Remote Office Infrastructure – A Microsoft Hyper-V Server 2012 Reference Architecture on Dell PowerEdge VRTX

Dell Reference Architecture solution guidance to design remote office infrastructure using Dell PowerEdge VRTX chassis.

Balamurugan B

Dell Engineering
June 2013

Revisions

Date	Description
June 2013	Initial release

Table of contents

Dell Reference Architecture White Paper | Remote Office Infrastructure – A Microsoft Hyper-V Server 2012 Reference Architecture on Dell PowerEdge VRTX | Version A00

Executive summary

IT organizations are looking for simple, efficient and manageable solutions to build branch office infrastructures rather than managing several independent entities like servers, storage, and networking. Dell PowerEdge VRTX is an all-in-one solution, which has server, integrated storage, and networking, as well as an efficient management console, allowing you to easily build branch office infrastructures. The Dell PowerEdge VRTX is a shared infrastructure platform offering extensive performance and capacity with office-level acoustics in a single, compact tower chassis. It is an ideal solution for small and midsize businesses as well as branch offices of large enterprises.

The goal of this reference architecture on Dell PowerEdge VRTX a remote office converged hardware solution is to provide guidance to define and deploy an efficient branch office infrastructure solution with minimum management and support overhead. It also provides comprehensive design information on the core infrastructure services that are hosted on the virtualization technology of Microsoft.

This horizontal Microsoft-based reference architecture for remote office and branch office solutions has the following benefits:

- Provides virtualization at multiple layers to meet remote, branch, and small office infrastructure requirements.

- Robust availability and fault tolerance built into the platform to protect all virtualized applications.

- Comprehensive solution to rapidly deploy and provision the virtual infrastructure.

- Effective monitoring and management of physical, virtual and software resources.

- Integration of Comprehensive Data Protection with reliable backup and recovery solutions to protect remote office data.

1   Introduction

Dell PowerEdge VRTX is a new shared infrastructure product focused on 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 the small and remote offices.

Dell PowerEdge VRTX combines Blade servers with new Mezzanine boards, planar, I/O and storage on to a single converged infrastructure. It also incorporates a suite of redundant subsystems such as dual PCIe switches, CMC, PERC, and SAS expanders into 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 as such, have limited to no IT staff and depend on a central corporate location for most IT support. These central locations face various challenges like infrastructure management, provisioning new servers, guaranteeing data protection, and providing business continuity with maximum uptime.

This solution provides a reference architecture to deploy a Microsoft Hyper-V based branch office on the converged hardware to reduce the hardware foot-print (server, storage, and network) to optimize the resource utilization and the centralized management.

The Virtualization Reference Architecture is intended to provide customers with an example of virtualization solution configurations that can be considered when designing a remote office solution on Dell PowerEdge VRTX chassis. This white paper is not intended to be an exhaustive study of specific architectures for every environment.

1.1   Objective

The objective of this reference architecture is to provide necessary guidance to deploy virtualization-enabled remote office/branch office environments (ROBO) using Windows Server 2012 on Dell PowerEdge VRTX chassis. It also discusses about different lifecycle management components for infrastructure management and data-protection.

1.2   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.

2   PowerEdge VRTX

Dell PowerEdge VRTX is a new modular infrastructure targeted for Remote Office/Branch Office (ROBO) installations. It combines existing blade servers (M520, M620) with new mezzanine boards, Planar, I/O and Storage infrastructure designed for high availability by integrating a suite of redundant subsystems.

