

# Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with PowerVault MD3200i Storage

---

A Solution Guide

By Shabana M

Dell Virtualization Solutions Engineering

[www.dell.com/virtualization/businessready](http://www.dell.com/virtualization/businessready)

Feedback: [solutionfeedback@dell.com](mailto:solutionfeedback@dell.com)

October 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, PowerVault 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 PowerVault MD3200i ..... 4
  - 4.3 Dell PowerConnect 5424 Series Switches..... 4
  - 4.4 PowerVault DL2100 Backup Solution ..... 5
  - 4.5 Solution Capabilities..... 5
- 5 Specifications ..... 7
- 6 Reference Architecture ..... 9
  - 6.1 Design Principles ..... 9
  - 6.2 Consolidation Solution..... 10
  - 6.3 High Availability Solution ..... 10
  - 6.4 High Availability + Backup Solution ..... 11
- 7 Network Architecture ..... 12
  - 7.1 Different Types of Network Traffic ..... 14
    - 7.1.1 Virtual Machine Network/Cluster Public Network ..... 15
    - 7.1.2 Private / CSV and Live Migration Networks ..... 16
  - 7.2 iSCSI SAN Architecture ..... 16
  - 7.3 Storage Connectivity ..... 16
  - 7.4 PowerVault MD3200i Network Requirements ..... 17
  - 7.5 Storage Configuration ..... 17
- 8 Management ..... 18
  - 8.1 Systems Management ..... 18
  - 8.2 Storage Management ..... 18
  - 8.3 Virtualization Management - Microsoft Systems Center Essentials 2010 ..... 18
- 9 Deployment..... 19
- 10 References ..... 23
- Appendix: Reference Network Configuration..... 24
- PowerConnect 5424 Configuration File for *Public, Virtual Network, Live Migration, Private and CSV*: . 25
- PowerConnect 5424 Switch Configuration File for the iSCSI Network: ..... 25

## 2 Introduction

Business-Ready Configurations for Virtualization are a family of reference architectures offered by Dell that are designed to aid with the ordering, deployment, and maintenance of a virtualization infrastructure. The reference architectures described in this document are targeted at small and medium business virtualization environments, although others may also benefit. The architecture discussed in this document includes Dell™ PowerEdge™ R510 servers, Dell PowerVault MD3200i series storage, Dell PowerConnect™ switches, and Microsoft® Hyper-V technology. Based on extensive design and engineering work, this solution can quickly and confidently be deployed into production environments. This infrastructure is optimally configured to run virtualized workloads and is designed to provide redundancy with no single point of failure. This document covers network architecture, storage configuration and the best practices necessary for deploying and configuring a Microsoft Hyper-V installation.

## 3 Audience and Scope

The intended audience for this white paper is IT administrators, IT managers, and channel partners who are planning to deploy or resell Microsoft. This guide can be used to plan, scope, and procure the required components to set up a robust virtualized infrastructure based on Microsoft Hyper-V. It is assumed that the reader has a basic understanding of server virtualization (Hyper-V preferred), iSCSI storage, and networking concepts.

## 4 Overview

This section provides an overview of the hardware components used for the Business-Ready Configuration followed by the key capabilities offered. For networking, this architecture leverages the benefits of PowerConnect 5424 for both iSCSI and other VM network traffic. Dell's PowerVault DL2100 disk-to-disk backup solution with CommVault® provides backup and recovery services. Microsoft System Center Essentials 2010 (here after referred to as SCE 2010) provides an integrated solution for managing physical and Hyper-V R2 virtual infrastructures.

Three reference architectures are discussed in this white paper - an entry level configuration for **Consolidation**, a **High Availability** configuration, and a **High Availability + Backup** configuration. Hardware included in each configuration is shown in Table 1.

	Consolidation	High Availability	High Availability + Backup
<b>Virtualization Hosts (Hyper-V R2)</b>	1 x PowerEdge R510	2 x PowerEdge R510	4 x PowerEdge R510
<b>External Storage</b>	None	1 x MD3200i (8 x 500 GB NLSAS drives)	1 x MD3200i (12 x 500GB NLSAS drives)
<b>Network Infrastructure</b>	None	2 x PowerConnect 5424	4 x PowerConnect 5424
<b>Backup Device</b>	None	None	DL2100 (CommVault)
<b>Virtualization Management</b>	Hyper-V Manager	1 x PowerEdge R410 with SCE 2010	1 x PowerEdge R410 with SCE 2010
<b>Key features</b>	<ul style="list-style-type: none"> <li>Ideal for consolidation or deployment of up to 15* Virtual Machines</li> </ul>	<ul style="list-style-type: none"> <li>Ideal for consolidation or deployment of up to 30* Virtual Machines</li> <li>High Availability of Virtual Machines using Failover Clustering</li> <li>iSCSI based shared storage array (up to 4 TB of storage capacity)</li> <li>Single console management using SCE 2010</li> </ul>	<ul style="list-style-type: none"> <li>Ideal for consolidation or deployment of up to 60* Virtual Machines</li> <li>High Availability of Virtual Machines using Failover Clustering</li> <li>iSCSI based shared storage array (up to 6 TB of storage capacity)</li> <li>Can scale up by adding MD1220 or MD1200 (maximum of 96 hard drives) for additional capacity</li> <li>Single console management using SCE 2010</li> <li>Disk-to-Disk backup with optional disk-to-tape backup using DL2100</li> </ul>

Table 1: Three pre-configured Solution Hardware Overviews

\* Actual results may vary.

