

4.3 Dell PowerConnect 5424 series Switch

Dell PowerConnect 5400 series switches deliver 24 ports (PowerConnect 5424) or 48 ports (PowerConnect 5448) of wire-speed Gigabit Ethernet with advanced security and enterprise management features to help meet the needs of organizations of all sizes. To provide availability at the network layer, redundant Ethernet switches can be used in combination with NIC teaming on the Hyper-V R2 hosts to provide protection against the failure of a switch or other network device. The

PowerConnect 5400 series supports VLANs and up to eight Gigabit ports can be combined into a Link Aggregate Group (LAG), providing an aggregated bandwidth of 8 Gbps.

4.4 PowerVault DL2200 Backup Solution

The Dell PowerVault DL2200 Backup to Disk Appliance is an integrated solution of hardware and software powered by CommVault. It has the following key features:

Simplified and quick deployment experience: PowerVault DL2200 is factory installed with CommVault Simpana Software. The DL2200 provides configuration and management wizards which help users to configure and manage the appliance in a quick and easy way.

Integrated Tape Support: The PowerVault DL2200 is available with the Dell PowerVault TL2100, TL4000, or ML6000 tape library integrated into the full solution. Users can implement backup to disk for quick availability and then transfer to tape from the same management console for offsite disaster recovery.

Built-in Deduplication: The PowerVault DL2200 Powered by CommVault has a built-in compression and deduplication capability that stores blocks once, eliminating redundant segments across consolidated backup sets.

NOTE: Although Dell's PowerVault DL2200 Powered by CommVault is discussed here, the PowerVault DL2200 Powered by Symantec is also available.

4.5 Solution Capabilities

This section provides a brief description of the key solution capabilities offered in this Business-Ready Configuration.

- Virtual Machine Live Migration
 - The High Availability (HA) and (HA) + Backup configurations support Live migration of virtual machines between two or more Hyper-V R2 hosts. Live migration is a new feature in Hyper-V R2 in which a virtual machine on one host can be migrated to another host for load balancing, or for physical host maintenance without causing downtime or interruption to the running virtual machines.
- Virtual Machine High Availability using Failover Clustering
 - In the HA and HA + Backup reference architectures, the Hyper-V R2 host servers are configured in a Failover Cluster to provide high availability for the virtual machines. Failover clustering is when one server fails and the virtual machines previously hosted on that node are failed over and restarted on the surviving server node without administrator intervention.
- Support for Cluster Shared Volumes (CSV)
 - CSV in Hyper-V R2 provides the capability to host multiple virtual machines on a single storage volume and to migrate those virtual machines independently among the servers

in the cluster. Using CSV, multiple nodes can read and write to the volume irrespective of which server node owns the volume.

- Management using Dell OEM Microsoft System Center Essentials 2010 Solution
 - Provides a unified Solution with a single console for managing physical and virtual environments, monitoring and reporting alerts.
 - SCE 2010 allows downloading and importing management packs for monitoring the hardware related events in the environment. Management packs are available for Dell PowerEdge Servers, and for PowerVault and EqualLogic storage arrays.
 - Dell Server PRO Management Pack 2.0 integrates with Dell OpenManage to monitor events such as loss of power supply redundancy, exceeding temperature thresholds, server storage battery errors, and internal disk failures. The PRO Tips generate support actions such as the live migration of all virtual machines off of the alerting host.
 - SCE 2010 has many features which help users deploy, manage, and monitor their virtual environment. Template-based virtual machine creation is one such feature. Administrators can build templates for creating virtual machines and store them in the SCE 2010 library. This speeds deployment of virtual machines in the virtual environment.
 - Monitoring the resource utilization: SCE 2010 monitors the resource utilization of the Hyper-V R2 host servers and the running virtual machines, which helps users to manage load balancing the virtual machines the virtual environment.

5 Specification

The following table provides the detailed hardware specifications for the three reference architectures provided in this document.

	Consolidation Solution	High Availability Solution	High Availability + Backup Solution
Hardware			
Solution ID	1100031.1	1100042.1	1100051.1
Hyper-V R2 Server	1 x R510	2 X R510	4 x R510
Storage Device	Local Storage (8 Internal HDD)	(1) PS 4000	(1) PS4000 (Can scale up to two arrays)