Figure 1   PowerEdge VRTX Simplified Block Diagram
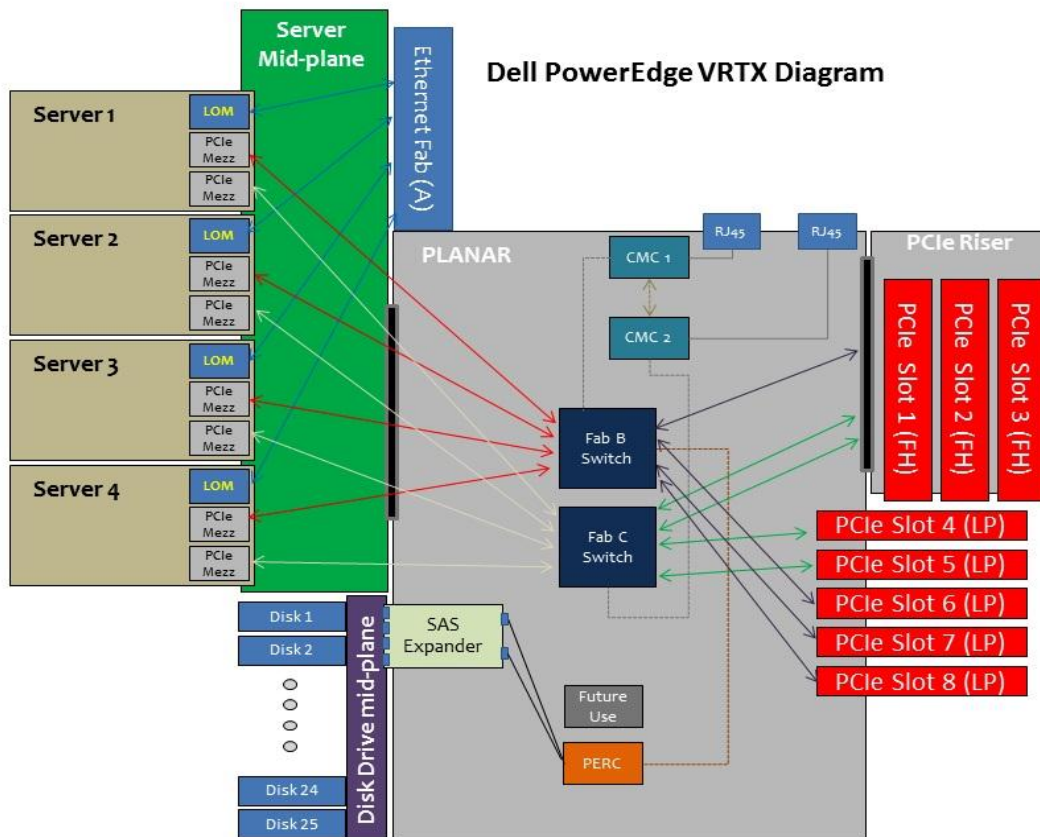
As shown in the figure 1, the system supports up to four Blade servers (half-height) which connect to the infrastructure through a mid-plane. The system Planar provides dual switches to direct traffic from the blades to the other PCIe-based elements: the PERC controllers and PCIe slots. The Chassis Management Controller (CMC) plugs into the planar and serves as the main entity to control switch paths, sequence the power-on of blades, PCIe cards etc. PERC controller (H710P-mini) docks into the planar and control the storage subsystem, which can have up to two expanders and supports either 2.5" or 3.5" disk options. Other than a simple switch for Fabric A Ethernet, the platform uses a shared PCIe infrastructure through PCIe Switches embedded in the system, enabling use of standard PCIe cards in standard PCIe slots.

3   Design Considerations

Below are the design considerations associated with this virtualized reference architecture which involves branch office fabric and fabric management.

- The architecture is designed to include PCIe network adapters for each of the servers to provide high availability and additional bandwidth.

- It is assumed that the organization already has a core network to act as a medium of communication between the designed solution with the Head office or any external network.

- Active Directory® Domain Services (AD DS) must be available on the network as the design involves a read-only domain controller (RODC) at the branch for the local authentication.

- The RODC should be installed and configured on a Virtual Machine before configuring the failover cluster.

- This Baseline architecture is designed for Microsoft workloads such as Exchange, Microsoft SQL server, SharePoint, etc.

- Sufficient power and cooling to support the solution must be present.

- Designed with single end to end backup solution to protect the branch office data. This comprehensive solution reduces the complexity in scheduling backup jobs and data management

4 Reference Architecture overview

This reference architecture is designed to provide a enterprise class, cost effective, virtualization infrastructure to enable remote offices, branch offices, and small businesses to rapidly and efficiently deploy and migrate workloads. The solution was also designed to integrate simple and highly-efficient infrastructure lifecycle management capability, data-protection services, and all necessary infrastructure services, such as domain services and databases.

