



Microsoft Hyper-V Server 2012 R2 on Dell PowerEdge VRTX

A Dell Reference Architecture

Dell Global Solutions Engineering
August 2015

A Dell Reference Architecture

Revisions

Date	Description
June 2015	Initial release
July 2015	Updated network switch module to R1-2210 (10GbE switch)

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.

Copyright © 2015 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



Table of contents

Revisions.....	2
1 Scope	4
2 Audience.....	5
3 Introduction	6
3.1 Solution Specification	7
3.1.1 PowerEdge VRTX	8
3.1.2 Microsoft Windows Server 2012 R2 Hyper-V	11
4 Design Principles.....	12
4.1 Data Center in a box.....	12
4.2 A Configuration that Balances Performance and Utilization	12
4.3 High Availability	13
4.4 Reduced Complexity	13
5 Solution Architecture	15
5.1 Dell Blade Server Architecture for PowerEdge VRTX	16
5.2 Network Architecture.....	17
5.2.1 Cluster Network Design.....	18
5.3 Storage Architecture.....	21
5.3.1 Virtual Disk Configuration	22
5.4 Hyper-V Cluster Design	23
6 Efficient System Management with VRTX	25
6.1 Monitor and Receive Alerts Centrally Through Dell OpenManage Essentials	25
6.2 Geographical View	25
6.3 Dell PowerEdge VRTX CMC	26
6.4 System Center 2012 R2 Virtual Machine Manager.....	28
7 Infrastructure Verification.....	30
7.1 Storage Verification	30
8 Solution Summary	32
A Additional Resources.....	33



1 Scope

This document provides guidance to deploy virtualization-enabled remote office/branch office environments (ROBO) using Windows Server 2012 R2 on Dell PowerEdge VRTX chassis. It also discusses lifecycle management components used for infrastructure management.



2 Audience

IT administrators and IT managers who have purchased, or are planning to purchase virtualization infrastructure for a remote office, branch office, small office, or a small business can use this document to understand the design elements, hardware and software components, and the overall architecture of the solution.



3 Introduction

Dell PowerEdge VRTX is a shared infrastructure product designed to address the needs of remote, branch, and small offices. The solution combines servers, shared storage, I/O in the form of PCIe slots, and networking into a 5U tower chassis that is suitable for small and remote offices.

Dell PowerEdge VRTX combines blade servers with new mezzanine boards, planar, I/O, and storage within a single converged infrastructure. The shared infrastructure solution also incorporates a suite of redundant subsystems such as dual PCIe switches, CMC, PERC, and SAS expanders in a compact package that can hold up to four server nodes and up to 25 x 2.5" hot-swap drives offering up to 3TB of RAM, 64 processing cores, and 48TB of storage.

This cost-effective converged infrastructure allows large and medium-size enterprise remote offices and SMB offices to take advantage of consolidation, virtualization, and high availability for branch office workloads. Remote sites may have limited to no IT staff and depend on a central corporate location for most IT support. These central locations face various challenges such as infrastructure management, provisioning new servers, guaranteeing data protection, and providing business continuity with maximum uptime.

Dell PowerEdge VRTX's reference architecture deploys a Microsoft Hyper-V based branch office on a converged hardware design that reduces the hardware footprint (server, storage, and network) and optimizes resource utilization and centralized management.

This reference architecture is intended to provide customers with sample virtualization solution configurations that can be considered when designing a remote office solution on Dell PowerEdge VRTX chassis. This reference architecture is not intended to be an exhaustive study of specific architectures for every environment.



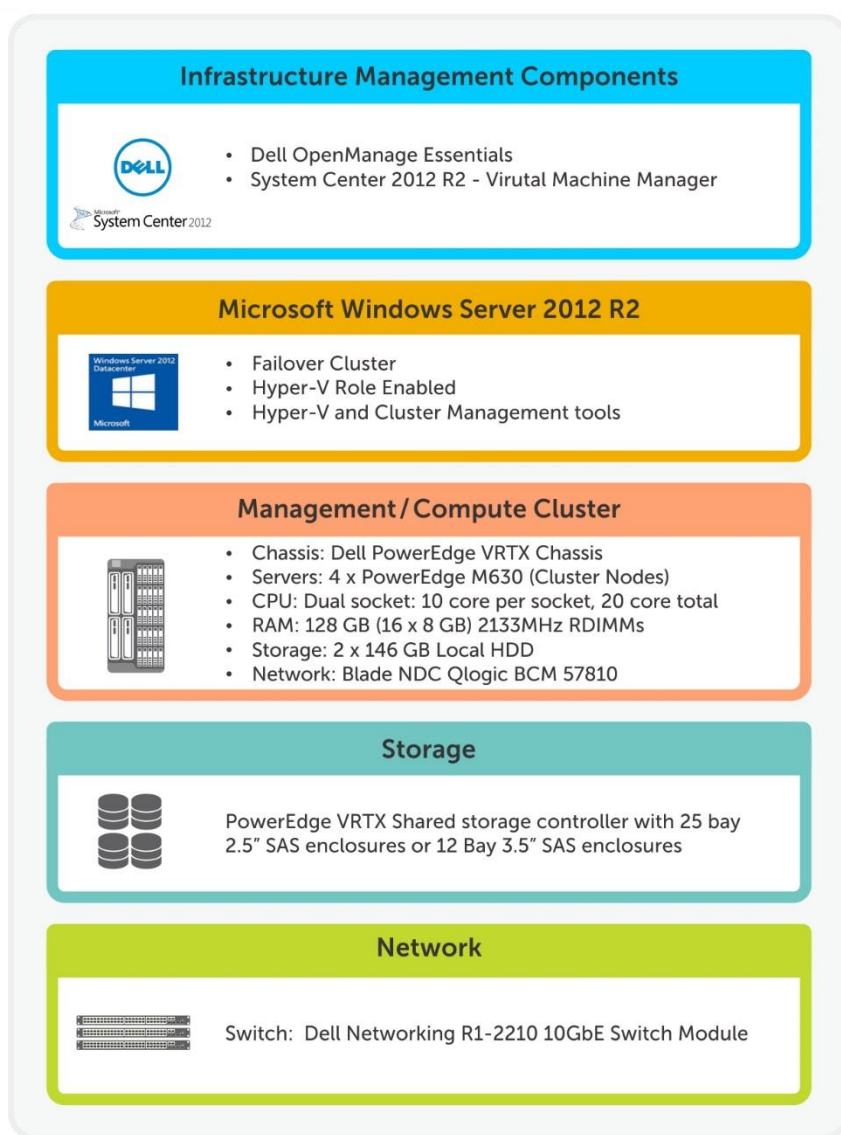
3.1 Solution Specification

Table 1 Dell PowerEdge VRTX solution specification

Component	Description	Role
Compute	4x PowerEdge M630 Blade Servers	Infrastructure to host Management VMs and the compute VMs.
Storage	PowerEdge VRTX Shared storage controller with 25 bay 2.5" SAS enclosures or 12 Bay 3.5" SAS enclosures	Provide shared storage for the Hyper-V clusters
Network Switches	Top of Rack Switches in the customer environment if required	External connectivity
Switch module	Dell Networking R1-2210 10GbE switch module	Support VM, Live Migration, Management and Cluster network
Microsoft Windows Server 2012 R2	Provides Hyper-V virtualization platform with high availability	Enable Virtualization and Host High availability
Microsoft System Center 2012 R2	System Center 2012 R2 <ul style="list-style-type: none">- Virtual Machine Manager- Operations Manager- Data Protection Manager	Enables infrastructure Monitoring, Management and business continuity
Dell OME	Dell Hardware management software.	Discovery, Monitoring and update for Dell Server, storage and network infrastructure