## 4.1 Dell PowerEdge R510 Server

The PowerEdge R510 follows the 11th Generation PowerEdge portfolio. The 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 or a twelve hard drive chassis. In the eight hard drive chassis option, the PowerEdge R510 provides up to 16TB 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 *integrated Dell Remote Access Controller (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 such as system deployment, system

updates, hardware configuration, and diagnostics from a single intuitive interface called the Unified Server Configurator (USC).

## 4.2 Dell PowerVault MD3200i

The Dell PowerVault MD3200i is the latest addition to the PowerVault MD series of arrays for iSCSI SANs. It is designed for deployment of entry-level storage consolidation in a virtualized environment with increased performance, scalability, and high availability. The MD3200i consists of two internal RAID controllers, each consisting of four 1 Gb Ethernet data ports for iSCSI communication and one 1 Gb Ethernet port for out of band management, 2GB of battery-backed cache per controller, increased storage capacity, and connectivity for an increased number of hosts. This improvement from the previous generation of MD series of arrays provides up to twice the performance. Dell PowerVault MD 32xx storage arrays provide the following capabilities:

**Scalability:** Up to 32 hosts can be connected to a single MD32xxi series iSCSI array. Storage capacity can be expanded up to a total of 96 hard drives by connecting additional PowerVault MD1200 and/or MD1220 enclosures. The MD3200i supports adding both MD1200 and MD1220 enclosures and thus allows a mix and match of drive types allowing an optimum tiered data environment.

**Reliability:** Dell PowerVault MD3200i arrays have hot-swappable redundant components, a choice of RAID types, and hot-spare disks. An MD3200i array consists of two RAID controllers in active-active mode, each providing I/O access to the virtual disks it owns. When a controller failure happens, the virtual disks owned by that controller are transferred to the other surviving controller which takes over the I/O requests from hosts transparently.

**Management Efficiency:** PowerVault MD3200i series arrays are managed by the enhanced Module Disk Storage Manager software (MDSM). MDSM is wizard-based management utility that simplifies array configuration. This software also includes an enterprise management window that monitors multiple systems (including previous generation MD3000i arrays) through a single interface. MDSM provides two types of array management - in band management and out of band management. Dell recommends using out-of-band management for managing the array.

## 4.3 Dell PowerConnect 5424 Series Switches

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 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 DL2100 Backup Solution

The Dell PowerVault DL2100 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:** The PowerVault DL2100 is factory installed with CommVault Simpana Software. The DL2100 provides configuration and management wizards to help users configure and manage the appliance in a quick and easy way.

**Integrated tape support:** The PowerVault DL2100 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 or archival storage.

**Built-in deduplication:** The PowerVault DL2100 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 DL2100 Powered by CommVault is mentioned in this document, the PowerVault DL2100 Powered by Symantec is also available.

## 4.5 Solution Capabilities

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

### Virtual Machine Live Migration

The **High Availability (HA)** and **High Availability (HA) + Backup** configurations support live migration of Virtual Machines (VM) between two or more Hyper-V R2 hosts. Live migration is a 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 migrate those Virtual Machines independently amongst 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 - Microsoft System Center Essentials 2010

- Provides a unified solution with a single console for managing physical and virtual environments, monitoring, and reporting alerts. Virtual Machine Manager 2008 R2 technology is included in SCE2010 and provides a single console and management solution for managing both physical and virtual servers.
- 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, 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.
- Template-based Virtual Machine creation allows administrators to build templates for Virtual Machines and store them in the SCE 2010 library. This speeds deployment of Virtual Machines in production environments.
- Physical server to Virtual Machine conversion
- Live migration of Virtual Machines.
- Resource utilization monitor - determines the usage of the Hyper-V R2 host servers and the running Virtual Machines to enable load balancing.



## 5 Specifications

The following table provides the detailed hardware specifications for the three reference architectures provided in this document. Solution ID's refer to the online configuration available at [www.dell.com/businessready](http://www.dell.com/businessready) (this provides a comprehensive Bill of Materials as well as a budgetary price).