Figure 2   Remote office / Branch office
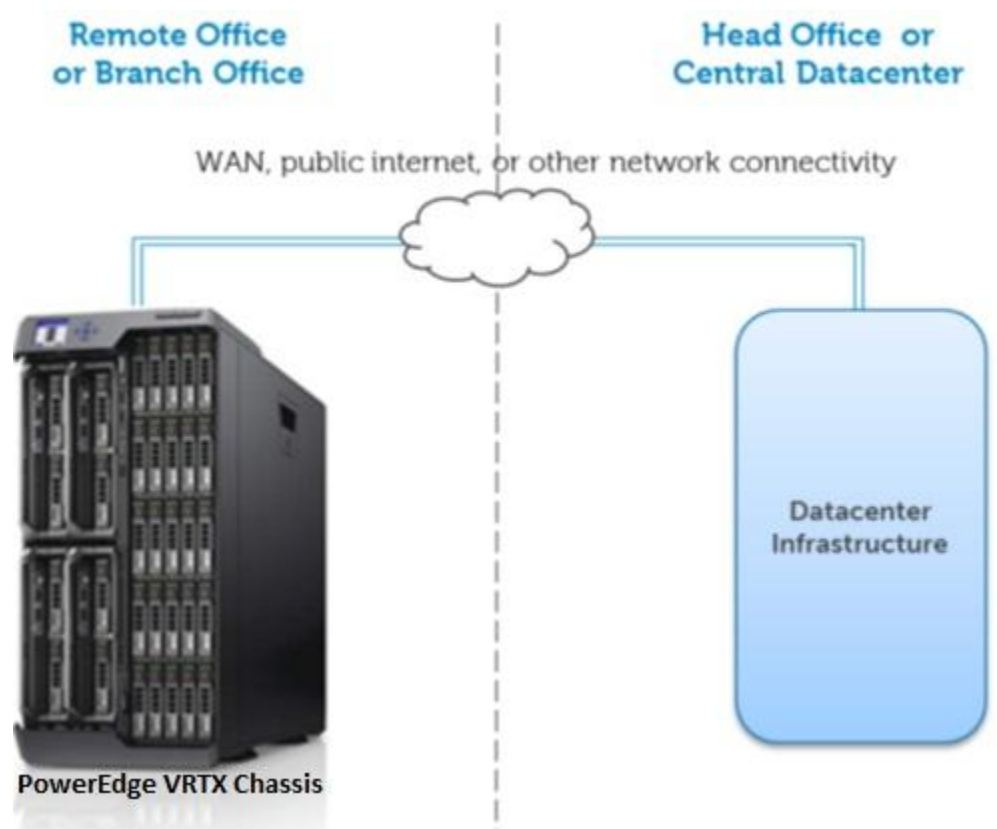
4.1 Solution Components

The reference architecture is designed to leverage the benefits of virtualization with Dell Remote office Branch office infrastructure to provide flexibility, high availability and efficiency. This section provides a high-level overview of major building blocks of the reference architecture, including Microsoft Hyper-V, Dell PowerEdge VRTX, PowerEdge M620 server, Dell Networking 5524 Switch and the management components, as illustrated in Figure 3.

Figure 3   Reference Architecture Components overview

INFRASTRUCTURE MANAGEMENT COMPONENTS
- AD-DNS
- Dell OME
- System Center 2012 SP1– Virtual machine Manager
- System Center 2012 SP1– Operations Manager
- System Center 2012 SP1– DPM

MICROSOFT WINDOWS SERVER 2012
- Failover Cluster
- Hyper-V Role Enabled
- Hyper-V and Cluster Management tools

MANAGEMENT/COMPUTE CLUSTER
- Chassis : Dell PowerEdge VRTX Chassis
- Servers : 4 X PowerEdge M620 (Cluster Nodes)
- CPU : Dual socket : 8 core per socket, 16 core total
- RAM: 128 GB (16x8GB) 1333 MHz RDIMMs
- Storage : 2 x 146GB Local HDD
- Network :Blade NDC BCM 57810 , BCM 5720 DP PCIe NIC

STORAGE
-PowerEdge VRTX Shared storage controller with 25 bay 2.5" SAS enclosures or 12 Bay 3.5" SAS enclosures

NETWORK
- Switch : 2 X Dell PowerConnect 5524
- Ports 1Gb Connectivity to match the host design

The below table 1 provides the role served by each of the above components

Table 1   Component Role

Component	Description	Role
Compute	4x PowerEdge M620 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 Switch	Network pass-through module on Fabric A of the PowerEdge VRTX chassis and Two Dell Networking 5524 top of the rack switches to provide network high availability for the infrastructure.	Support VM, Live Migration, Management and Cluster network
Microsoft Windows Server 2012	Provides Hyper-V virtualization platform with high availability	Enable Virtualization and Host High availability
Microsoft System Center 2012 SP1	System Center 2012 SP1 - 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

The section below provides an overview of some of the critical components of this solution, including Hyper-V cluster, Dell PowerEdge VRTX, and the Dell Networking 5524 switch. Readers can skip the sections on products with which they are familiar.

Microsoft Windows Server 2012:   Microsoft® Windows Server® 2012 provides the Hyper-V® virtualization platform; it helps to consolidate workloads enabling IT managers to optimally utilize 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 provides failover clustering that enables high availability and live migration of virtual machines between physical servers.

