

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

Feedback: 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 DL2200 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 <i>Public, Virtual Network, Live Migration, Private and CSV</i> : .	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. 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 described in this document are targeted at small and medium business virtualization needs, although others may also benefit from these architectures. These architectures include Dell™ PowerEdge™ R510 servers, Dell PowerVault MD3200i series storage, Dell PowerConnect™ switches, and Microsoft® Hyper-V technology. Based on extensive design and engineering work, customers can quickly and confidently deploy this proven engineering architecture 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 architecture, storage configuration and the best practices necessary for deploying and configuring the solution.

3 Audience and Scope

The intended audiences for this white paper are IT administrators, IT managers, and channel partners who are planning to deploy or resell Microsoft Virtualization themselves or for their customers. The reference architecture provides an overview of the recommended servers, storage, software, and services. It 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 PowerVault MD3200i Series iSCSI storage 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

As mentioned previously, the reference architecture discussed in this white paper is centered on Microsoft's latest virtualization platform, Microsoft Windows Server 2008 R2 with Hyper-V (referred to here on as Hyper-V R2), PowerEdge R510 servers, and Dell PowerVault MD3200i storage. The architecture leverages the benefits of PowerConnect 5424 for both iSCSI and other network traffics, Dell PowerVault DL2200 disk-to-disk backup solution with CommVault®, and Dell OEM Microsoft System Center Essentials 2010 (here after 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 white paper - an entry level configuration for Consolidation, High Availability configuration, and High Availability + Backup configuration. A high-level overview of the hardware included in each configuration is shown 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 MD3200i (8 x 500 GB NLSAS drives)	Up to 1 x MD3200i (12 x 500GB NLSAS drives)
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 up to 15* Virtual Machines 	<ul style="list-style-type: none"> Ideal for consolidation or deployment of up to 15* Virtual Machines High Availability of Virtual Machines using Failover Clustering iSCSI based shared storage array (up to 4 TB of storage capacity) Single console management using SCE 2010 	<ul style="list-style-type: none"> Ideal for consolidation or deployment of up to 45* Virtual Machines High Availability of Virtual Machines using Failover Clustering iSCSI based shared storage array (up to 6 TB of storage capacity) Can scale up by adding MD1220 or MD1200 (maximum of 96 hard drives) for additional capacity 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 Overviews

* Note: Virtual Machine count mentioned in the above table is based on Virtual Machines with an average of 3GB of RAM and 133GB of shared storage space. 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 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

Dell PowerVault MD3200i is the new addition to PowerVault MD series of arrays for iSCSI SANs. It is designed for deployment of entry-level storage consolidation in a virtualized environment with better performance, scalability, and high availability. MD3200i consists of two 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 for iSCSI SAN will provide twice the performance over the previous Dell PowerVault MD3000 series of arrays. In addition to the new features described above, Dell PowerVault MD 32xx storage arrays provide the following capabilities:

Scalability: Up to 32 hosts can be connected to a single MD32xxi series array in iSCSI SAN. Storage capacity can be expanded up to a total of 96 hard drives by connecting additional PowerVault MD1200 and/or MD1220 enclosures. 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 and it 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 software which provides ease in configuring the array. Besides the array management facility, another feature 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. It is recommended to use out-of-band management for managing the array.

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 mentioned in this document, the PowerVault DL2200 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 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. Virtual Machine Manager 2008 R2 technology is added to the product, so that Essentials 2010 provides a single console and management solution for managing both your 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.
 - SCE 2010 has many features which help users in deploying, managing, and monitoring 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. Other features offered by SCE2010 are physical server to Virtual Machine conversion and live migration of Virtual Machines.
 - 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 in 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	1209506.1	1209559.1	1209569.1
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
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	48 GB (6 x 8GB, DDR3)	48 GB (6 x 8 GB, DDR3)	48 GB (6 x 8 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) MD3200i	(1) MD3200i
Drives	None	8 X 500GB, NLSAS	12 X 500GB, NLSAS
Storage Capacity	None	4 Terabyte per device	6 Terabyte 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

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

6.1 Design Principles

The following design principles of Optimal Configuration, High Availability and Live Migration, Redundancy and Virtualization Management were used and carefully considered while designing the 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 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. **Redundancy with no single point-of-failure:** For HA and HA + Backup Solution, redundancy is incorporated in every aspect of the solution, including networking and storage to avoid any single point of failure.
5. **Virtualization Management Capabilities suited for SMB:** Management of the reference configuration discussed in this solution guide is done using SCE 2010, which is best suited for the SMB customer. 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 used with three 500GB SATA hard drives in RAID 5 configuration. To increase the internal storage, it can scale up to eight 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 PowerVault MD3200i 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 for deploying 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.