	Consolidation Solution	High Availability Solution	High Availability + Backup Solution
<b>Hardware</b>			
Solution ID	<a href="#">1209506.1</a>	<a href="#">1209559.1</a>	<a href="#">1209569.1</a>
Hyper-V R2 Server	1 x R510	2 X R510	4 x R510
Storage Device	Local Storage (8 Internal HDD)	(1) MD3200i	(1) MD3200i (Can add MD1200 / MD1220 for additional capacity)
Infrastructure Server and Management Server	N/A	(1) x R410	(2) x R410
<b>Network Configuration</b>			
Network Configuration	N/A	2 x PowerConnect 5424 Network Switches	4 x PowerConnect 5424 Network Switches
Back up Device	N/A	N/A	DL2100 Disk Based Backup Solution

## Hyper-V R2 Server Configuration

<b>Rack Server Model</b>	R510	R510	R510
<b>Processor</b>	(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
<b>Memory</b>	48 GB (6 x 8GB, DDR3)	48 GB (6 x 8 GB, DDR3)	48 GB (6 x 8 GB, DDR3)
<b>Onboard Network Adapter</b>	(1) Integrated Broadcom 5716 Dual port GbE adapter	(1) Integrated Broadcom 5716 Dual port GbE adapter	(1) Integrated Broadcom 5716 Dual port GbE adapter
<b>Add-in Controllers</b>		(2) Broadcom 5709 Dual Port GbE PCI/e Card	(2) Broadcom 5709 Dual Port GbE PCI/e card
<b>Internal Storage Controller</b>	PERC 6/i	PERC 6/i	PERC 6/i
<b>Internal Hard Drive</b>	(3) X 500GB, SATA, RAID5 Up to (8) X 500GB, SATA	(2) x 146 GB in RAID1	(2) x 146GB in RAID1
<b>Hypervisor</b>	Microsoft Hyper-V R2 Enterprise Edition	Microsoft Hyper-V R2 Enterprise Edition	Microsoft Hyper-V R2 Enterprise Edition

## External Storage Configuration

<b>Storage Device</b>	None	(1) MD3200i	(1) MD3200i
<b>Drives</b>	None	8 X 500GB, NLSAS	12 X 500GB, NLSAS
<b>Storage Capacity</b>	None	4 Terabytes per device	6 Terabytes 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 6.3	Open Manage 6.3	Open Manage 6.3

Table 2: Specification for Reference Architecture Configuration

## 6 Reference Architecture

### 6.1 Design Principles

The following design principles of Optimal Configuration, High Availability, Live Migration, Redundancy and Virtualization Management were used and carefully considered while implanting these reference configurations.

**Choosing the Optimal Configuration:** The solution components selected are specifically designed for the requirements of a typical SMB environment. SCE 2010 was chosen for the management software due to its extensive functionality and ease of implementation.

**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.

**Isolated and Redundant Network Architectures:** The network configuration used provides isolation for different networks in the virtual infrastructure and avoids a single point of failure.

**Redundancy with No Single Point-of-Failure:** For the HA and HA + Backup Solutions, redundancy is incorporated in every aspect of networking and storage to avoid any single point of failure.

**Virtualization Management Capabilities Suited for SMB:** SCE 2010 supports management of up to 50 server operating system 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 for providing 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 eight drive chassis is configured with three 500GB SATA hard drives in a RAID 5 configuration. Additional drives may be added to increase the internal storage. The built-in Hyper-V R2 Manager can be used to manage this single server environment.

## 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. Both servers are connected to a single PowerVault MD3200i storage array in an iSCSI SAN. The network infrastructure consists of two PowerConnect 5424 switches and a PowerEdge R410 running as the SCE 2010 management server for deploying the virtual infrastructure. Other infrastructure services required include Active Directory, domain name resolution (DNS), and DHCP server. For additional capacity, MD1200 and/or MD1220 enclosures can be connected to this configuration (a maximum of 96 hard drives are supported). Refer to the PowerVault MD3200i documentation for best practices in adding an MD12xx storage array for additional capacity. In Figure 1, cabling connections are only shown on one R510 server. Cabling connections for the other R510 server are similar.




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 shown below. The recommended solution consists of four PowerEdge R510 servers running Hyper-V R2 configured in a failover cluster directly attached to a PowerVault MD3200i storage array. 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). For additional capacity, MD1200 and/or MD1220 enclosures can be connected to this configuration (a maximum of 96 hard drives are supported). Refer to the PowerVault MD3200i documentation for best practices in adding an MD12xx storage array for additional capacity.