Dell PowerEdge VRTX : This is an ideal solution for small and midsize businesses as well as remote and branch offices of large enterprises, All four server nodes have access to the low-latency internal shared storage that is ideal for virtualization and clustering and a separate networking device and PCIe resources

that are shared across the compute nodes within the chassis. The PowerEdge VRTX hardware consolidates server, storage, and networking. It provides single console for easier management.

Below is the list of PowerEdge VRTX Features:

- Powerful: Up to 4 enterprise-class PowerEdge server nodes
- Small: Just 12" x 19" and about 175 lbs.
- Two form factors: Deskside or Rackable
- Plenty of storage: Up to 48TB of storage with 12 x 3.5" or 25 x 2.5" hot plug hard drives
- Fast networking: 8 1Gb Ethernet port pass-through module or optional 1GbE internal switch module
- Fully expandable input/output:8 PCIe slots (3 full height/full length and 5 low profile/half length)
- High availability and reliability: Redundant power supplies and fans
- Simple, easy-to-use management: Entry "Express" and advanced "Enterprise" levels of chassis-level management
- Office level acoustics & power: Standard 100V – 240V AC power – just plug in the wall

Dell Networking 5524 Switch: The Dell Networking 5524 Switch is designed to offer secure, fixed-port Gigabit Ethernet switching solutions which deliver full wire-speed switching performance. With 24 built-in copper Gigabit Ethernet ports in a 1U form factor, the Dell Networking 5524 Switch has a total switching capacity of up to 176 Gbps to support demanding network environments. The switches also offer simple management and scalability via a 40Gbps high-availability stacking architecture that allows managing up to eight switches from a single IP address, and sharing the dual SFP+ across the stack for uplinks to the next layer in the network. The switches are designed for Energy Efficient Ethernet (802.3az), which will reduce per port power consumption up to 50% if the link is idle and further reduction if ports are inactive.

The switch provides enterprise class features like, auto speed negotiation, flow control, port mirroring, broadcast storm control, spanning tree and rapid spanning tree, etc. The Dell Networking 5524 Switch also supports enhanced VLAN support such as Voice VLAN and Guest VLAN. The Dell Networking 5524 Switch supports enterprise level manageability features like, QoS(Quality of Service), multicast support, link aggregation, and dynamic VLAN configuration. The switch supports a web-based management interface, and also an industry-standard CLI. The switch also supports LLDP (Link Layer Discovery Protocol) which allows for troubleshooting and enhanced network management over multivendor environments. The switch also provides USB drive support for auto configuration of switches without the need of TFTP to transfer configuration files

The switch provides enterprise class features such as:

- Auto speed negotiation, flow control, port mirroring, broadcast storm control, spanning tree and rapid spanning tree,
- Enhanced VLAN support such as Voice VLAN and Guest VLAN.
- Enterprise level manageability features like, QoS, multicast support, link aggregation, and dynamic VLAN configuration.
- A web-based management interface, and also an industry-standard CLI.
- LLDP (Link Layer Discovery Protocol).

- USB drive support for auto configuration.

For more information on Dell Networking 5524 Switch, visit [the Dell Networking 5524 Switch page on Dell.com](#).

Dell PowerEdge M620: Designed for workloads, such as email, database and virtual environments, the Dell™ PowerEdge™ M620 blade server is an ideal blend of density, performance, efficiency and scalability. The M620 delivers unprecedented memory density and superb performance with no compromise on enterprise-class features.

The PowerEdge™ M620 blade server delivers exceptional memory density and broad network flexibility for large and small data centers. It dramatically boosts application performance with the latest Intel Xeon processor E5 family, developed specifically for server environments. Built with 32-nanometer process technology and up to 8 cores per processor, it enables super-fast processing for compute-intensive tasks. Take advantage of the tremendous virtualization density of the PowerEdge™ M620 with its exceptional 24-DIMM memory capacity and up to 16 cores, which give outstanding throughput in a compact, half-height blade chassis. The Dell PowerEdge M620 provides comprehensive internal storage options, including several drive types and storage controllers to choose from. The M620 supports up to 1.2TB of internal storage.

The Dell PowerEdge M620 offers balanced, scalable I/O capabilities, including integrated PCIe 3.0 capable mezzanine card slots. Dell Select Network Adapters, Dell's network daughter cards, helps to choose the right network fabric without using up a valuable mezzanine card slot. Pick the speed, technology, vendor, and other options such as switch independent partitioning, which helps to share and manage bandwidth on 10GbE connections.

Microsoft System Center 2012 SP1: System Center 2012 SP1 provides a common management toolset for the remote office infrastructure and application services. Some of these components may also be installed at the head-office or at the central datacenter but also having these components available locally allows the remote/branch office to continue operation in case of loss of network connectivity between the two sites.