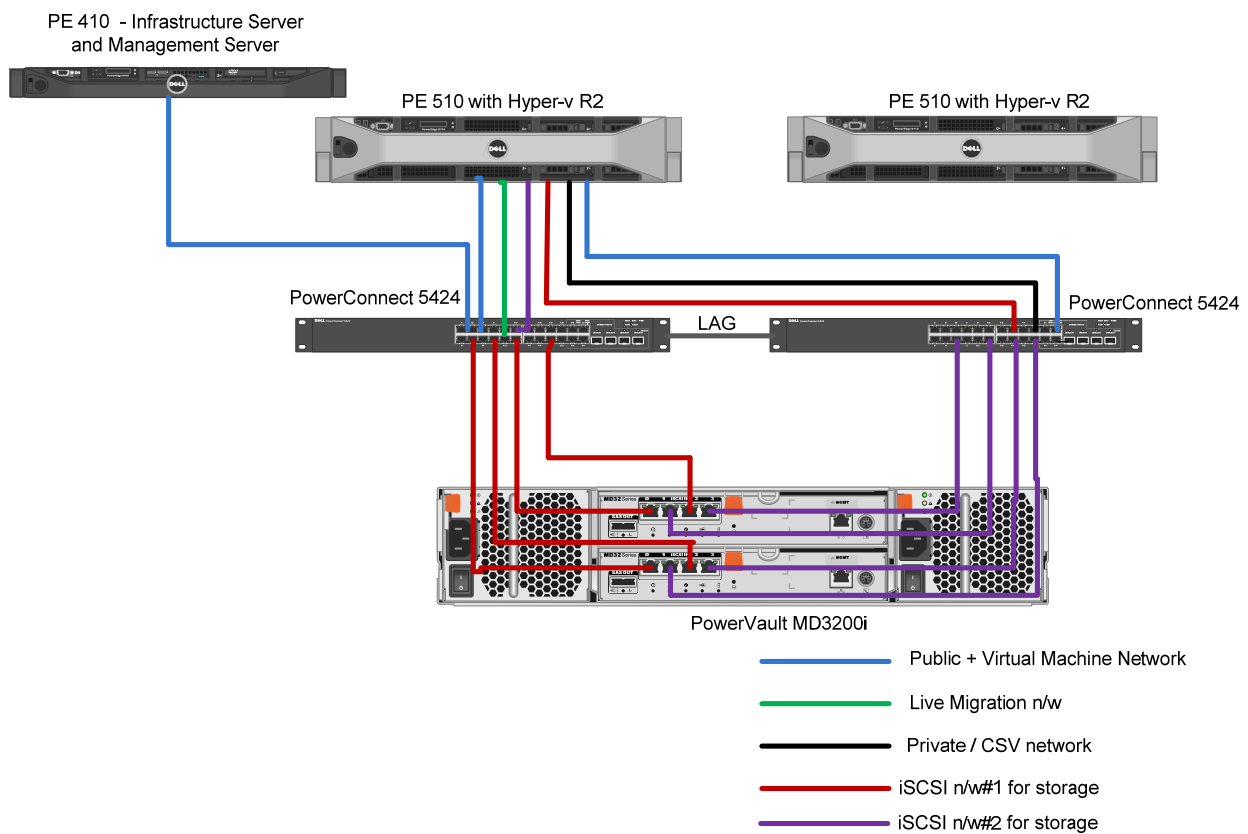


Figure 1: Reference Architecture diagram for the High Availability Solution

6.4 High Availability + Backup Solution

The reference architecture diagram for the High Availability + Backup configuration is as shown in the figure below. The solution consists of four PowerEdge R510 servers running Hyper-V R2 configured in a failover cluster connected to 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 in 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.