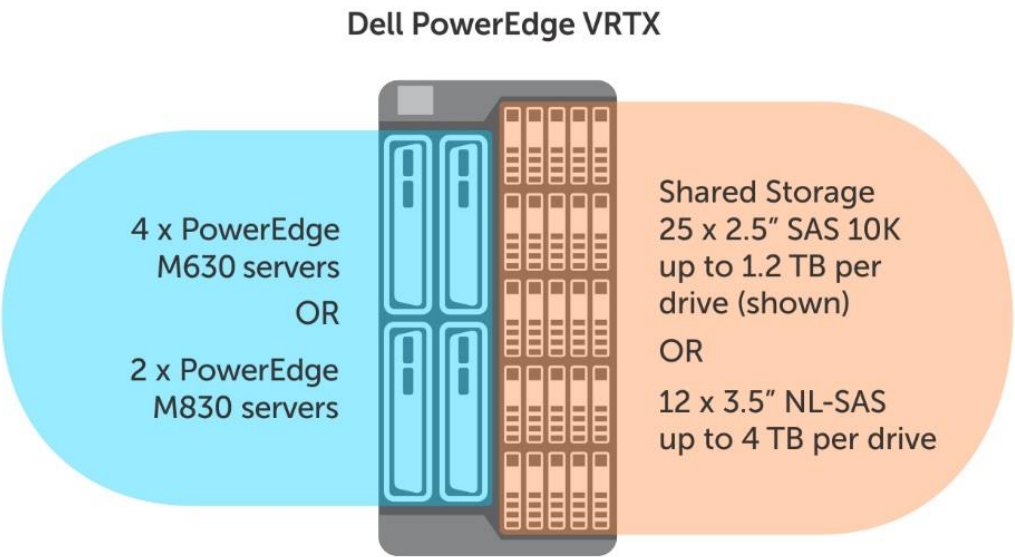
Figure 1 Solution components



3.1.1 PowerEdge VRTX

PowerEdge VRTX is a shared infrastructure platform that combines server, storage, and networking into an easy-to-deploy modular chassis. The platform is designed to address and resolve IT concerns of SMBs. PowerEdge VRTX addresses the needs of virtualized workloads and projected storage capacity and performance.

Figure 2 Configuration options for PowerEdge VRTX



PowerEdge VRTX uses I/O cards in the PCIe industry-standard format to provide shared storage of up to 48 TB in 3.5-inch HDD bay chassis. The storage can be shared among four server nodes and managed through the PowerEdge VRTX Chassis Management Controller (CMC). Using CMC, virtual disks can be created and assigned to single or multiple server nodes. (Virtual disks can be created and assigned to multiple server nodes only if clustering-aware software is installed.) The CMC web console can also be used to assign PCIe slots to server nodes. Up to four PCIe slots can be assigned to a single server node. These mappings can be reassigned later, but the servers involved in the reassignment need to be power cycled.

Table 2 Overview of PowerEdge VRTX infrastructure

Feature	Description
Server Compatibility	PowerEdge M520, M620, M630, M820, M830 server nodes
Form Factor	Stand-alone tower or 5U rack enclosure
Number of Servers	Up to 4 PowerEdge M520/M620/M630 servers or Up to 2 PowerEdge M820/M830 servers
I/O Cards	8 PCIe slots (supporting Ethernet, FC, GPU)
Power Supplies	Up to 4 PSUs (PSU and AC redundant options)



Chassis Storage	Up to twelve 3.5-inch NL-SAS, SAS HDDs/SSDs or Up to twenty five 2.5-inch NL-SAS, SAS HDDs/SSDs
Raid Controller	Shared PowerEdge Raid Controller (PERC 8; single Shared ¹ PERC with optional second redundant shared PERC8)
Management	1 or 2 Chassis Management Controllers
Network	1 GbE pass-through module or 1 GbE internal switch module (8 external ports) 10 GbE internal switch module (8 external ports)

3.1.1.1 Dell PowerEdge M630

Designed for workloads such as email, database, and virtual environments, the Dell PowerEdge M630 blade server is an ideal blend of density, performance, efficiency, and scalability. The M630 delivers unprecedented memory density and superb performance with no compromise of enterprise-class features.

The PowerEdge M630 blade server delivers value for both large and small data centers. It dramatically boosts application performance with the latest Intel Xeon processor E5 family, developed specifically for server environments. PowerEdge M630 takes advantage of the tremendous virtualization density with its 24-DIMM memory capacity and up to 36 cores, which give outstanding throughput in a compact, half-height blade chassis. The Dell PowerEdge M630 provides comprehensive internal storage options, with several drive types and storage controllers to choose from. The M630 supports up to 1.2TB of internal storage.

The Dell PowerEdge M630 offers balanced, scalable I/O capabilities, including integrated PCIe 3.0 capable mezzanine card slots. Dell's network daughter cards choose the right network fabric without using up a valuable mezzanine card slot. You select the speed, technology, vendor, and other options such as switch independent partitioning that can share and manage bandwidth on 10GbE connections.

3.1.1.2 Dell Networking R1-2210 10GbE switch module

As a consequence of changing application deployment methods over the years and increased VM movement between hosts, traffic patterns in the data center have changed from being north-south to predominantly east-west. The VRTX 10GbE switch module addresses this by offering Layer-2 switching and thus helps in keeping inter-server traffic within the chassis.

The VTRX 10GbE Switch Module, model R1-2210, is a layer 2 switch that supports 10/100/1000/10000 Mb Ethernet. All internal ports (te 1/1 – te 4/4) communicate together as if plugged into a rack switch. East-west traffic and north-south traffic are converged within the same chassis. The R1-2210 switch also offers

¹ The reference architecture uses redundant Shared PERC 8 controller.



increased bandwidth for inter-server communication and reduces cabling required to Top of Rack Switches compared to Dell VRTX PTM I/O Module.

The external ports allow communication to any upstream devices on the network like any typical switching/bridging device.

Key Features of the R1-2210 include:

- 10Gb server-to-server and external connectivity
- 16 internal 10Gb ports (SFP+)
- 4 external 10Gb ports (SFP+)
- 2 external 1Gb ports (RJ45)

3.1.2 Microsoft Windows Server 2012 R2 Hyper-V

Microsoft Windows Server 2012 R2 provides the Hyper-V virtualization platform; it helps to consolidate workloads enabling IT managers to optimize the use of their available hardware resources. The failover clustering feature of Microsoft Windows helps to increase the availability of the applications and services. If one node in the cluster fails the service will failover to the other node in the cluster to provide the service. Windows Server 2012 R2 provides failover clustering that enables high availability and live migration of virtual machines between physical servers.



4 Design Principles

While designing this reference architecture based on VRTX on Hyper-V, Dell has followed certain specific design tenets that address customer requirements with a remote office use case.