Microsoft System Center 2012 SP1 is an integrated management platform that helps to efficiently manage remote datacenters and IT environments with the following components:

SCVMM - Virtual Machine Manager is a management solution for the virtualized datacenter, enables to configure and manage virtualization host, networking, and storage resources in order to create and deploy virtual machine on the remote office.

SCOM - Operations Manager provides infrastructure monitoring that is flexible and cost-effective, helps ensure the predictable performance and availability of vital applications, and offers comprehensive monitoring for the remote office infrastructure.

DPM - Data Protection Manager enables comprehensive data protection solution for the remote office data either locally or at a remote location. It also has the capability of configuring a secondary protection on a different storage to provide high availability.

Dell Open Manage Essentials: OME is utilized for discovery, inventory, and hardware level monitoring of servers, PowerEdge VRTX chassis, and Dell Networking 5524 network switches. Each of these components are configured to send SNMP traps to the centralized OME console to provide a one-to-many monitoring interface for major hardware components. OME provides a comprehensive inventory of solution component thought WS-MAN and SNMP inventory calls. For instance, reporting is available to provide server firmware versions or solution warranty status.

OME can be used as the single point of monitoring for all hardware components within an enterprise. Agent-free management ensures that processor cycles are not consumed by agent execution or intensive inventory collection. In addition, an agent-free environment enhances system security.

5   Reference Architecture

This solution consists of a PowerEdge VRTX chassis populated with PowerEdge M620 blade servers running Microsoft Hyper-V, shared storage, and networking. The figure 4 depicts high-level design of the reference architecture.

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

Figure 4   Reference Architecture for PowerEdge VRTX with Hyper-V
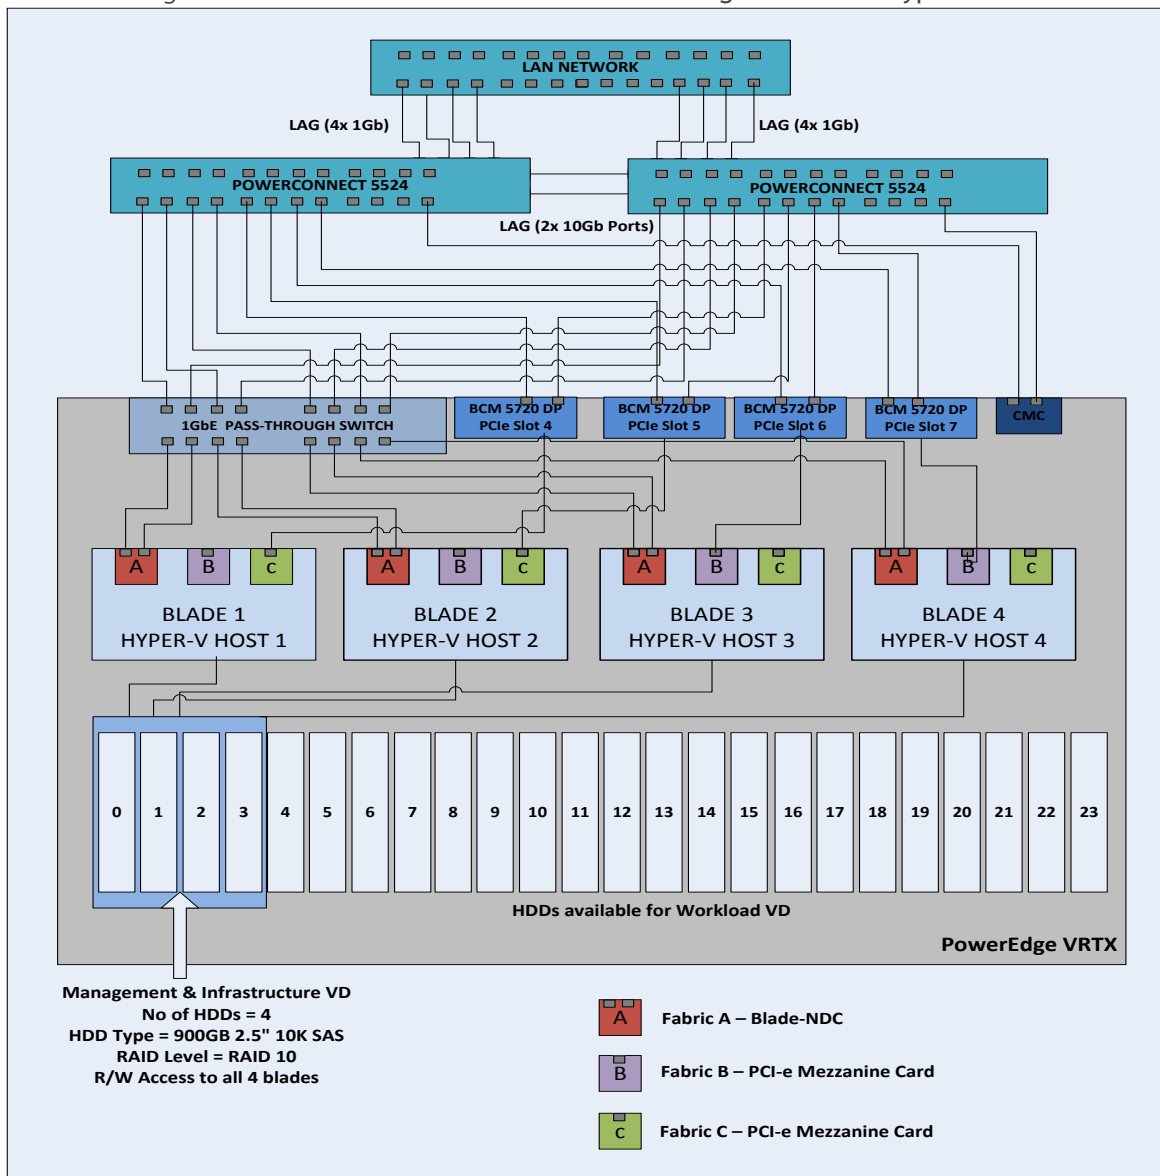
In the above design, each of Dell Blade server LOMs (LAN On Motherboard) directly maps to Fabric A and the mezzanine cards on the slot B and slot C are not used for the traditional I/O expansion via external fabrics. Instead they are used for connecting to the PCIe infrastructure (Broadcom BCM 5720 DP PCIe Card). The PowerEdge VRTX is designed to use PCIe slots. This is done via two PCIe switches on the PowerEdge VRTX Chassis – one for Fabric B and one for Fabric C. Each PCIe switch ties into one of the two mezzanine cards on the server nodes to provide connectivity to the PCIe slot as mentioned in the below table 2.