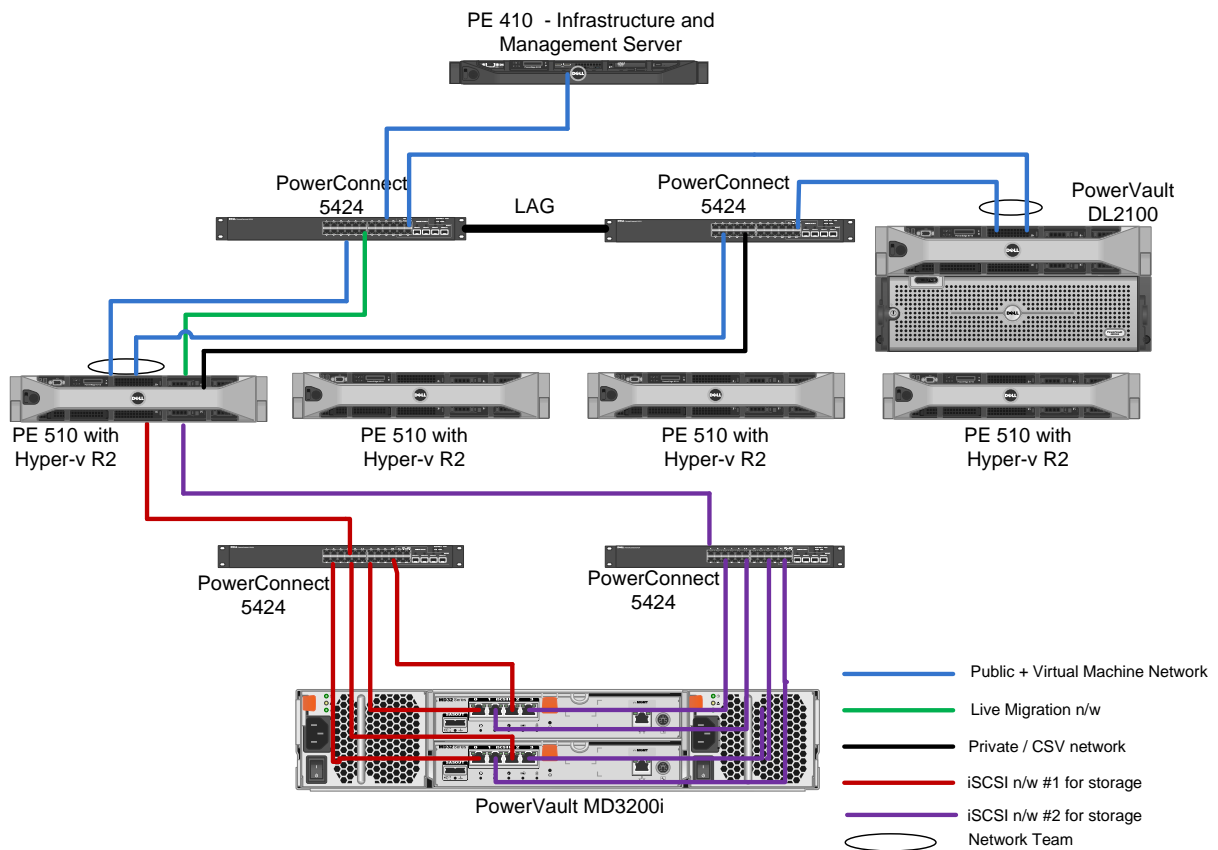


Figure 2: Reference architecture for High Availability + Backup Solution

7 Network Architecture

The Hyper-V R2 virtual infrastructure is comprised 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. In this section we describe 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 & 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, you can use two or four Dell PowerConnect 5424 switches. The PowerConnect switches are interlinked using a 4 port LAG (Link Aggregation Groups) which is used by all the networks besides the iSCSI traffic, as shown in figures 3 and 4, which provides optimal bandwidth for the traffic between the switches and failover capabilities should a link fail. The PowerConnect 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 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 used for the Virtual Machine, Cluster Public, Private/CSV, and Live Migration networks separated by VLANs. The second two switches are dedicated for iSCSI SAN traffic. 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.