4.1 Data Center in a box

Given the remote office use cases VRTX appeals to, it is important that the solution be self-contained. For example the management components, whether SCVMM or OpenManage, should not require additional hardware outside the VRTX. Following are some of the inherent features of PowerEdge VRTX and design choices in the reference architecture that support this design tenet.

- The PowerEdge VRTX hardware consolidates server, storage, and networking and provides a self-contained solution that is easy to deploy and manage. All four server nodes have access to the low-latency internal shared storage that is ideal for virtualization and clustering, a networking card, and PCIe resources which are shared across the compute nodes within the chassis.
- Dell Chassis Management Controller (CMC) is leveraged to enable a single management interface for chassis infrastructure management. In the case of a remote office or a branch office, some of these components may also be installed at the head office or central data center, but having these components available locally allows the remote/branch office to continue operation in case of loss of network connectivity to the head office or data center. The single interface provided by the Power Edge VRTXs Chassis Management Controller enables a local administrator at the remote or branch office to manage the complete solution of servers, shared storage, network module, and PCIe devices with significantly lower running curve than a traditional solution.
- For the remote office use case, as compute resources are limited within the PowerEdge VRTX chassis, management components like SCVMM or Open Manage are hosted along with the workload applications to optimize use of the shared infrastructure.
- This reference architecture includes R1-2210 switch module catering to all I/O traffic to and from the VRTX chassis. The R1-2210 is capable of accepting traffic from up-to 16 server side ports from Fabric A, with output comprising of four external ports. This switch is the only VRTX Fabric A device that can support the full four-lane count from Fabric A. This reference architecture, however, uses a dual-port network card on the servers thus offering 2:1 oversubscription ratio of internal to external connectivity. The low oversubscription ratio ensures that the switch module is suitable for performance intensive environments.
- Including the switch within the chassis allows customers to completely eliminate Top of the Rack switches in most cases and save significant costs

4.2 A Configuration that Balances Performance and Utilization

Designing a balance configuration is important for customers, allowing them to deliver efficiency and effectiveness without leaving any money on the table. This Dell reference architecture has been designed as an “eco system” where all technology choices and recommendations are carefully evaluated in the context of the overall design.



- The choice of PowerEdge M630 servers in the reference architecture strikes a balance between density, capacity, and performance in comparison to M830 servers
- The reference architecture leverage 2.5 inch 10K SAS drives for internal storage. It provides a balance between performance and capacity needs of the office environments. As shown in Table 3, choosing 2.5 inch drives helps to achieve the sweet spot for IOPS and capacity.

Table 3 RAW capacity and IOPS (no RAID penalty) from VRTX supported drives²

	2.5		3.5	
	1 TB / 7.2K	1.2 TB / 10K	4 TB / 7.2K	600 GB / 15K
RAW capacity	24 TB	28.8 TB	44 TB	6.6 TB
IOPS	2400 (100 per spindle)	3360 (140 per spindle)	1100 (100 per spindle)	2310 (210 per spindle)

4.3 High Availability

Loss of productivity because of unplanned downtime is a grave concern for most IT customers. This reference architecture is designed to provide comprehensive hardware redundancy both at the topology and element levels. The design choices made in the reference architecture ensure that enough resources are available at all layers of the infrastructure.

- Each Dell PowerEdge M630 blade server in VRTX is configured with a dual port QLogic BCM 57810 blade NDC.
- The four node cluster comprising of Dell PowerEdge M630 servers provides high availability for the workloads by managing component or node failure.

4.4 Reduced Complexity

Most customers see reducing complexity as a key requirement for managing remote office environments given their limited IT staffing resources. Many of today's tools can concentrate on a single element domain, such as servers or VMs and make it complex for administrators to manage and monitor the systems remotely. Dell VRTX enables customers to manage and monitor systems in remote offices as one single entity.

Dell PowerEdge VRTX's system management is designed to help customers:

- Reduce the time spent on routine management of servers, storage, and networking. With PowerEdge VRTX, servers, storage, and networking are managed with a single, unified tool, which helps to reduce management complexity, save time, and reduce potential for error. For example,

² The table provides theoretical capacity and IOPS without RAID overhead, and considers one hot spare.



setting up Hyper-V cluster on VRTX takes just 23 steps versus 92 steps with a regular four-tower or rack server.

- Assess the health and status of remote systems, and take action on them as needed. Dell Open Manage Essentials provides a consolidated view across multiple geo locations and makes remote management simple and easy.

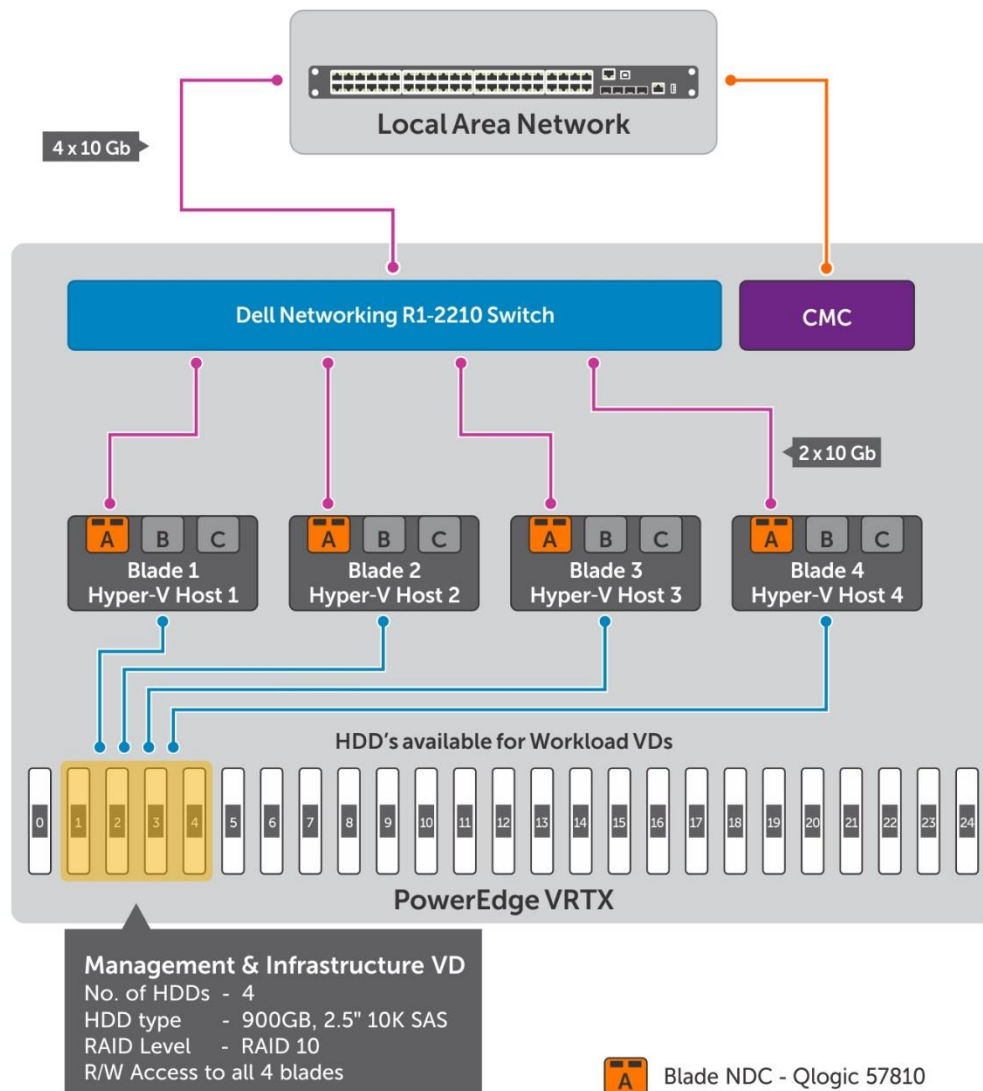


5 Solution Architecture

This solution consists of a PowerEdge VRTX chassis populated with PowerEdge M630 blade servers running Microsoft Hyper-V, shared storage, and networking. Figure 3 depicts the high-level design of the reference architecture.