Table 2   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

As per design, the SPERC is also exposed through the Fabric B PCIe Switch which is mapped to the mezzanine slot B on all four Blade servers to access the shared storage on the PowerEdge VRTX chassis. Subsequent sections provide more detailed information on various components involved in the solution.

6 Infrastructure Architectural details

The below sections detail about the hardware and software components of branch office infrastructure which includes: server, storage, network, Hyper-V cluster.

6.1 Dell Blade Server Architecture for PowerEdge VRTX

Each PowerEdge M620 server is configured with one Broadcom 57810 network daughter card (NDC) in fabric A and one PCIe mezzanine card each in fabric B and C. Instead of using traditional I/O modules for Fabric B or Fabric C, Dell has designed the PowerEdge VRTX to use PCIe slots. This is done via two PCIe switches – one for Fabric B and one for Fabric C. Each PCIe switch ties into one of the two mezzanine cards on the server nodes

Two PCIe mezzanines are required for all server nodes to enable consistency of PCIe mapping and user experience with the system. In the below diagram Each PCIe Switch is internally connected to the PCIe slots available on the PowerEdge VRTX.

The Figure 5 illustrates how the fabrics are populated in PowerEdge M620 and how the I/O modules are utilized.

Figure 5   I/O Connectivity of M620 on PowerEdge VRTX Chassis
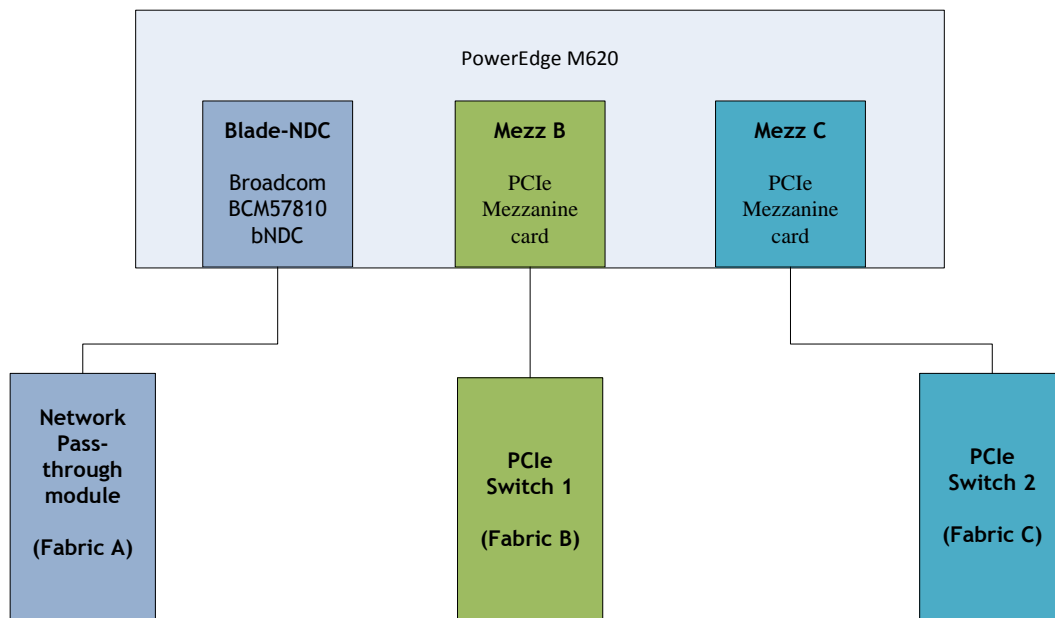
Each compute node in the cluster uses a Broadcom 57810 Blade-NDC to connect to the Fabric A (network pass-through module) and PCIe Mezzanine cards populated on the Fabric B and Fabric C to access the shared storage and the additional Network PCIe Broadcom 5720 DP card.

6.2 PowerEdge VRTX PCIe Slot Configuration and Slot Mappings

In this reference architecture, by leveraging the flexible PCIe Switch design, each server node on the chassis is mapped to a Broadcom 5720 Dual Port Adapter to provide network high availability and additional bandwidth.

The below Table provides detailed information on the PCIe slot mapping and the PCIe adapter placement.

Table 3   PCIe slot Mapping

| Blade Slot	PCIe Card	PCIe Slot
Slot 1	Broadcom 5720 DP 1Gb Base-T Adapter	Slot 4
Slot 2	Broadcom 5720 DP 1Gb Base-T Adapter	Slot 5
Slot 3	Broadcom 5720 DP 1Gb Base-T Adapter	Slot 6
Slot 4	Broadcom 5720 DP 1Gb Base-T Adapter	Slot 7

6.3 Network Architecture

Dell PowerEdge VRTX chassis has two options available for the network fabric. A Pass-Though Module or an Ethernet Switch Module, to build high availability a redundant network architecture is designed to handle network failure by mapping the host ports from pass-through network Fabric to a redundant Dell Networking 5524 switch to support Management, live migration, and cluster interconnects traffic.

Out of the two Ethernet pass-through ports associated to each blade, one port connects to one of the Networking 5524 switches and the other port connects to the other. Similarly, out of the two Broadcom 5720 NIC ports associated to each blade, one port connects to one of the Networking 5524 switches and the other port connects to the other.

These two top of the rack switches are connected to each other through an inter-switch link. Also, both switches have uplinks to the customer LAN infrastructure. Figure 6 provides an overview of how the network interfaces on each PowerEdge VRTX server is connected to the Networking 5524 network switches.