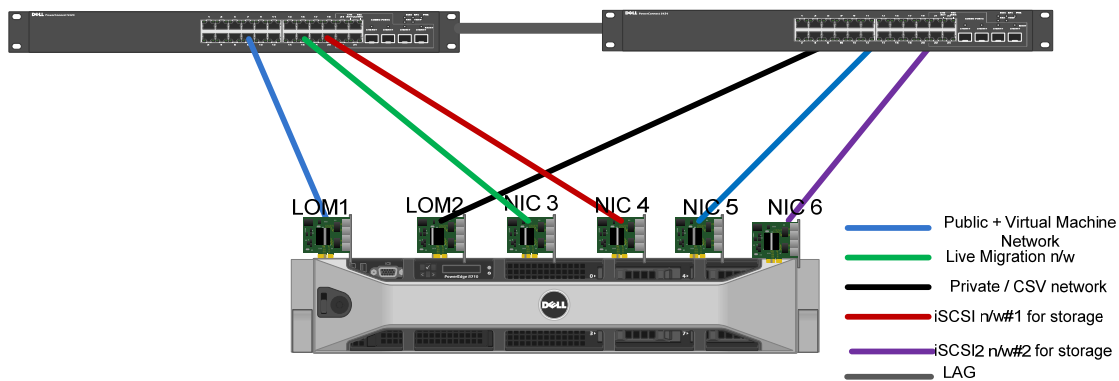


Figure 3: Single server network connection in High Availability Configuration

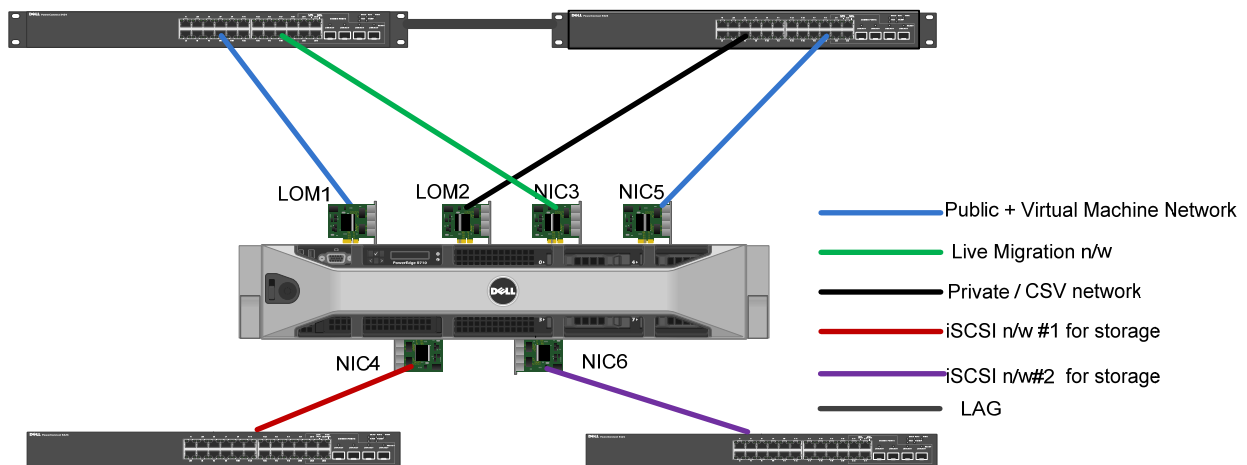


Figure 4: Single server network connection in High Availability + Back up Configuration

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 are isolated using VLANs. As mentioned earlier, 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. Table 3 gives a sample VLAN configuration that can be used for different types of network used in HA and HA + Backup configuration. This section will discuss more on the details for configuring different types of network and VLAN in these configurations.

iSCSI Network #1 and #2 (VLAN ID 10, 11): Used for iSCSI communication between Hyper-V R2 hosts and 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 the both the configurations, the iSCSI network includes two subnets to enable load balancing capabilities of the MD3200i.

Cluster Public Network or Management (VLAN ID 20): - Used by the cluster for public or external communication. For management purposes, 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 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 cluster public and Virtual Machine network.

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.