Note: This reference architecture is a sample using a Qlogic network adapter for the host network; it doesn't restrict usage of other network vendor adapters supported on PowerEdge VRTX

Figure 3 Reference architecture for PowerEdge VRTX with Hyper-V



In the above design, each of Dell Blade server LOMs (LAN on Motherboard) directly maps to Fabric A. Fabric A goes to the 10 GbE Dell Networking R1-2210 switch module which provides 20Gbps throughput between the servers in the chassis. The switch module also has 4 x 10GbE uplink ports providing 40Gbps uplink to the Top of Rack switch in the customer environment.

Note: The mezzanine cards on the slot B and slot C are not used for network traffic in this reference architecture. However the mezzanine cards are required for connectivity to SPERC.

Table 4 shows PCIe switch mappings in a Dell VRTX chassis.

Table 4 PCIe Switch Mappings

PCIe Switch	PCIe slot Connectivity	Mezzanine Slot
Fabric B PCIe Switch	PCIe slot 3, PCIe slot 6 ,PCIe slot 7 , PCIe Slot 8 and a Shared PowerEdge RAID Controller (SPERC)	Mezz B
Fabric C PCIe Switch	PCIe slot 1 , PCIe slot 2 , PCIe slot 4 , PCIe slot 5	Mezz C

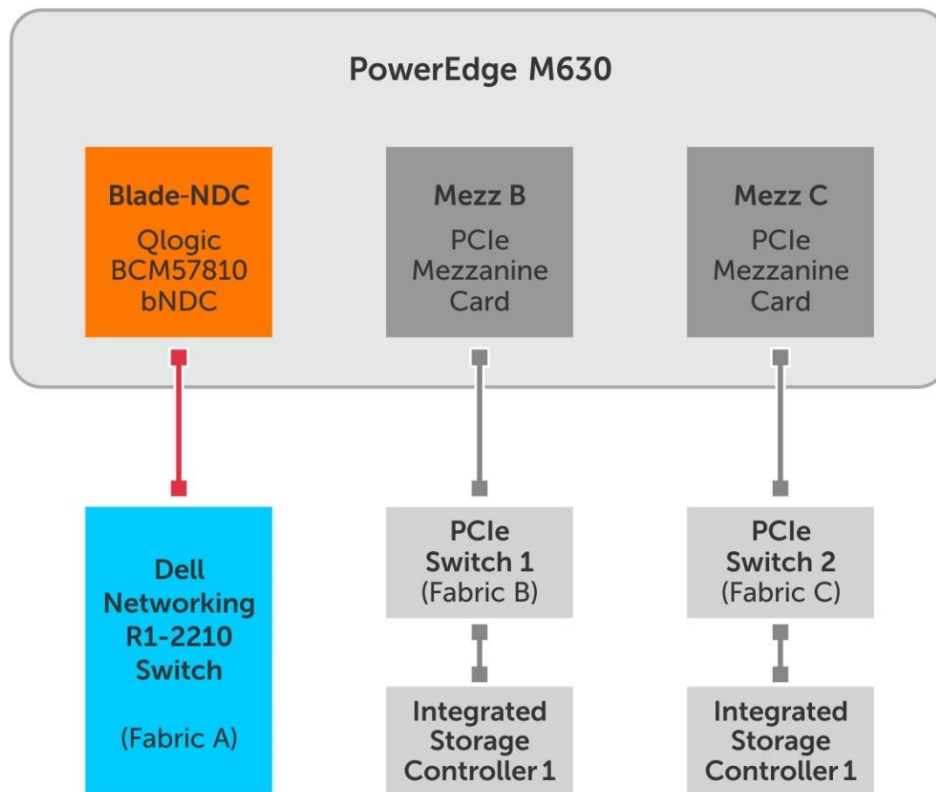
The SPERC is mapped to the blades in the VRTX chassis through Fabrics B and C PCIe switches. Subsequent sections provide more detailed information on various components involved in the solution.

5.1 Dell Blade Server Architecture for PowerEdge VRTX

Each PowerEdge M630 server is configured with one QLogic BCM 57810 network daughter card (NDC) in Fabric A. Figure 4 illustrates how the fabrics are populated in PowerEdge M630 and how the I/O modules are utilized.



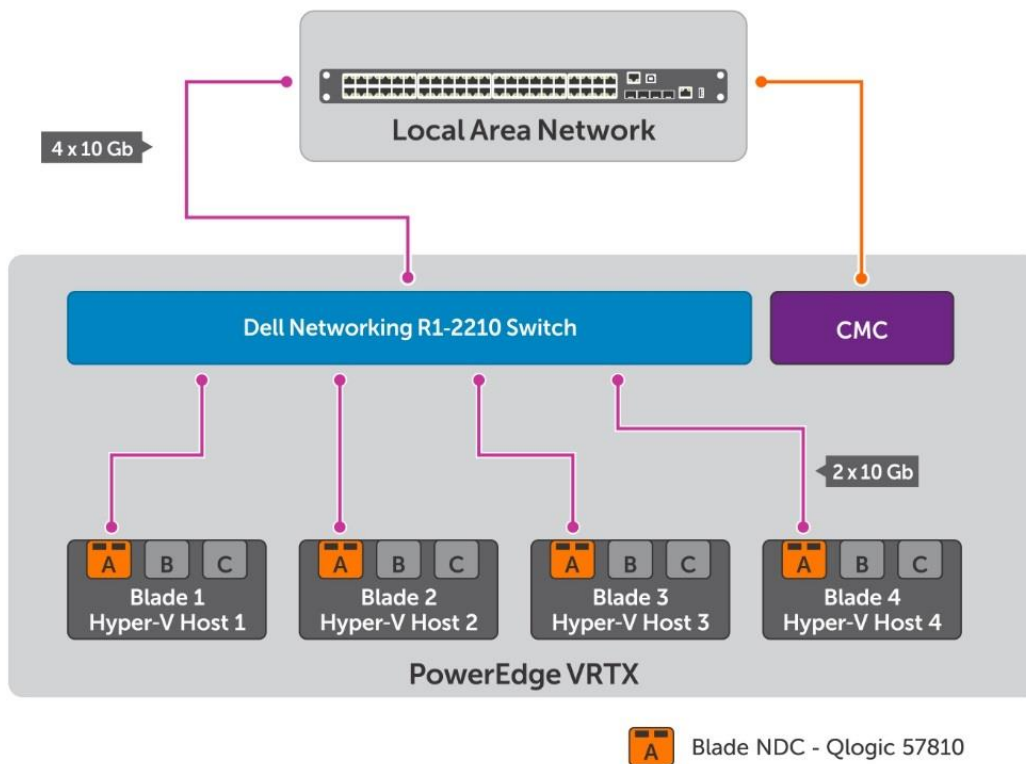
Figure 4 I/O connectivity of M630 on PowerEdge VRTX chassis



5.2 Network Architecture

In this design, the physical network is configured to eliminate the need for TOR switches within the Reference Architecture and also to provide high internal and external speeds. The east-west traffic can go up to 20Gbps and the uplinks can hold up to 40Gbps based on the configuration of the network team and uplink switch ports. This reference architecture includes R1-2210 switch module catering to all I/O traffic to and from the VRTX chassis.