Infrastructure Server and Management Server	N/A	(1) x R410	(2) x R410
Network Configuration			
Network Configuration	N/A	2 x PowerConnect 5424 Network Switches	4 x PowerConnect 5424 Network Switches
Back up Device	N/A	N/A	DL2200 Disk Based Backup Solution Powered by CommVault
Hyper-V R2 Server Configuration			
Rack Server Model	R510	R510	R510
Processor	(2) x Intel Xeon (Westmere) E5630, 2.53Ghz, 12M Cache	(2) x Intel Xeon (Westmere) E5630, 2.53Ghz, 12M Cache	(2) x Intel Xeon (Westmere) E5630, 2.53Ghz, 12M Cache
Memory	32 GB (8 x 4GB, DDR3)	32 GB (8 x 4 GB, DDR3)	32 GB (8 x 4 GB, DDR3)
Onboard Network Adapter	(1) Integrated Broadcom 5716 Dual port GbE adapter	(1) Integrated Broadcom 5716 Dual port GbE adapter	(1) Integrated Broadcom 5716 Dual port GbE adapter
Add-in Controllers		(2) Broadcom 5709 Dual Port GbE PCI/e Card	(2) Broadcom 5709 Dual Port GbE PCI/e card
Internal Storage Controller	PERC 6/i	PERC 6/i	PERC 6/i
Internal Hard Drive	(3) X 500GB, SATA, RAID5 Up to (8) X 500GB, SATA	(2) x 146 GB in RAID1	(2) x 146GB in RAID1
Hypervisor	Microsoft Hyper-v R2 Enterprise Edition	Microsoft Hyper-v R2 Enterprise Edition	Microsoft Hyper-v R2 Enterprise Edition
External Storage Configuration			
Storage Device	None	(1) PS4000	(2) PS4000
Drives	None	16 X 600GB, 10K SAS	16 X 600GB, 10K SAS

Storage Capacity	None	9.6 TB per device	9.6 TB per device
Infrastructure and Management Server			
Rack Server Model	None	(1) R410	(2) R410
Processor	None	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) x Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache
Memory	None	8 GB	8 GB
Network Controller	None	(1) Broadcom 5709 Dual Port GbE PCI/e card	(1) Broadcom 5709 Dual Port GbE PCI/e card
Operating System	None	Microsoft Windows Server 2008 R2 Standard Edition	Microsoft Windows Server 2008 R2 Standard Edition
Virtualization Management Software	Native Hyper-V Manager	Dell OEM Microsoft System Center Essentials (SCE) 2010 with Dell PRO Pack Management Pack 2.0	Dell OEM Microsoft System Center Essentials (SCE) 2010 with Dell PRO Pack Management Pack 2.0
Server Management	Open Manage 6.2.1	Open Manage 6.2.1	Open Manage 6.2.1

Table 2: Specification for Reference Architecture Configuration

6 Reference Architecture Diagram

This section describes the reference architectures for this solution. As mentioned before, there are three pre-configured solutions for this Business-Ready Configuration: the Consolidation Solution, the High Availability Solution, and the High Availability + Backup Solution.

6.1 Design Principles

The following design principles were used while designing these configurations.

1. **Choosing optimal configuration for Small and medium Business:** The solution components for this Business-Ready Configuration are specifically designed for the requirements of an SMB environment. The hardware components chosen are suited for virtualization in a SMB. SCE 2010 is chosen as the management software which is an integrated solution for managing the physical and virtual environments and is designed for SMB customers.
2. **High Availability and Live Migration Capabilities:** Configuring Hyper-V R2 hosts in Microsoft Failover Clustering enables virtual machines to move between hosts using live migration with minimal downtime. This also protects virtual machines from having a single point of failure.

3. **Isolated and redundant network architectures:** The network configuration used in the reference architecture provides isolation for different networks in the virtual infrastructure and avoids a single point of failure.
4. **Virtualization Management Capabilities suited for SMB:** Management of the reference configuration discussed in this document uses SCE 2010, which is best suited for the SMB customer. It supports management of up to 50 server operating systems environments and 500 client operating system environments. It is a unified solution with a single console for managing both physical and virtual environments.

6.2 Consolidation Solution

This configuration can be used to provide an entry-level virtualization solution for a small business environment. The configuration consists of a single PowerEdge R510 with Microsoft Hyper-V R2. In this configuration, the internal storage is used for storing the virtual machine files. In the reference architecture discussed here, the R510 with an 8 drive chassis is used with three 500GB SATA hard drives in RAID 5 configuration. To increase the internal storage, it can scale up to 8 internal disks. The built-in Hyper-V R2 Manager can be used to manage the virtual environment for this Solution.

6.3 High Availability Solution

The High Availability Solution consists of two Hyper-V R2 hosts deployed on PowerEdge R510 servers configured in a failover cluster. These are connected to a single EqualLogic PS4000 storage array in an iSCSI SAN. The network infrastructure consists of two PowerConnect 5424 switches and a PowerEdge R410 running SCE 2010 management server, which is used to deploy the virtual infrastructure. Other infrastructure services include Active Directory, domain name resolution (DNS), and DHCP server.

In Figure 1, cabling connections are only shown on one R510 server. Cabling connections for the other Hyper-V R2 server are the same.