The Table 3 below shows sample VLANs that can be configured in the HA and HA + Backup configuration for each of the six networks 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

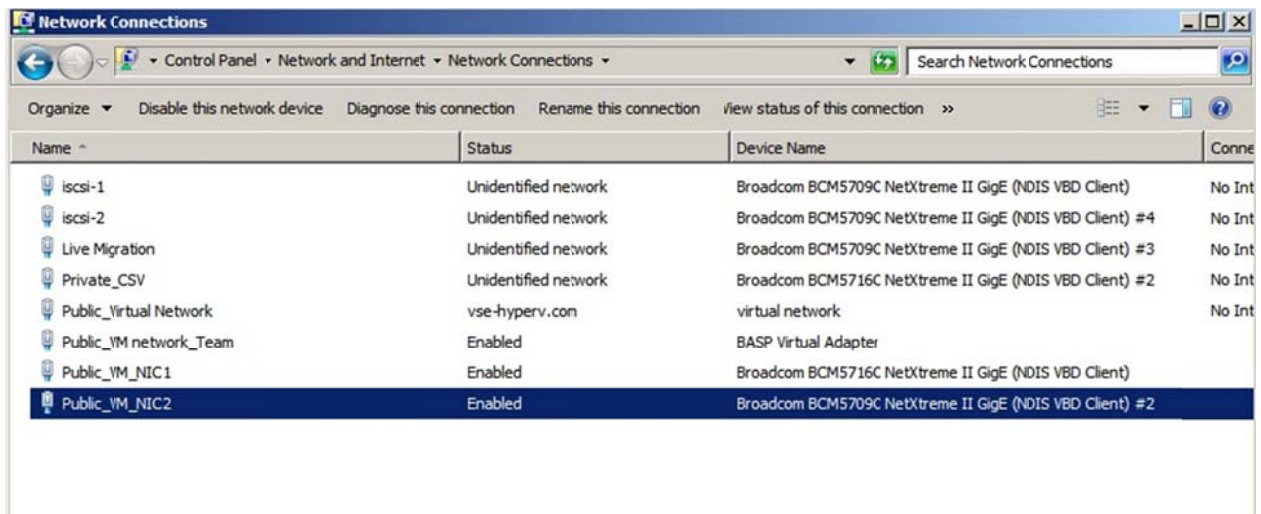


Figure 5: The Network Connections Manager showing different networks configured in the Hyper-V R2 hosts.

7.1.1 Virtual Machine Network/Cluster Public Network

These two network traffic types are configured using two Broadcom network adapters. The Broadcom Advanced Configuration Suite (BACS) is used to configure Smart Load Balancing (SLB) with failover capabilities for these networks. 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 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 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 respectively. Both networks 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 will discuss 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 PowerVault MD3200i 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 and PowerVault MD3200i iSCSI target.

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 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 NICs from each of the Hyper-V R2 hosts are connected to the switches ports configured for VLAN 10 and 11 respectively. For best performance, all data ports from MD3200i should be used for the iSCSI traffic. Two ports from each of the MD3200i controllers should be connected to each of the PowerConnect switches. Hence two of the four ports in one controller will be configured in one VLAN and the other two ports in other VLAN. The iSCSI ports on the MD3200i array should be configured in the respective VLANs. Note that the LAG between the two switches is configured in trunk mode to accept packets from all VLAN except VLAN 10 and 11 (iSCSI networks) in this configuration.

HA + Backup Configuration: The only difference between the HA configuration and the 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.

As mentioned earlier, the iSCSI network uses two subnets for storage connectivity to enable load balancing capabilities of the MD3200i. On the Hyper-V R2 hosts, the Microsoft software iSCSI initiator is configured to establish multiple sessions such that there is one session for each initiator IP address and each target IP address on the controller, so that all hosts can see four paths to each volume. Multiple paths are managed using a multi path driver that is provided as part of the 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. We have highlighted some of the important recommendations, however for more information for an IP SAN network in PowerVault™ MD3200i, and refer the IP SAN Best Practices document in the following link:

<http://www.dell.com/downloads/global/products/pvaul/en/ip-san-best-practices-en.pdf>

- 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 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 the Flow control is turned on.
- 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 iSCSI data ports on each RAID controller in MD3200i, network switches, and the NICs used for iSCSI traffic.