Figure 5 Network architecture



Each Dell PowerEdge M630 blade server is configured with a dual port QLogic BCM 57810 bNDC running at 10GbE via Dell Networking R1-2210 switch module.

5.2.1 Cluster Network Design

This network design simplifies data center management by consolidating all communication (management, live migration and cluster heartbeat) onto a single fabric by using OS native NIC Teaming, a new feature included in the Windows Server 2012 R2 operating system that allows the teaming <http://technet.microsoft.com/en-us/cloud/private-cloud> of any two NICs regardless of vendor.

A fault-tolerant network team is created by using the four network ports (2 blade-NDC ports) available on each of the PowerEdge M630 servers using the settings described in the following table:

Table 5 Network team settings

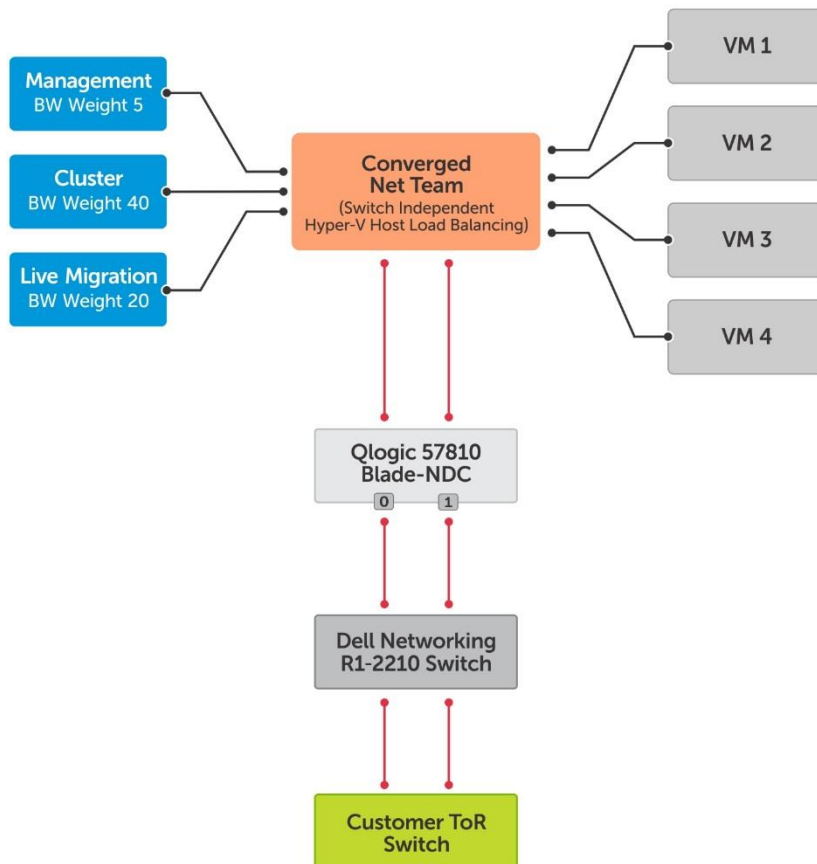
Network Team Settings	
Teaming mode	Switch independent
Load balancing algorithm	Dynamic
Standby adapter	None

The "Switch independent" teaming mode is used because Windows Server 2012 R2 simplifies the configuration of the team. No special switch configuration changes are necessary. The load-balancing algorithm used is "Dynamic". Dynamic mode is a new feature of Windows Server 2012 R2. This load balancing algorithm balances traffic based on flowlets. It can spread traffic both inbound and outbound. Hyper-V port algorithm forces all VM traffic only through one physical NIC, whereas Dynamic mode spreads traffic across all the NICs regardless of workload.

Using the teamed network, a virtual switch is provisioned in Hyper-V, and separate virtual network adapters were created for management, live migration and cluster heartbeat with a specific VLAN tag and bandwidth weights. Figure 6 shows the network design and its configuration.



Figure 6 Cluster Network Design



Tables 6 and 7 describe in detail usage of the different traffic types associated with the Hyper-V cluster in the network architecture

Table 6 Network traffic description

Traffic Type	Use
Management	Supports virtualization management traffic and communication between the host servers in the cluster
Cluster / CSV	Supports internal cluster network communication between the servers in the cluster.
Live Migration	Supports migration of VMs between the host servers in the cluster.
VM	Supports communication between the VMs hosted on the cluster and external systems

The VLAN IDs are assigned to the virtual network interfaces to segregate the traffic on the host between the cluster, live migration, management, and other types of traffic that are described in Table 8.



Table 7 Network VLAN configuration

Traffic Type	VLAN
Management	10
Cluster / CSV	20
Live Migration	30
VM	40

Hyper-V QoS in Windows Server 2012 R2 enables minimum bandwidth, which guarantees a specified minimum bandwidth for a traffic flow (identified by a Hyper-V Virtual Switch port number). Traffic flows that have no weight assigned are handled by the default flow. You should also assign a weight to the default flow. The sum of all the assigned weights (including the weight of the default flow) is the basis of the calculation for each traffic flow minimum bandwidth percentage.

Table 8 provides a sample QoS bandwidth allocation for the different Hyper-V cluster virtual network adapters.

Table 8 Network QoS bandwidth allocation

Traffic Type	Minimum Bandwidth Weight
Management	5
Cluster / CSV	40
Live Migration	20
DefaultFlowMinimumBandwidthWeight	10

5.3 Storage Architecture

The PowerEdge VRTX chassis offers an integrated, shareable storage with a new I/O virtualization technology that allows multiple compute nodes to talk to the same storage controllers. This technology provides faster access time and requires less setup than a traditional Storage Area Network (SAN) environment.

The PowerEdge VRTX chassis supports either 12 x 3.5-inch hot-swap drives or 25 x 2.5-inch hot-swap drives. The drives are supported by a shared PowerEdge RAID Controller, or SPERC. The VRTX storage controller presents four virtual adapters, one for each blade in the chassis. These virtual adapters appear to the compute node almost identically as a standard Dell PERC card. The shared storage controller offers RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60, as well as options like hot-spare.

The VRTX CMC provides the ability to change the virtual adapter mapping to the servers, as well as the ability to map a virtual disk to a virtual adapter. The VRTX chassis has a default assignment for the virtual adapters as mentioned in table below.