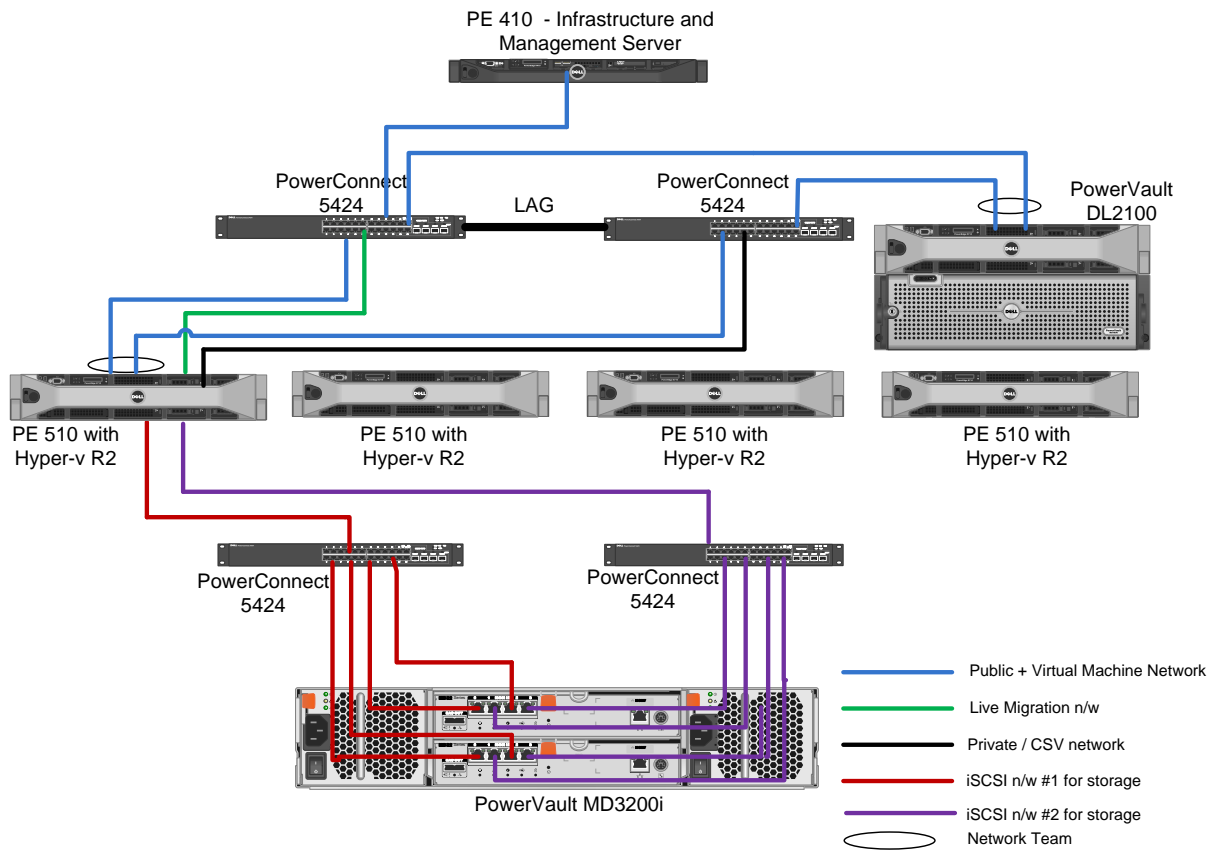


Figure 2: Reference Architecture for High Availability + Backup Solution

## 7 Network Architecture


The Hyper-V R2 virtual infrastructure consists of six network traffic types: 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 of the physical network, the PowerConnect switches, and the virtual. Gigabit Ethernet is used for all network traffic types.

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 & NIC 4*: First and second ports on the first dual-port Broadcom NetXtreme II 5709 Gigabit Ethernet controller.
- *NIC 5 & 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, use either two or four Dell PowerConnect 5424 switches. The PowerConnect switches are interlinked using a 4 port LAG (Link Aggregation Group) which is used by all the networks with the exception of iSCSI traffic (see Figures 3 and 4). The LAG provides optimal bandwidth for the traffic between the switches and failover capabilities should a link fail. The PowerConnect 5424 switch supports a maximum of eight Gigabit ports combined into a LAG, providing an aggregated bandwidth of 8 Gbps. The ports participating in the LAG are configured as trunk ports and are configured as members for VLANs assigned for VM/Public network, Private/CSV network, and Live Migration network.

In the **High Availability** configuration, all network traffic 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 used for the Virtual Machine, Cluster Public, Private/CSV, and Live Migration networks separated by VLANs. The second two switches are dedicated to iSCSI SAN traffic.



**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

In the Microsoft best practices guide, network traffic is broken down into separate networks, and each of these network types is isolated using VLANs. Table 3 gives a sample VLAN configuration that can be used for different types of network configurations used in the HA and HA + Backup configurations. This section will discuss more of the details for configuring different types of network traffic.

**iSCSI Network #1 and #2 (VLAN ID 10, 11):** Used for iSCSI communication between the Hyper-V R2 hosts and the PowerVault MD3200i storage array in the **High Availability** configuration. In the **High Availability + Backup** configuration, two network switches are used for dedicated iSCSI communication. In this case no VLANs are needed to isolate this traffic. In both of the configurations, the iSCSI network includes two subnets to enable the load balancing capabilities of the MD3200i.