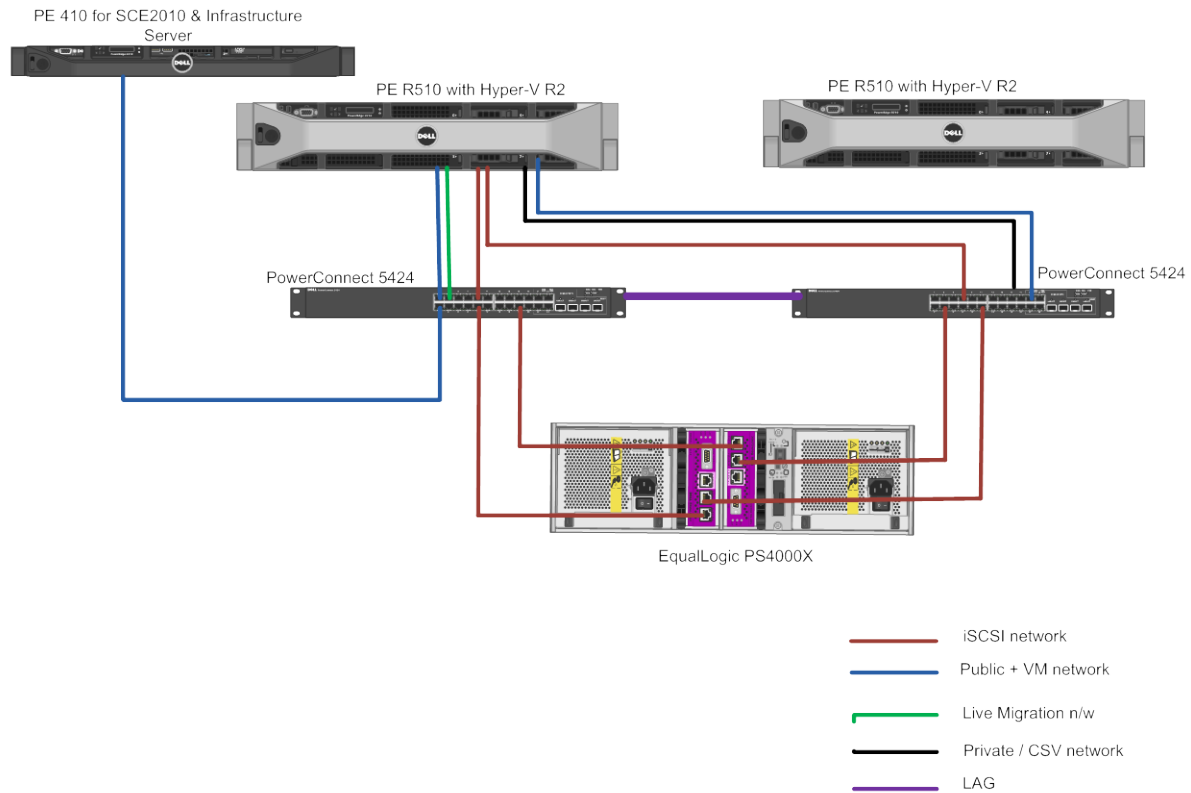


Figure 1: Reference Architecture diagram for the High Availability Solution

6.4 High Availability + Backup Solution

The reference architecture diagram for the High Availability + Backup configuration is as shown in Figure 2 below. The solution consists of four PowerEdge R510 servers running Hyper-V R2 configured in a failover cluster. This configuration connects two EqualLogic PS4000 arrays. The network infrastructure consists of four PowerConnect 5424 switches. Two switches are used for dedicated iSCSI traffic while the other two switches isolate network traffic. There are two PowerEdge R410 servers used in this configuration, with the first R410 server used as a management server while the second optional R410 server is used for infrastructure services (AD, DNS and DHCP). The reference architecture diagram shows how the hosts are connected to a single PS4000 array. For scalability, two PS4000 arrays can be connected in this configuration. Refer to the EqualLogic documentation for best practices in adding a second PS array to the group.