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 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 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 so as to balance controller utilization.

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. It is recommended to configure the array in RAID 10 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 less performance than RAID 10 in random I/O situations and requires additional overhead in the case of a drive failure scenario. 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 however is impacted by the write performance overhead for the double parity calculations it requires.

8 Management

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

8.1 Systems Management

Dell OpenManage 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 www.dell.com/openmanage.

The Dell Unified Server Configurator (USC) which ships with all 11G 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 configurations.

8.2 Storage Management

PowerVault MD3200i series arrays are managed by the enhanced MDSM software. It offers two different management paths - in-band array management and out-of-band array management. In the reference configuration discussed in this document out of band management is used by the MDSM software. Besides the array management facility, the enterprise management window feature monitors multiple systems, including previous generation MD3000i arrays, through a single interface. MDSM detects and alerts you to problems, then launches an automatic Recovery Guru to help you troubleshoot and resolve the problem.

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

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

Windows Server 2008 R2 Enterprise natively ships with the ability to locally manage Hyper-V R2 servers with limited management functionality like power cycle, migrate, etc. using the Hyper-V Manager. The Microsoft Failover Cluster Manager snap-in is another natively supported management feature, which also helps users to manage the Virtual Machines with limited functionality.

Dell OEM Microsoft System Center Essentials 2010 (SCE 2010) helps users to manage Virtual Machines, as well as the physical systems, using the same console. SCE 2010 is a Microsoft System Center family of 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), suitable for small and medium sized business. It 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, which helps administrators in monitoring the Hyper-V R2 systems health and 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.

9 Deployment

This section provides guidelines for the initial deployment and configuration of the 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 as per the architecture.
- Interlink the network switches.

2. Configure PowerConnect switches

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

3. Perform initial configuration of MD3200i

- Initialize MD3200i 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 the controllers.
- Connect the out of band management on each 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 for management.
- 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) assigned for iSCSI communication.
- Enable Jumbo frames on both the RAID controllers by setting the MTU to 9000 bytes on all 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 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. 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 (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 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/CSV, and live migration traffic. If there is no VLAN configured, the sample VLAN IDs and the subnet provided in the document can be utilized.
 - c. Configure a Hyper-V virtual switch on the teamed BASP 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 MDSM software on the hosts from the Resource CD provided along with the MD3200i on the hosts. 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 Failover Cluster Manager Feature for Remote Administration of Failover Cluster.
5. Configure 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 it to that during installation. Otherwise SQL Server Express Edition will be installed.
 - Refer to Microsoft documentation for planning, deploying, and managing using SCE 2010 at <http://technet.microsoft.com/en-us/library/ff603627.aspx>

7. Configure iSCSI initiator on the Hosts

- On the Hyper-V R2 servers, perform the following:
 - Using 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 port in the MD3200i controller. Together there will be eight sessions - four per initiator IP address to the target MD200i.

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

- 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 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 at [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)

The procedure for deploying failover Clustering and Hyper-V has the following 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 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 live migration network, give the priority to the NIC connected to VLAN 50 assigned for Live migration network.

10. 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.
- For deploying Virtual Machines on the Hyper-V R2 Failover cluster using SCE 2010, refer to the Microsoft documentation at: <http://technet.microsoft.com/en-us/systemcenter/essentials/ff631054.aspx>

11. 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 are enabled.
- Configure the discovery schedule.
- Refer to Dell and Management Systems Management Solutions link for downloads and user guide in installing management pack at 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)
- 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
- Planning for Hyper-V: PowerVault™ MD3200 and MD3200i Series of Arrays
<http://www.dell.com/downloads/global/products/pvaul/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/pvaul/en/powervault-md3200-series-hyper-v-implementation.pdf>
- PowerVault MD3200 and MD3200i: Array Tuning Best Practices
<http://www.dell.com/downloads/global/products/pvaul/en/powervault-md3200i-performance-tuning-white-paper.pdf>
- Installation and Configuration for PowerVault DL2200
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 of 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, CSV network, Live Migration network, and Virtual Machine Network, and the other 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 PowerVault iSCSI network need to be configured on the two switches.

PowerConnect 5424 configuration file used 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 used 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
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
```