**Cluster Public Network or Management (VLAN ID 20):** Used by the cluster for public or external communication. This network is also used by the SCE 2010 management server to connect to the Hyper-V R2 hosts in the cluster.

**Virtual Machine Network (VLAN ID 20):** Used for Virtual Machine connectivity from the Hyper-V R2 host servers 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 cluster public and Virtual Machine network traffic.

**Cluster Private Network (VLAN ID 30):** Used by the cluster for intra-node communication.

**Network for CSV (VLAN ID 30):** 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):** Used for the live migration traffic between Hyper-V R2 hosts in the failover cluster.



Table 3 below shows sample VLANs that can be configured in the HA and HA + Backup configuration for each of the six network types discussed above.

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
iSCSI Network #1	10	172.10.1.X
iSCSI Network #2	11	172.11.1.X

Table 3: Sample VLAN for each Network Types

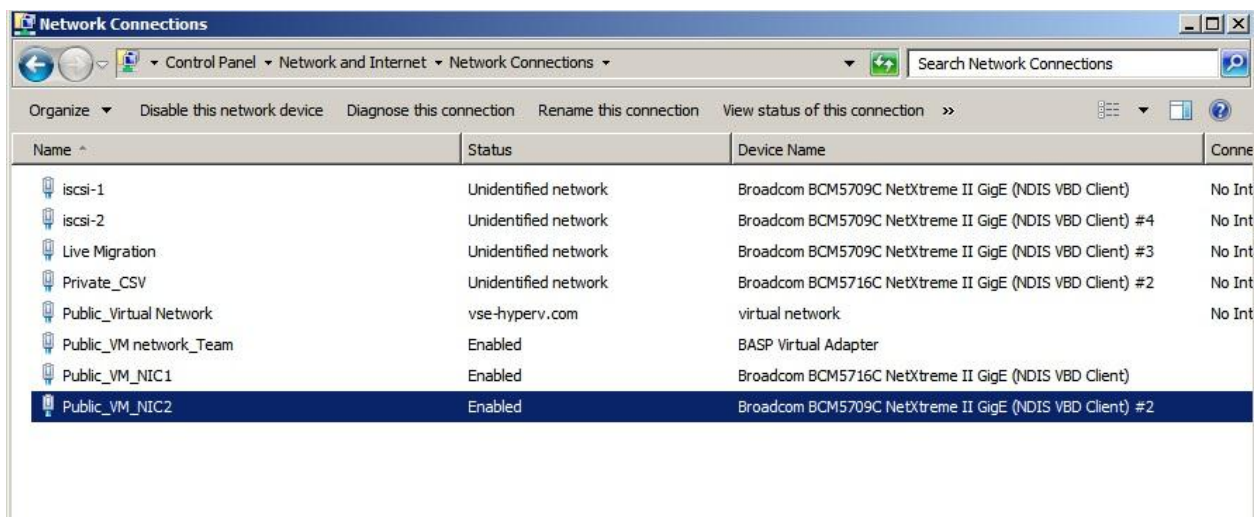


Figure 5: The Network Connections Manager - Networks Configured in the Hyper-V R2 Hosts.

### 7.1.1 Virtual Machine Network/Cluster Public Network

The Virtual Machine Network and the Cluster Public Network are configured using two Broadcom network ports. The Broadcom Advanced Configuration Suite (BACS) is used to configure Smart Load Balancing (SLB) with failover capabilities. One network port is configured as the active member of the team, while the other is a stand-by network port to avoid a single point of failure. These NICs are physically connected to the ports configured for VLAN ID 20. This network also supports the SCE 2010 management server for discovering, monitoring, and managing the physical and virtual systems. In Figure 5, the network listed as the “Public\_Virtual Network” is the virtual network created in external mode for the Cluster Public Network and the Virtual Machine Network. It uses the Broadcom Advanced Server Program (BASP, installed with BACS) Virtual teamed adapter for the network connection. In the reference configuration, LOM 1 and NIC 5 are used to create the network team.

### 7.1.2 Private / CSV and Live Migration Networks

The *Cluster Private/CSV Network* port (LOM2) and *Live Migration Network* port (NIC3) can be configured for failover. Both networks are connected to different switches for redundancy. NICs for the *Cluster Private/CSV Network* are connected to the switch ports configured for VLAN 30 and the *Live Migration Network* NICs are connected to switch ports configured for VLAN 50.

## 7.2 iSCSI SAN Architecture

This section discusses the best practices for configuring the iSCSI SAN using the Microsoft software iSCSI initiator provided in Windows Server 2008 R2 and the PowerVault MD3200i iSCSI target.