Table 9 Virtual adapters mapping

Virtual Adapter	Server Slot Mapping
Virtual Adapter 1	SLOT-01
Virtual Adapter 2	SLOT-02
Virtual Adapter 3	SLOT-03
Virtual Adapter 4	SLOT-04

It is necessary to set the assignment mode settings for the virtual adapters to multiple assignment. This allows a virtual disk to be assigned to multiple virtual adapters at a time.

5.3.1 Virtual Disk Configuration

Storage virtual disk is one of the critical components in hosting a Hyper-V cluster. As per the reference architecture design it requires minimum of three VDs, the first for cluster quorum configuration, the second for cluster shared volume (CSV) to host the management infrastructure VHDs, and the last for the cluster shared volume for compute/workload VM VHDs. As per the design, dedicated storage drives are used for the management VDs and the compute VDs.

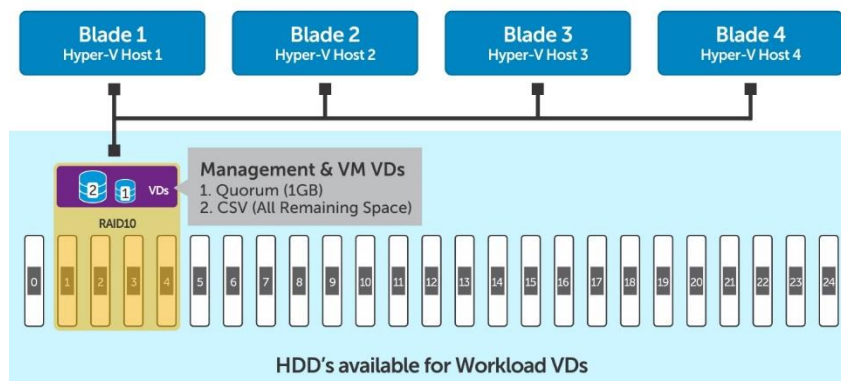
All the virtual disks are given read-write access to all four physical nodes in VRTX chassis to enable failover cluster as well as live migration. There is one dedicated global hot spare disk configured to handle physical hard drive failure in shared storage of the VRTX Chassis.

In 3.5-inch disk PowerEdge VRTX chassis the cluster quorum and management infrastructure CSV is configured on the first 2 disks with RAID 1 to provide greater protection. The remaining disks are left for the workload VMs, the raid configuration; size and the setting are dependent on the workload.

Similarly, in 2.5-inch PowerEdge VRTX chassis the cluster quorum and management infrastructure CSV is configured on the first 4 disks with RAID 10 along with the respective settings mentioned in Table 11.

Figure 7 details how the storage is configured in this reference architecture.

Figure 7 Virtual disk configuration on 2.5-inch PowerEdge VRTX chassis



The table below provides detailed information on the virtual disk configuration used in this reference architecture:

Table 10 Virtual disk settings

Virtual Disk	Disk type	No of Disks	RAID Level	Minimum Capacity	Read Policy	Write Policy	Stripe Element Size
Management & Infrastructure Virtual Disks	Cluster Quorum	4 x 1.2 TB 10k SAS Disks	RAID 10 (25disk Chassis)	1 GB	Adaptive Read Ahead	Write Back	64KB
	CSV (Management Infra)			Remaining space on the same Raid level used for the Quorum disk	Adaptive Read Ahead	Write Back	64KB
Compute & Workload Virtual Disks	CSV (Compute infra)	20 Disks * (25 disk chassis)	RAID 10* (25disk Chassis)	Depending upon the workload requirement	Adaptive Read Ahead*	Write Back*	64KB*
Global Hotspare		1 Disk					

* Configurable as per application best practices.

5.4 Hyper-V Cluster Design

Recommended practice is to have two or more Hyper-V servers in a dedicated cluster for the fabric management virtual machines, and then have separate clusters for the fabric itself. In small infrastructure deployments where the compute environment is limited, fabric management virtual machines could be hosted on the fabric itself for better utilization and management.

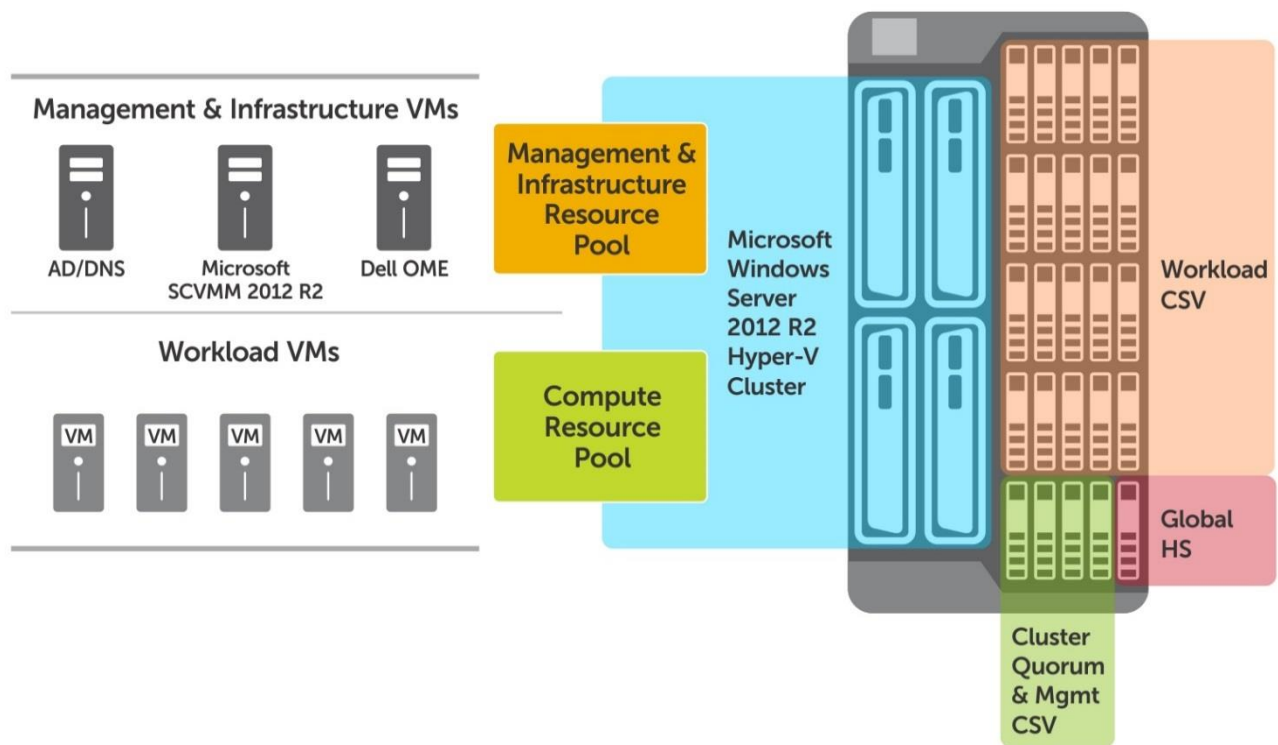
In this solution, a Hyper-V failover cluster was designed for the purpose of virtualizing the common infrastructure to increase availability and enable access to applications during planned or unplanned downtime. This design enables better resource utilization as well. There are separate highly available virtual network adapters configured for cluster heartbeat, live migration and management as discussed earlier.