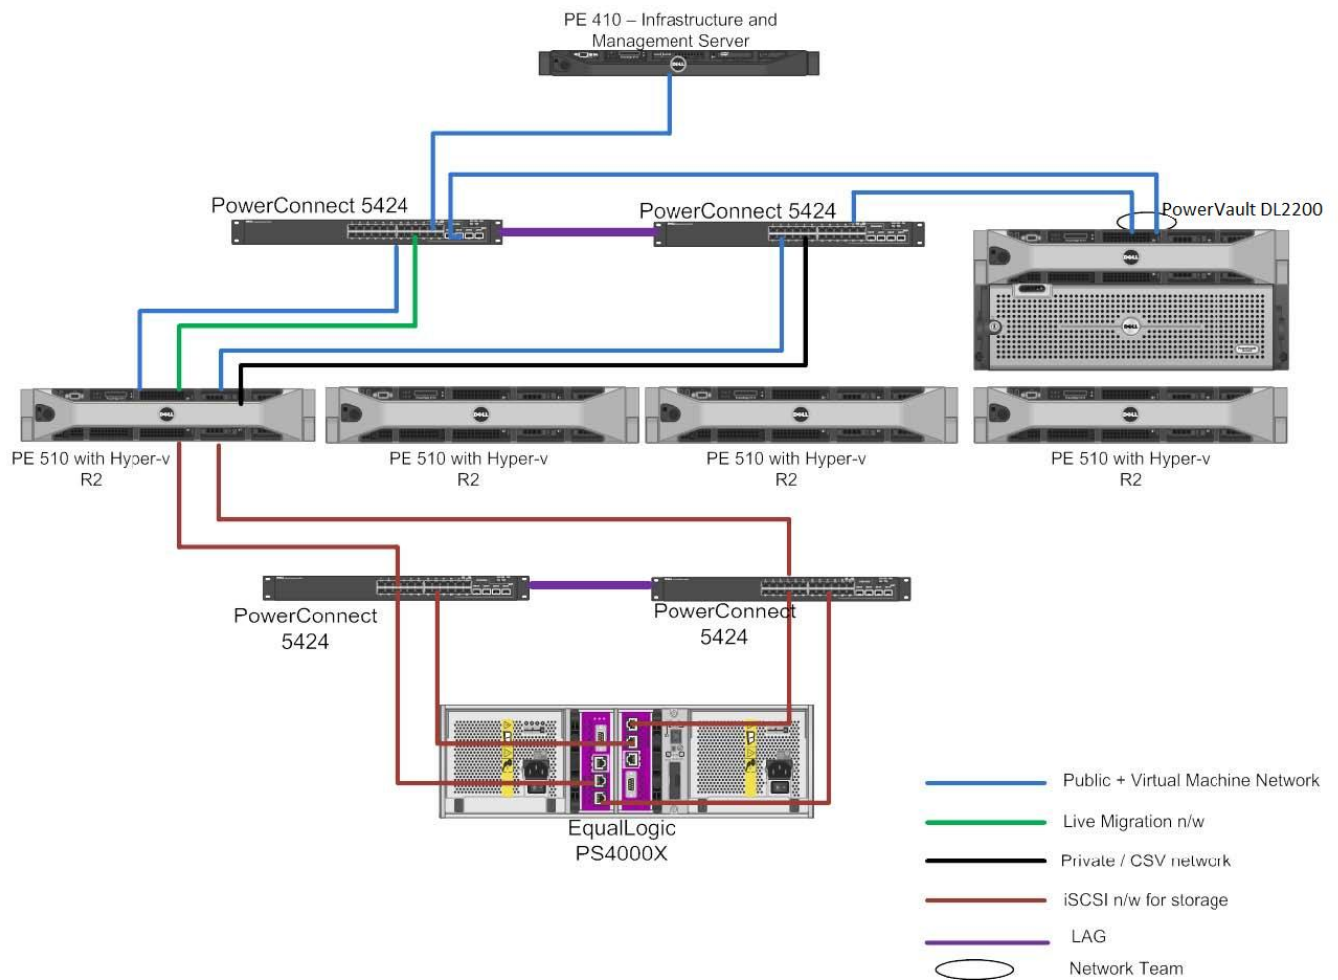


Figure 2: Reference architecture for High Availability + Backup Solution

7 Network Architecture

The Hyper-V R2 virtual infrastructure is comprised of six network traffic types: the virtual machine network traffic, cluster public network traffic, cluster private network traffic, network traffic for the cluster shared volume, live migration network traffic, and iSCSI traffic. This section describes the configuration for the physical network, PowerConnect switches, and the virtual networks which provide load balancing and high availability. Gigabit Ethernet is used for all network traffic types in the reference configurations.

In both High Availability and High Availability + Backup configurations, each host has six network ports per R510 server, two LAN On Motherboard (LOM) gigabit Ethernet ports, and two dual port network adapters.

- *LOM #1, LOM #2*: First and second LOM ports on the 5716 dual port Gigabit Ethernet controller integrated on the system board.
- *NIC 3 and NIC 4*: First and second ports on the first dual-port Broadcom NetXtreme II 5709 Gigabit Ethernet controller.
- *NIC 5 and NIC 6*: First and second ports on the second dual-port Broadcom NetXtreme II 5709 Gigabit Ethernet controller.

To support the different traffic types in a virtualized environment, you can use two or four Dell PowerConnect 5424 switches. The PowerConnect switches are interlinked using a 4 port Link Aggregation Groups (LAG), as shown in figure 3 and 4. This provides optimal bandwidth for the traffic between the switches, as well as failover capabilities should a link fail. The PowerConnect switch supports a maximum of eight Gigabit ports combined into a LAG, providing an aggregated bandwidth of 8 Gbps.

In the High Availability configuration, all network traffics types share both physical switches and are isolated using VLANs. In the High Availability + Backup configuration, four switches are used. The first two switches are dedicated for iSCSI SAN traffic. The second two switches are used for the virtual machine Public Cluster, Private/CSV, and Live Migration networks, which are separated by VLANs. Figures 3 and 4 show how the network connections are setup for a single Hyper-V R2 host in a High Availability and High Availability + Backup configuration.




Figure 3: Single server network connection in High Availability Configuration




Figure 4: Single server network connection in High Availability + Backup Configuration

7.1 Different types of network traffic

Following Microsoft Best Practices, network traffic is broken down into separate networks, and each network type is isolated using VLANs. As mentioned previously there are six different network traffic types in reference configurations for Hyper-V R2. Below are the network traffic types and a brief description of each.

iSCSI network (VLAN ID 10): This network type is used for iSCSI communication between Hyper-V R2 hosts and an EqualLogic PS4000 storage array. In the High Availability and High Availability + Backup configurations, two network switches are used for dedicated for iSCSI communication. No VLANs are needed to isolate this traffic.

Cluster Public Network or Management (VLAN ID 20): Used by a cluster for public or external communication. This network is also used by the SCE 2010 management server to connect to the cluster for management purposes.