**HA Configuration:** The two 1Gb Ethernet network ports for iSCSI communication on the Hyper-V R2 hosts are physically connected to the same PowerConnect switches used for all the traffic types, but are separated by VLANs. For iSCSI traffic, VLAN 10 is configured on one switch and VLAN 11 is configured on the other switch. Hence, all six of the different traffic types share the same physical switches, but they are isolated using VLANs. The two NIC ports from each of the Hyper-V R2 hosts are connected to the switch ports configured for VLAN 10 and 11 respectively. For best performance, all data ports from the MD3200i should be used for the iSCSI traffic. Two ports from each of the MD3200i controllers should be connected to the PowerConnect switches. Hence two of the four ports in one controller will be configured in one VLAN while the other two ports are attached to the other VLAN. Note that the LAG between the two switches is configured in trunk mode to accept packets from all VLANs except VLAN 10 and VLAN11 (iSCSI networks) in this configuration.

**HA + Backup Configuration:** The HA + Backup configuration uses two additional iSCSI-dedicated PowerConnect 5424 switches. Since the traffic is isolated physically, there is no need to configure VLANs for iSCSI traffic in this configuration.

The iSCSI network uses two subnets for storage connectivity to enable the load balancing capabilities of the MD3200i array. On the Hyper-V R2 hosts, the Microsoft software iSCSI initiator is configured to establish multiple sessions: one session for each initiator IP address and one for each target IP address on the controller so that all of the hosts can see four paths to each volume. Multiple paths are managed using the multipath driver that is provided as part of MDSM. For additional information on configuring a PowerVault MD3200i Array, refer to the PowerVault MD3200i documentation at <http://support.dell.com/support/edocs/systems/md3200i/>.

## 7.3 Storage Connectivity

Each controller in an MD3200i array has four 1Gb Ethernet ports for iSCSI traffic and one 1Gb Ethernet port for out-of-band management. The reference configuration uses out-of-band management to isolate management traffic from the iSCSI traffic. For better performance, all the four ports in the controller are used for iSCSI traffic. These ports are connected to the two switches - two from each controller is connected to each of the PowerConnect 5424 switches to provide load balancing between the controllers.

## 7.4 PowerVault MD3200i Network Requirements

In addition to the guidelines discussed previously, there are certain IP SAN networking specific recommendations for connecting MD3200i arrays to the network.

- Spanning-Tree protocol (STP) should not be used on switch ports that connect end nodes (iSCSI initiators or array network interfaces). If STP or Rapid STP (preferable to STP) are necessary, enable the port settings (FastLink or Port Fast) available on some switches that let 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 MD3200i will auto-configure to the PowerConnect switch when Flow Control is turned on.
- Disable Unicast Storm Control on each switch that handles iSCSI traffic if the switch provides this feature. The use of broadcast and multicast storm control is encouraged on switches that support that functionality.
- Enable Jumbo Frames on the iSCSI data ports on each RAID controller in MD3200i, network switches, and the NICs used for iSCSI traffic.

For more information on configuring an IP SAN network with the PowerVault™ MD3200i, please refer to: <http://www.dell.com/downloads/global/products/pvaul/en/ip-san-best-practices-en.pdf>

## 7.5 Storage Configuration

The section provides recommendations for configuring the Storage Array.

**Load Balancing:** Multipathing is a technique that allows more than one physical path to be used to transfer data between a host and an external storage device. The MDSM contains the Device Specific Module (DSM) for Microsoft Multipath I/O solution (MPIO). MPIO provides for load balancing the I/O between the four ports in the RAID controller and also manages the multiple paths between the hosts and storage for availability. The load balancing policies supported by MD3200i include *round robin with subset*, *least queue depth*, and *weighted paths*. The default selection for each volume in the MD3200i is *least queue depth*, but this can be changed to *round robin* to balance iSCSI traffic over multiple physical NIC ports and improve performance.

The Dell PowerVault MD3200i storage system supports active/active controllers, with each controller being able to process I/O simultaneously. Each virtual disk is owned by a single controller and all I/O to the virtual disk is only possible through the owning controller. To take advantage of both the controllers for I/O access, distribute virtual disks among the controllers to balance I/O access.

**Array RAID Configuration:** The storage array RAID configuration is highly dependent on the workload in your virtual environment. Dell PowerVault MD32xx Series arrays support four RAID types; RAID 5, RAID 6, RAID 10, and RAID 50. RAID 10 is the Dell recommended configuration for the best throughput and performance. In general, RAID 10 provides the best performance, at the expense of storage

capacity. RAID 50 generally provides more usable storage, but has lower performance than RAID 10 in random I/O situations and requires additional overhead in the case of a drive failure. RAID 5 provides the most storage capacity at the expense of slightly lower performance and availability. RAID 6 provides better protection than RAID 5 with minimal disk cost but is impacted by the write performance overhead of double parity calculations.

## **8 Management**

This section covers the different aspects of managing and deploying the reference architecture.

### **8.1 Systems Management**

Dell OpenManage Server Administrator manages the Hyper-V R2 servers in this reference architecture. Dell Open Manage Server Administrator is a web based tool that can be used for managing and monitoring the server hardware, system health, and performance. For information on Dell OpenManage and its capabilities, see [www.dell.com/openmanage](http://www.dell.com/openmanage).