Quorum configuration in a failover cluster determines the number of failures that the cluster can sustain. If an additional failure occurs, the cluster must stop running. Quorum and Cluster Shared Volumes (CSVs) in a Windows Server 2012 R2 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume.

With CSVs, clustered roles can fail over quickly from one node to another without requiring a change in drive ownership, or dismounting and remounting a volume. CSVs help simplify managing a potentially large number of LUNs in a failover cluster in this design. Separate CSVs are used for the management VMs and the workload VMs to store the VHDs. Figure 8 explains how the Hyper-V cluster is designed in the VRTX solution as well as how the VMs are placed on the same infrastructure.

Figure 8 Hyper-V Cluster architecture



6 Efficient System Management with VRTX

Businesses with remote offices or branch offices can benefit from streamlined management tasks offered by PowerEdge VRTX.

6.1 Monitor and Receive Alerts Centrally Through Dell OpenManage Essentials

Dell OpenManage Essentials (OME) is a tool to monitor the health of your cluster. Dell OME works by communicating directly to the systems through iDRAC8 with Lifecycle controller, monitoring everything from temperature to hardware faults.

Dell OpenManage Essentials (OME) is installed on an individual VM to monitor the hardware components of this reference architecture. With OpenManage Essentials, a web-based and one-to-many systems management application for Dell systems and other devices, one can:

- Discover and inventory the systems.
- Monitor the health of the system.
- View and manage system alerts.
- Perform system updates.
- View hardware inventory and compliance reports.

OpenManage Essentials provides pre-defined reports for all discovered and inventoried devices. With these reports, one can:

- Consolidate information about devices in the environment.
- Filter report data based on the devices by clicking the Filter by: drop-down list. And also it enables to add a new group of devices from the dashboard by clicking Add New Group from the Filter by: drop-down list.
- Export data for use in another application in the XML file format.

6.2 Geographical View

Managing and monitoring systems across multiple remote sites distributed globally from multiple interfaces could be complex and challenging. Dell Open Manage Essentials solves that problem by providing a consolidated view of all the systems across multiple geos.



Figure 9 Geographical view



The Geographical View feature can be used to map the locations of PowerEdge VRTX systems that are geographically dispersed around the world. This allows the capability to monitor the PowerEdge VRTX infrastructure at different remote locations from a single map-view. It requires no additional configuration operations for VRTX systems; all map locations are entered and managed through the OME console

It provides two options to define and map the locations of VRTX systems:

- Pin each VRTX system to the map by entering the address in the console manually
- Pin all VRTX systems in a single process by importing a list of their addresses in a CSV file

6.3 Dell PowerEdge VRTX CMC

The PowerEdge VRTX uses a chassis management controller (CMC) to manage the system. This includes management of server nodes, the power, cooling and the storage and PCIe subsystems. The CMC has its own microprocessor and memory and is powered by the modular chassis into which it is plugged.

CMC enables an IT administrator to:

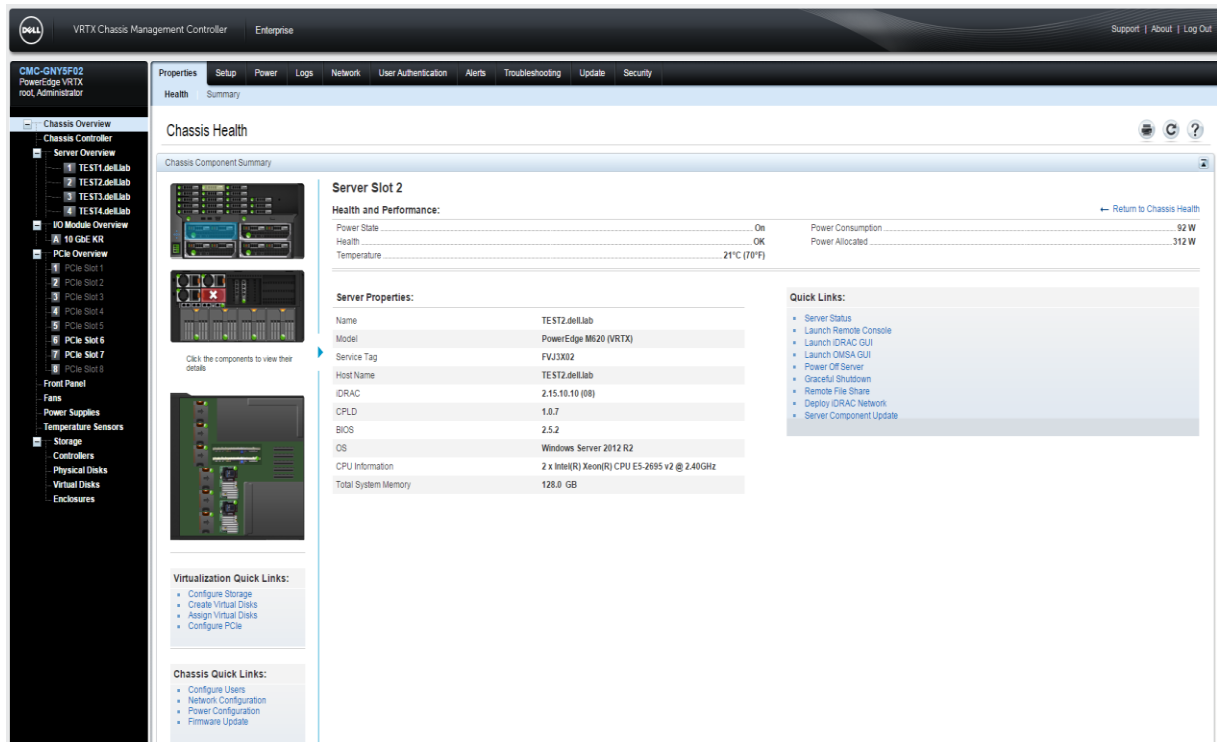
- View inventory
- Perform configuration and monitoring tasks
- Remotely turn on and turn off chassis and servers
- Enable alerts for events on servers and components in the server module
- View and manage storage controller and hard disk drives in the VRTX chassis



- Manage the shared PCIe subsystem in the VRTX chassis
- Provide a one-many management interface to the iDRACs and I/O modules in the chassis

The PowerEdge VRTX chassis can either be with a single CMC, or with redundant CMCs. In redundant CMC configurations, if the primary CMC loses communication with the chassis or the management network, a standby CMC takes over the chassis management.

Figure 10 PowerEdge VRTX Chassis Management Controller(CMC)



As show in figure above, CMC provides multiple System Management functions for servers. Power and thermal management are the primary functions of CMC, which are listed as follows:

Enclosure-level real-time automatic power and thermal management.

- CMC monitors system power requirements and supports the optional Dynamic Power Supply Engagement (DPSE) mode. This enables CMC to enable or set power supplies in standby, dynamically depending on load and redundancy requirements to improve power efficiency.
- CMC reports real-time power consumption, which includes logging high and low points with a time stamp.
- CMC supports setting an optional enclosure maximum power limit (System Input Power Cap), which alerts and takes actions such as limiting the power consumption of servers, and/or preventing the turning on of new servers to keep the enclosure under the defined maximum power limit.