Virtual Machine Network (VLAN ID 20): Used for virtual machine connectivity from the Hyper-V R2 host servers to virtual machines, applications, and services to the rest of the network. In the reference configuration, the virtual machine network shares a teamed network connection with the cluster public network, avoiding single point of failure for both public and virtual machine network traffic. Both networks use the same VLAN. Different VLANs can also be used to isolate the two.

Cluster Private Network (VLAN ID 30): The private network traffic used by a cluster for intra node communication.

Network for CSV (VLAN ID 30): This network type is used by the server nodes in the failover cluster for cluster shared volumes. The reference configuration uses the same VLAN ID as that of Cluster Private Network.

Network for Live Migration (VLAN ID 50): This Network is used for the live migration traffic between Hyper-V R2 hosts in the failover cluster.

Type	Sample VLAN	Sample Subnet
Cluster public/management	20	172.20.1.X
Cluster private/CSV	30	172.30.1.X
Virtual machine	20	172.20.1.X
Live migration	50	172.50.1.X

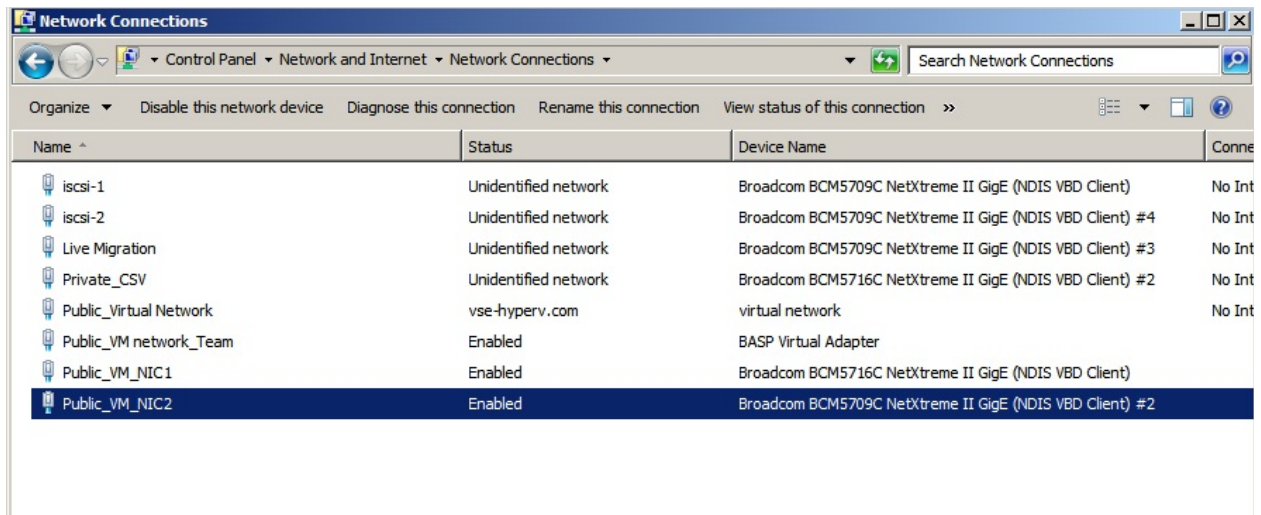


Figure 5: The Network Connections Manager showing different networks configured in the Hyper-V R2 hosts.

7.1.1 Virtual Machine Network/Cluster Public Network

These two network traffic types are configured consisting using two Broadcom network adapters and the Broadcom Advanced Configuration Suite (BACS) to perform Smart Load Balancing (SLB) with failover capabilities. One network adapter is configured as the active member of the team, while the other is a stand-by network adapter used to avoid a single point of failure for both the virtual machine and cluster public networks. These NICs are physically connected to the ports configured for VLAN ID 20. As mentioned earlier, this network is also used by the SCE 2010 management server for discovering, monitoring, and management of the physical and virtual systems. In Figure 5, the network listed as the “Public_Virtual Network” is the virtual network created in external mode and is used for the Cluster public and Virtual Machine networks. It uses the Broadcom Advanced Server Program (BASP, installed with BACS) Virtual teamed adapter for network connection. In the configuration, LOM 1 and NIC 5 are used to create the network team.

7.1.2 Private / CSV and Live Migration Networks

The private Cluster / CSV and Live Migration networks can be configured in such a way that whenever there is a failure of one NIC, the other NIC will allow the failed network traffic to flow through it. For the cluster private / CSV and live migration networks on each host, LOM 2 and NIC 3 are used. Both

network teams are connected to different switches to avoid a single point of failure. NICs for the private /CSV networks are connected to the switch ports configured for VLAN 30 and the NIC for live migration is connected to switch port configured for VLAN 50.

The following section discusses the iSCSI network architecture used in the reference configuration.

7.2 iSCSI SAN Architecture

The iSCSI SAN is comprised of iSCSI traffic to the Dell EqualLogic PS4000 Series array. This section discusses the best practices for configuring the iSCSI SAN using the Microsoft Software iSCSI initiator provided in Windows Server 2008 R2.

HA Configuration: The two 1Gb Ethernet network ports used for iSCSI communication on the Hyper-V R2 hosts are physically connected to the same PowerConnect switches and are used for both traffic types, but are separated by VLANs. Hence, all six of the different traffic types share the same physical switches, but they are isolated using VLANs. The two NICs are connected to switch ports configured for VLAN 10. Since the PS4000 series array controller ports provide load balancing, both of the NICs should be configured to use the same subnet.