The Dell Unified Server Configurator (USC) which ships with all 11G PowerEdge servers enables easy deployment and configuration of the host servers. The USC can be used for BIOS, firmware, RAID configuration, and driver updates.

### **8.2 Storage Management**

PowerVault MD3200i series arrays are managed by the enhanced MDSM software. Two different management paths are available: in-band array management and out-of-band array management. For the reference configurations discussed in this document, out-of-band management is preferred. Additionally, the enterprise management window feature monitors multiple systems including previous generation MD3000i arrays through a single interface. MDSM detects and alerts problems, then launches an automatic Recovery Guru to help troubleshoot and resolve any issues.

In addition to MDSM, SCE 2010 can monitor the Dell PowerVault array by installing a Management Pack on the PowerEdge R410 management server. Dell PowerVault Management Pack will collect alerts and information and present it to the SCE 2010 console. This enables a single console to monitor the server, storage, and Virtual Machines in their environment.

### **8.3 Virtualization Management - Microsoft Systems Center Essentials 2010**

Windows Server 2008 R2 Enterprise natively ships with the ability to locally manage Hyper-V R2 servers with limited management functionality (i.e. power cycle, migrate, etc.) using the Hyper-V Manager. The Microsoft Failover Cluster Manager snap-in is another natively supported feature that enables management of the Virtual Machines with limited functionality.

Dell OEM Microsoft System Center Essentials 2010 (SCE 2010) helps administrators manage Virtual Machines, as well as the physical systems, using the same console. SCE 2010 includes 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), suitable for small and medium sized businesses. 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, which assists administrators in monitoring the Hyper-V R2 systems health and recommends implementation tips.

SCE 2010 can be deployed on an existing infrastructure server that hosts DNS, DHCP and Active Directory functions if already present. Optionally, it can be deployed on a separate R410 server if those services have not been implemented.

## 9 Deployment

This section provides guidelines for the initial deployment and configuration of the proposed solution.

Managing a PowerVault MD3200i array, PowerConnect 5424 switches, 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.
- Interlink the network switches.

### 2. Configure PowerConnect Switches

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

### 3. Perform Initial Configuration of the MD3200i Array

- Initialize the MD3200i array by assigning a management IP address to each of the controllers for out-of-band management. Connect the serial cable and assign management IP addresses for both of the controllers.
- Connect the out-of-band management port on each array RAID controller to the existing infrastructure for hardware management.
- Install MDSM software from the PowerVault MD3200i Resource CD on a management station. Using MDSM, discover the storage array.
- Configure the four 1G Ethernet ports for iSCSI using MDSM. Assign IP addresses for the iSCSI ports on each controller. As mentioned earlier, two ports on each controller will be in one subnet (iSCSI network #1) and other two ports will be in the other subnet (iSCSI network #2).
- Enable Jumbo frames on both the array RAID controllers by setting the MTU to 9000 bytes on all of the iSCSI host ports using MDSM.
- For information on configuring storage, see the Dell PowerVault documentation from <http://support.dell.com/support/edocs/systems/md3200i/en/index.htm>

#### 4. Configure Infrastructure Server

- Deploy Microsoft Windows Server 2008 R2 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](http://support.dell.com/support/edocs/software/win2008/WS08_R2/en/IIG_HyperV/IIG_HypV.pdf). This guide provides updates on the latest known issues.

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 (a utility is available during the boot process).
3. Install the latest supported version of the Dell OpenManage 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 an updated list of KB articles and hot fixes at [http://technet.microsoft.com/en-us/library/ff394763\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ff394763(WS.10).aspx)
6. Configure the network adapters.
  - a. Configure NIC teaming using the adapters assigned to the *Virtual Machine Network* and the *Cluster Public Network* using BACS.
  - b. Set the static IP addresses for the NICs devoted to the *Cluster Public Network*, the *Private/CSV Network*, and the *Live Migration Network*. If there is no VLAN configured, the sample VLAN IDs and the subnets provided in this document can be used.
  - c. Configure a Hyper-V virtual switch on the teamed BASP virtual adapter.