- CMC monitors and automatically controls the functions of cooling fans and blowers on the basis of actual ambient and internal temperature measurements.
- CMC provides comprehensive enclosure inventory and status or error reporting.

CMC provides a mechanism for centralized configuration of the following:

- PowerEdge VRTX enclosure's network and security settings
- Power redundancy and power ceiling settings
- I/O switch and iDRAC network settings
- First boot device on the server module
- Checks I/O fabric consistency between the I/O module and servers. CMC also disables components, if necessary, to protect the system hardware
- User access security
- Storage components
- PCIe slots

CMC can be configured to send email alerts or SNMP trap alerts for warnings or errors such as temperature, hardware misconfigurations, power outages, fan speeds, and blowers.

6.4 System Center 2012 R2 Virtual Machine Manager

The Microsoft System Center 2012 R2 Virtual Machine Manager is enterprise class management software that enables administrators to host Hyper-V Cluster to provision, monitor and manage the branch office virtual infrastructure. It provides the fundamental services for creating and managing virtual machines, optimizing resources, and rapidly deploying infrastructure easily and effectively. It manages both physical and virtual environments from a single management console and thus avoid the complexity of using multiple consoles typically associated with managing an IT infrastructure.

The key capabilities of SCVMM 2012 R2 include:

- Hyper-V life cycle management (deploy, configure, manage, patch)
- Network Management & Storage Management
- Enterprise class management suite for managing Hyper-V environments
- Support for intelligent placement of virtual machines
- Dell Server PRO Management Pack integrates with Microsoft System Center Virtual Machine Manager (SCVMM) and SCOM to provide customers with effective management of Dell physical servers that host virtual workloads
- Native capability for Physical to Virtual migration and Virtual to Virtual migrations
- Integration with Failover Clustering to support HA and Quick Migration of virtual machines
- Easy automation capabilities leveraging Windows PowerShell

SCVMM also helps to proactively balance the load of VMs across a cluster and also decides the right number of hosts to run the workload to optimize the resource and power. It also helps to abstract server, network and storage resources. The VMM monitoring pack helps the Operations Manager to monitor the



availability of VMM, and it also monitor the availability, health, and performance of all virtual machines and virtual machine hosts.



7 Infrastructure Verification

The failover capability of the virtualized infrastructure is verified. This includes both hypervisor host failure for unplanned downtime and live migration of VMs from one host to another for planned downtime.

After the infrastructure is deployed it is essential to verify whether the infrastructure functions as designed. Basic infrastructure verification was done on the infrastructure by running generic workload. The infrastructure is verified for host failure and ability of the infrastructure to handle heavy network and I/O stress for a long period of time. During stress, virtual machines were migrated across hosts to verify ability of virtual machines to be moved without disruption.

7.1 Storage Verification

For a virtualization infrastructure, it is important to ensure that each component within the infrastructure is sized appropriately to run the virtualized applications. Hence, the performance of the VRTX integrated, shared storage plays an important role in the overall solution performance. While it's not practical to size every type of application that may run on a general-purpose infrastructure, it is still important to determine the overall performance that the storage subsystem can offer.

To this extent, the storage verification of the solution infrastructure was carried out to understand how the underlying storage performs. An Online Transaction Processing (OLTP) IO load profile was simulated using IOMeter. Table 11 provides the IO profile details.

Table 11 IO Profile for storage verification.

IO Profile Type	Read	Write	Access Size	IO Pattern
Database – OLTP	70%	30%	8k	100% Random

For IOMeter verification, two VMs per cluster node were created and the storage verification tests were performed within the VM. Table 12 provides the VM configuration used for the IOMeter verification.

Table 12 VM Configuration for IOMeter verification

VMs per Host	Virtual CPU/VM	Memory /VM	OS VHDX	Data VHDX
2	10	20 GB	40 GB	40GB

As shown in Table 12 a 40GB fixed VHDX was used for IOMeter verification. These VHDX were stored in a CSV volume as showing in Table 10.

Table 13 provide data generated from IOMeter runs on the PowerEdge VRTX architecture.



Table 13 IOMeter performance data from one PowerEdge VRTX architecture

Queue Depth	1	4	8
IOPs	1701	4135	5699
MBps	13	33	44
Read MBps	9	23	31
Write MBps	4	10	13
Transactions per Second	1701	4135	5699
Average Response Time (ms)	5	8	11
Average Read Response Time (ms)	6.52	10.86	15.85
Average Write Response Time (ms)	0.45	0.45	0.45
Average Transaction Time	4.70	7.74	11.23

Table 13 shows the number of IOPs the storage subsystem in a VRTX environment can deliver. These figures are only indicative and it is always recommended to size the storage infrastructure by using real-world workload scenario that is planned for the deployment.



8 Solution Summary

The Dell PowerEdge VRTX shared infrastructure platform resolves the complexities involved in managing infrastructure for application deployments and enables IT administrators to combine servers, storage and networking into an easy-to-deploy chassis. Potential errors are reduced with the use of unified and simplified systems management practices.

The design points emphasize the use of the latest hardware and software features presented in PowerEdge VRTX enable high consolidation of applications, optimum performance for branch office centric applications and high availability of infrastructure. The infrastructure HA features ensure that the application services are continuously available and deliver optimal performance.

The shared storage architecture is simple to configure and manage using the Chassis Management Controller (CMC) interfaces. Creating multiple virtual disks for VM and application data achieves both data isolation and optimal performance.

The converged virtual switch used for the VM network access enables high quality-of-service (QoS) for different classes of traffic and ensures that the demands for network bandwidth are met. Hence, the bandwidth requirements for VM live migration and cluster traffic are given priority, when needed, to enable seamless movement of VMs during optimization and/or planned downtime of the hypervisor hosts.

To fully abstract the application VMs from their physical hosts, VMs must be placed dynamically. VMs must also be load balanced across servers depending upon their workload type and their relationship to other VMs. To achieve server load balancing, the cluster should be configured with SCVMM Dynamic Optimization (DO) to automatically live migrate VMs to less utilized hosts. SCVMM availability sets should be used to ensure that the VMs with similar application roles or high resource requirements do not coexist on the same hypervisor host.

By offering various design choices and guidance on choosing right components, this Dell reference architecture takes the “guess work” out of solution design and makes it simpler for customers to transform their branch office infrastructure and improve overall efficiency.



A Additional Resources

1. **Dell.com/support** is focused on meeting customer requirements with proven services.
2. **DellTechCenter.com** is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices and information about Dell products and installations.
3. Referenced or recommended Dell publications:
 - a. [Dell PowerEdge VRTX](#)
 - b. [Deploy Microsoft UC&C applications on PowerEdge VRTX](#)
 - c. [Microsoft Hyper-V Server 2012 R2 on Dell PowerEdge VRTX Deployment Guide](#)
 - d. [Configuring Fault Tolerance with a Dell Networking R1-2210 in a VRTX Chassis](#)
4. Referenced or recommended Microsoft publications:
 - a. [Configuring Dynamic Optimization in VMM](#)
 - b. [Configure Availability Sets in VMM](#)