HA + Backup Configuration: The only difference between HA and HA + Backup configuration is that in the HA + Backup configuration, the network adapters used for iSCSI traffic are connected to two separate PowerConnect 5424 switches for dedicated iSCSI traffic. Since the traffic is isolated physically, there is no need to configure VLANs for iSCSI traffic in this configuration.

In both configurations, a two-port LAG is configured between the switches to ensure high availability and bandwidth. The number of ports added to the LAG is equivalent to the number of active data ports in the PS array group. The Microsoft software iSCSI initiator should be set up to establish two sessions per target, one session for each initiator IP address, so that all hosts can see two paths to each volume. Multiple paths are managed using the EqualLogic host integration toolkit (HIT kit) MPIO driver for Windows Server 2008 R2. Additional information on the HIT kit and MPIO driver are available at <https://www.equallogic.com/support/>.

8 Configuring EqualLogic PS4000 Array

8.1 Storage Connectivity

Each controller in the PS4000 array has two 1Gb Ethernet ports for iSCSI traffic and one 10/100 Mb port for management traffic. The two 1Gb Ethernet ports for iSCSI traffic are connected to the two switches, one from each controller to each of the PowerConnect 5424 switches.

8.2 EqualLogic Network Requirements

In addition to the guidelines discussed previously, the EqualLogic array has specific recommendations for connecting PS Series arrays to the network. Some of the important recommendations have been

highlighted; however, for more information see the Dell EqualLogic PS Quick Start Guide at <https://www.equallogic.com/support/> (account registration may be required).

- Spanning-Tree protocol (STP) should not be used on switch ports that connect end nodes (iSCSI initiators or array network interfaces). If you want to use STP or Rapid STP (preferable to STP), you should enable the port settings (FastLink or Port Fast) available on some switches which lets the port immediately transition into the STP forwarding state upon link-up. This functionality can reduce network interruptions that occur when devices restart and should only be enabled on switch ports that connect end nodes.

Note: The use of Spanning-Tree for a single-cable connection between switches is encouraged, as is the use of trunking for multi-cable connections between switches.

- Enable Flow Control on each switch port and NIC that handles iSCSI traffic. The PS Series arrays will correctly respond to Flow Control.
- Disable unicast storm control on each switch that handles iSCSI traffic if the switch provides this feature. However, the use of broadcast and multicast storm control is encouraged on switches.
- Enable Jumbo Frames on the switches and the NICs used for iSCSI traffic.
- Disable iSCSI optimization on the PowerConnect 5424 switches used for iSCSI traffic.

9 Management

This section describes some aspects of managing and deploying the reference architectures.

9.1 Systems Management

Dell Open Manage Server Administrator can be used for managing the Hyper-V R2 servers used in the reference architectures. Dell Open Manage Server Administrator is a web-based tool which can be used for managing and monitoring the server hardware, system health, and performance. For information on Dell OpenManage and its capabilities, see <http://www.dell.com/openmanage/>

The Dell Unified Server Configurator (USC), which ships with all current generation PowerEdge servers, can be used for easy deployment and configuration of the host servers. The USC can be used for BIOS, firmware and driver updates. It can also be used for configuring the server hardware and RAID controller.

9.2 Storage Management

The Dell EqualLogic array provides a rich set of management features that are available at no additional cost and come with exceptionally easy-to-use management tools. **SAN HeadQuarters (SAN HQ)** is Dell's new management tool for consolidated performance and robust event monitoring across multiple virtualized EqualLogic SAN groups. It gathers and formats performance data and other vital group information for easy viewing. With SAN HQ, you have the ability to centrally monitor alerts, get

historical performance reporting, do trend analysis for capacity planning, and troubleshoot array alerts. For more information on Dell EqualLogic features and management capabilities, see <http://www.dell.com/equallogic/>.

In addition to SAN HQ, SCE 2010 can monitor the Dell EqualLogic Management Pack installed on the PowerEdge R410 management server. Dell EqualLogic Management Pack collects the alerts and information and presents them to the SCE console. Customers can thus use a single console to monitor the server, storage and virtual machines in their environment.

9.3 Virtualization Management using Dell OEM Microsoft Systems Center Essentials 2010 Solution

Windows Server 2008 R2 Enterprise natively provides the ability to locally manage Hyper-V R2 servers with limited management functionality such as power cycle and migrate, using the Hyper-V Manager. The Microsoft Failover Cluster Manager snap-in is another natively supported management feature. This snap-in also helps users to manage the virtual machines with limited functionality.

Dell OEM Microsoft System Center Essentials 2010 helps users manage virtual machines as well as the physical systems using the same console. SCE 2010 is a Microsoft System Center Family product and includes features from other system center products like System Center Operations Manager (SCOM), Systems Center Configuration Manager (SCCM), and System Center Virtual Machine Manager 2008 R2 (SCVMM R2), which are suitable for a small and medium sized business. SCE 2010 provides an easy way of deploying, monitoring, and managing virtual machines. Some of the key features of SCE 2010 include: rapid deployment of virtual machines using templates, intelligent placement of virtual machines based on host work load, and Dell PRO pack management packs with PRO tips. PRO tips help administrators monitor Hyper-V R2 system health and provide recommended implementation tips.