*Note: The virtual switch must have the same name across all Hyper-V R2 servers in the cluster.*
  - d. Configure the iSCSI network adapters on the host. Assign static IP addresses.
  - e. Use BACs to set up the iSCSI NICs - 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 MDSM software on the hosts from the MD3200i Resource CD. MDSM contains the DSM for MPIO. Choose the hosts option while installing MDSM so that the DSM for MD3200i gets installed on the Hyper-V R2 hosts.
9. Start the Microsoft iSCSI Software 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 the Failover Cluster Manager Feature for Remote Administration.
5. Configure the network adapter.
  - Set the static IP addresses.
6. Install and configure SCE 2010.
  - Install SCE 2010 management server. If there is a SQL server already deployed in the infrastructure, point SCE 2010 to that database during installation. Otherwise SQL Server Express Edition will be installed.
  - Refer to Microsoft documentation for planning, deploying, and managing SCE 2010 at <http://technet.microsoft.com/en-us/library/ff603627.aspx>

## 7. Configure the iSCSI Initiator on the Hosts

- Perform the following on the Hyper-V R2 host servers:
  - Using the Microsoft iSCSI software initiator snap-in, add the target portals using the Discover option and login to the iSCSI target establishing multiple sessions for multi-pathing.
  - Using the two iSCSI IP addresses on the hosts, establish one session for each of the iSCSI ports on the MD3200i controller. Together there will be eight sessions - four per initiator IP address to the target MD3200i array.

## 8. Configure LUNs and Provide Quorum to the Hyper-V R2 Servers

- Using MDSM, perform the following:
  - Configure host access on the array by detecting the iSCSI Qualified Name (IQN) of each of the Hyper-V R2 hosts.
  - Create a Disk Group with RAID 10.
  - Create virtual disks in each of the disk groups.
  - For better performance, PowerVault storage documentation recommends configuring up to four virtual disks per disk group.
  - Create a Host Group consisting of the Hyper-V R2 hosts and virtual disks configured for Virtual Machine files and quorum.
- On the Hyper-V R2 hosts, perform the following:
  - Launch the Disk Management mmc and force a rescan
  - Confirm that multiple paths are seen to the volumes.

## 9. Create the Failover Cluster Using Hyper-V R2 Servers as the Nodes

Refer to the Microsoft TechNet link for a 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)

Perform the following procedure for deploying failover Clustering and Hyper-V on only one of the R710 servers:

- 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 the remainder of the 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 the 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 the *Live Migration* network, give priority to the NIC connected to VLAN 50.

#### 10. Add Hyper-V R2 Hosts to the SCE 2010 Management Server

- From the SCE console, perform a discovery on the domain to detect the Hyper-V R2 hosts and install the agents.
- Designate the Hyper-V R2 node as host using SCE console. SCE will detect the failover cluster for deploying Virtual Machines.
- For deploying Virtual Machines on the Hyper-V R2 Failover cluster using SCE 2010, refer to the Microsoft documentation found here: <http://technet.microsoft.com/en-us/systemcenter/essentials/ff631054.aspx>

#### 11. Import Dell PRO Pack Management Software onto the SCE 2010 Management Server

- On the management server, import and install Dell PRO Pack Management Server 2.0.
- Enable PRO Tips.
- Configure the discovery schedule.
- Refer to the Dell Systems Management link below for downloads and the user guide: [http://www.dell.com/content/topics/global.aspx/sitelets/solutions/management/microsoft\\_sms\\_essentials?c=us&cs=555&l=en&s=biz](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\(W5.10\).aspx](http://technet.microsoft.com/en-us/library/cc732470(W5.10).aspx)
- 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](http://www.dell.com/content/topics/global.aspx/sitelets/solutions/management/microsoft_sms_essentials?c=us&cs=555&l=en&s=biz)
- Planning for Hyper-V: PowerVault™ MD3200 and MD3200i Series of Arrays  
<http://www.dell.com/downloads/global/products/pvail/en/powervault-md3200-series-hyper-v-planning.pdf>
- PowerVault MD3200 and MD3200i Series of Storage Arrays: Implementing Hyper-V  
<http://www.dell.com/downloads/global/products/pvail/en/powervault-md3200-series-hyper-v-implementation.pdf>
- PowerVault MD3200 and MD3200i: Array Tuning Best Practices  
<http://www.dell.com/downloads/global/products/pvail/en/powervault-md3200i-performance-tuning-white-paper.pdf>
- Installation and Configuration for PowerVault DL2100  
[http://documentation.commvault.com/dell/release\\_8\\_0\\_0/books\\_online\\_1/english\\_us/prod\\_info/dell\\_appliance.htm](http://documentation.commvault.com/dell/release_8_0_0/books_online_1/english_us/prod_info/dell_appliance.htm)

## Appendix: Reference Network Configuration

This appendix contains sample network configuration files for the PowerConnect 5424 switches used in the **High Availability + Backup** Configuration. Four switches are used in this configuration. Two of the switches are used for five networks: *Cluster Public*, *Cluster Private*, *CSV*, *Live Migration*, and *Virtual Machine*. The remaining two switches are dedicated to the 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 the other networks, separated by a VLAN. The VLAN for iSCSI and the recommended settings for PowerVault 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 the 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
iscsi target port 860 address 0.0.0.0
iscsi target port 3260 address 0.0.0.0
interface vlan 1
ip address 172.168.11.103 255.255.255.0
exit
username admin password 5f4dcc3b5aa765d61d8327deb882cf99 level 15 encrypted
snmp-server community Dell_Network_Manager rw view DefaultSuper
```