SCE 2010 can be deployed on the same Infrastructure Server (the PowerEdge R410) used for configuring the infrastructure services like DNS, DHCP, and Active Directory. Optionally, it can also be deployed on a separate R410 server if those services are not already present on the network.

Following are provides guidelines for the initial deployment and configuration of the solution.

Managing the EqualLogic PS4000 array, the PowerConnect 5424 switch, and the PowerEdge servers through iDRAC can be done using existing network infrastructure.

1. Perform physical configuration

- Rack and cable the components of the reference architecture.
- Connect power to the servers, storage and switches.
- Cable servers, storage and switches.
 - Cable the server and storage for Ethernet connectivity as per the architecture.
- Interlink the network switches.

2. Configure PowerConnect switches

- Perform initial configuration of the switch. Assign management IP address for the PowerConnect 5424 switches.
- Connect the management port to the existing hardware infrastructure management.

3. Perform initial configuration of PS 4000

- Configure the 1G Ethernet ports for iSCSI and the dedicated management port for array management. For information on configuring storage, see the Dell EqualLogic PS QuickStart Guide at <https://www.equallogic.com/support/> (account registration may be required).
- Create volumes for virtual machine storage. When creating a volume using the EqualLogic group manager user interface that will be shared across Hyper-V R2 hosts, select the “Allow simultaneous connections from initiators with different IQN names” option.

4. Configure the infrastructure server

- Deploy Microsoft Windows Server 2008 Standard on the infrastructure server.
- Install Active Directory Domain Services, DNS and DHCP if required.
- Configure a domain for the virtual infrastructure using the command *dcpromo*.

5. Perform initial configuration of the Hyper-V R2 Servers

It is recommended to review the *Microsoft® Windows Server® 2008 R2 With Hyper-V™ for Dell™ PowerEdge™ Systems Important Information Guide* available at http://support.dell.com/support/edocs/software/win2008/WS08_R2/en/IIG_HyperV/IIG_HypV.pdf. Among other things, this guide provides updates on the latest known issues and resolutions

On each Hyper-V R2 server perform the following:

1. Ensure hardware-assisted virtualization (Intel-VT) is enabled in the BIOS
2. Configure the iDRAC adapter (utility available during the boot process)
3. Install the latest supported version of the Dell Open Manage Server Administrator
4. Enable the Hyper-V Role
5. Install the KB articles related to Hyper-V R2. Refer to the Hyper-V update list for Windows Server 2008 R2 for updated list of KB articles and hotfixes.
[http://technet.microsoft.com/en-us/library/ff394763\(W5.10\).aspx](http://technet.microsoft.com/en-us/library/ff394763(W5.10).aspx)
6. Configure network adapters
 - a. Configure NIC teaming using the adapters assigned for public and virtual machine network using BACS.
 - b. Set the static IP addresses adapters for cluster public, private, and live migration traffic. If there is no VLAN configured that sample VLAN IDs and the subnet can be utilized.
 - c. Configure a Hyper-V virtual switch on the teamed BACP virtual adapter.
 - i. *Note: The virtual switch must have the same name across all Hyper-V R2 servers in the cluster.*
 - d. Configure iSCSI Network adapters on the host. Assign static IP address.
 - e. For the NIC configured for iSCSI traffic, using BACS, set the Jumbo MTU to 9000 and enable flow control.
7. Enable the MPIO feature using the Add Feature wizard of Windows Server 2008 R2.

8. Install EqualLogic Multipath I/O DSM from the EqualLogic Host Integration Tool (HIT) kit 3.3.2 on the hosts. For more information related to HIT kit refer to the installation guide.
9. Start the iSCSI initiator service.

6. Configure SCE 2010 management server

1. Configure the iDRAC adapter (utility available during the boot process).
2. Install the latest supported version of the Dell Open Manage Server Administrator.
3. Enable the Hyper-V Manager feature for Remote Administration.
4. Install Failover Cluster Manager Feature for Remote Administration of Failover Cluster.
5. Configure the network adapter.
 - Set the static IP addresses.
6. Install and configure SCE 2010.
 - Install SCE 2010 management server on this server. If there is a SQL server already deployed in the infrastructure point to that during installation. Otherwise SQL express edition will be installed.
 - Refer Microsoft documentation for planning, deploying and managing using SCE 2010. <http://technet.microsoft.com/en-us/library/ff603627.aspx>

7. Configure LUNs and provide quorum to the Hyper-V R2 server

- Using PS series group manager, perform the following:
 - Create a PS volume.
 - Create the access list for giving access to Hyper-V R2 hosts.
- On the Hyper-V R2 servers, perform the following:
 - Using Microsoft iSCSI software initiator snap-in, login to the iSCSI target establishing multiple sessions for multipathing.
 - Force a rescan via Disk Management.
 - Confirm that multiple paths are seen to the volumes.

8. Create the Failover cluster using Hyper-V R2 servers as the nodes

Refer to Microsoft Technet link for step by step description on how to set up failover clustering and Hyper-V. [http://technet.microsoft.com/en-us/library/cc732181\(WS.10\).aspx#BKMK_Install](http://technet.microsoft.com/en-us/library/cc732181(WS.10).aspx#BKMK_Install)

Deploying failover Clustering and Hyper-V involves the following basic steps. On only one of the Hyper-V R2 servers, perform the following:

- Initialize/format disks.
 - Assign a drive letter to the quorum disk.
- Create the cluster.
 - Run through the “Validate a Configuration Wizard” and ensure that no unexpected errors are present.
 - Select all the other Hyper-V R2 servers during the creation process.

- Enable CSVs on shared storage volumes configured for virtual machines.
- Configure the cluster networks.
 - Ensure Virtual Machine network is not available for use by cluster.
 - Using the Failover Cluster Manager, set the priority of the networks for use by Live Migration [Live Migration (highest), Cluster Private/CSV, Cluster Public/Management].
 - Configure the cluster network metrics on the private networks (Live Migration and Cluster Private/CSV).
 - For live migration network, give the priority to the NIC connected to VLAN 50 assigned for Live migration network.

9. Add Hyper-V R2 hosts to the SCE 2010 management server

- From the SCE console, perform a discovery on the domain for detecting the Hyper-V R2 hosts and install agents.
- Designate the Hyper-V R2 node as host using SCE console. SCE will detect the failover cluster for deploying virtual machines.

10. Import Dell PRO pack management pack onto the SCE 2010 management server

- On the management server import and install Dell PRO pack management server 2.0.
- Ensure that PRO tips is enabled.
- Configure the discovery schedule.
- Refer to Dell and Management Systems Management Solutions link for downloads and user guide in installing management pack:
http://www.dell.com/content/topics/global.aspx/sitelets/solutions/management/microsoft_sms_essentials?c=us&cs=555&l=en&s=biz

10 References

- Dell Solution Guides for Microsoft Hyper-V
<http://support.dell.com/support/edocs/software/HyperV/en/index.htm>
- Microsoft Hyper-V Documentation: Hyper-V Getting started Guide
[http://technet.microsoft.com/en-us/library/cc732470\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc732470(WS.10).aspx)
- Hyper-V Role in Windows Server 2008 R2
[http://technet.microsoft.com/en-us/library/cc753637\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc753637(WS.10).aspx)
- Deploying Pools and Tiered Storage in a PS Series SAN
<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5239>
- Deploying Thin Provisioning in a PS Series SAN
<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5245>
- Dell and Microsoft Systems Management Solutions
http://www.dell.com/content/topics/global.aspx/sitelets/solutions/management/microsoft_sms_essentials?c=us&cs=555&l=en&s=biz

Appendix A: Reference Network Configuration

This appendix contains sample network configuration files for the PowerConnect 5424 switch used in the High Availability + Backup Configuration. As mentioned in the reference architecture diagram, there are four switches used in this configuration. Two switches are used for five networks namely Cluster Public network, Cluster private and CSV network, Live Migration network and Virtual Machine Network, and the other two switches are dedicated for iSCSI network.

Configuring PowerConnect 5424 switches for the High Availability configuration is similar except that the iSCSI network also shares the same two switches used for other networks separated by VLAN. Hence the VLAN for iSCSI and the recommended settings for EqualLogic iSCSI network need to be configured on the two switches.

PowerConnect 5424 Configuration File for Public, Virtual Network, Live Migration, Private and CSV:

```
interface port-channel 1
switchport mode trunk
exit
vlan database
vlan 20,30,50
exit
interface range ethernet g(3-10)
switchport access vlan 20
exit

interface port-channel 1
switchport trunk allowed vlan add 20
exit
interface range ethernet g(11-14)
switchport access vlan 30
exit
interface port-channel 1
switchport trunk allowed vlan add 30
exit
interface port-channel 1
switchport trunk allowed vlan add 50
exit
interface range ethernet g(21-24)
channel-group 1 mode on
exit
iscsi target port 860 address 0.0.0.0
iscsi target port 3260 address 0.0.0.0
interface vlan 1
ip address 172.168.11.101 255.255.255.0
exit
username admin password 5f4dcc3b5aa765d61d8327deb882cf99 level 15 encrypted
snmp-server community Dell_Network_Manager rw view DefaultSuper
```

PowerConnect 5424 Switch Configuration File for iSCSI network:

```
spanning-tree mode rstp
interface port-channel 1
spanning-tree portfast
exit
interface range ethernet g(2-20)
spanning-tree disable
exit
interface range ethernet g(21-24)
spanning-tree portfast
exit
interface range ethernet all
flowcontrol on
exit
port jumbo-frame
```



```
interface port-channel 1
switchport mode trunk
exit
interface range ethernet g(17-24)
channel-group 1 mode on
exit
iscsi target port 860 address 0.0.0.0
iscsi target port 3260 address 0.0.0.0
no iscsi enable
interface vlan 1
ip address 172.168.11.103 255.255.255.0
exit
username admin password 5f4dcc3b5aa765d61d8327deb882cf99 level 15 encrypted
snmp-server community Dell_Network_Manager rw view DefaultSuper
```