Figure 6   Network Architecture
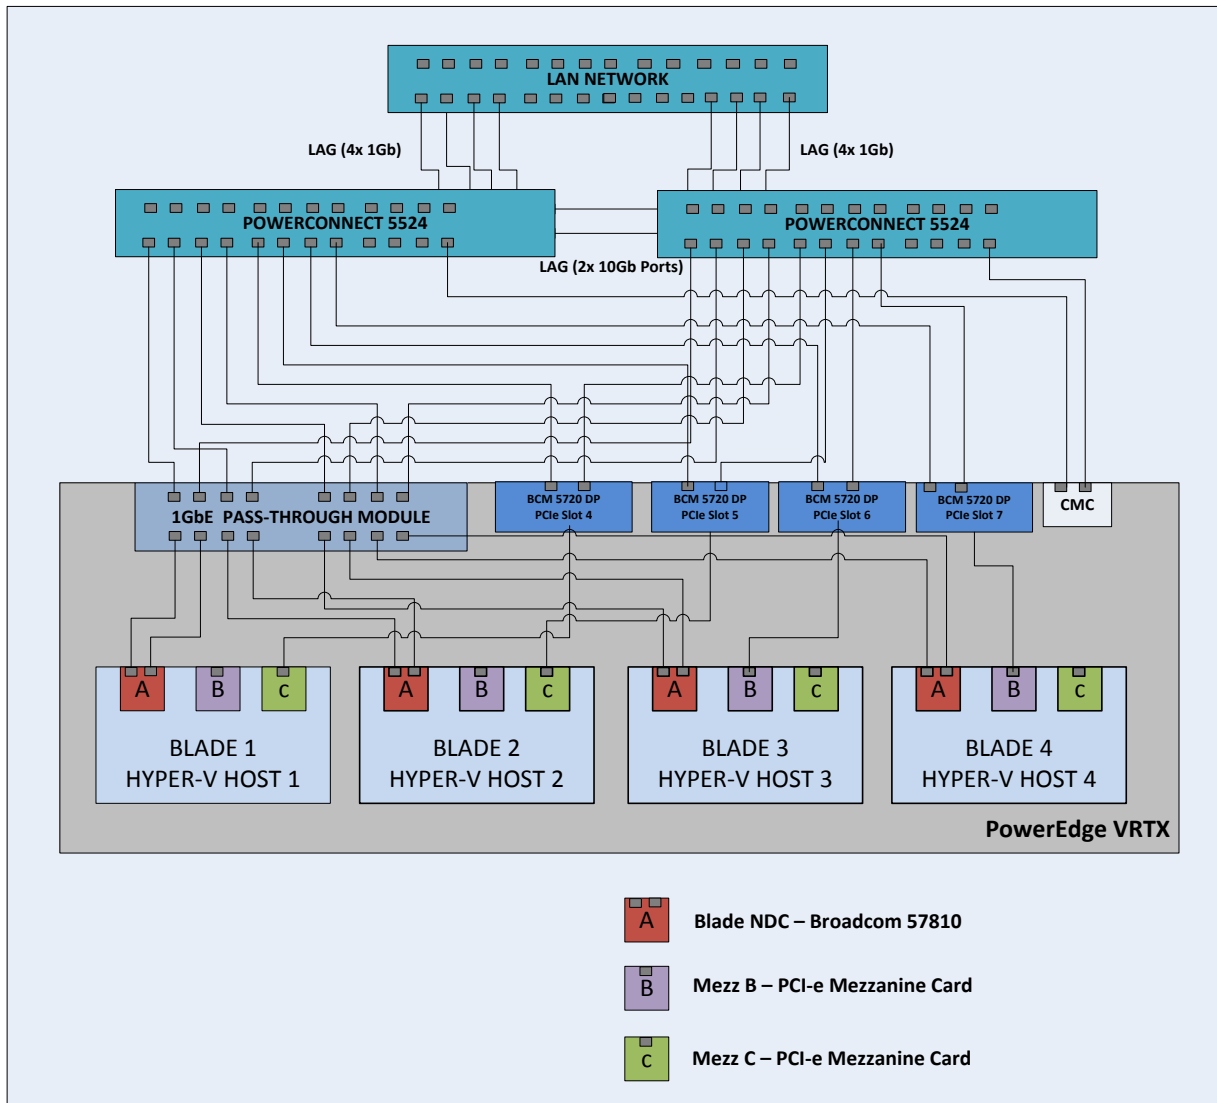
Each Dell PowerEdge M620 blade server is configured with a dual port Broadcom BCM57810 blade-NDC (running at 1GbE via 1GbE Pass-Through Module. Note that the BCM57810 blade NDC is 10GbE Capable for future network bandwidth growth) and another PCIe Network Broadcom BCM 5720 Dual Port adapter is configured to provide additional bandwidth and high availability for each the cluster node. All four network ports (two ports from the Network pass-through module and two from the PCIe adapter) are connected to the redundant Dell Networking 5524 switches placed outside of the PowerEdge VRTX enclosure.

The below section provides network configuration guidelines for setting up a redundant network configuration for the failover cluster and Hyper-V network.

6.3.1 Cluster Network design

This network design simplifies data center management by consolidating all communication (Management, Live Migration and Cluster heartbeat) onto a single fabric for ease of management by using the OS native NIC Teaming a new feature included in the Windows Server 2012 operating system that allows to team any two NICs irrespective of vendor.

For more information refer to [Windows Server 2012 NIC Teaming (LBFO) Deployment and Management guide](#)

A fault-tolerant network team is created by using the four network ports (2 Blade-NDC Ports and 2 PCIe network adapter ports) available on each of the PowerEdge M620 servers using the following settings mentioned in table 4.

Table 4   Network team settings

Network Team Settings	
Teaming Mode	Switch Independent
Load Balancing Algorithm	Hyper-V Port
Standby adapter	None

The "Switch Independent" teaming mode is used because Windows Server 2012 simplifies the configuration of the team. No special switch configuration changes are necessary. The team provides network fault tolerance as its ports are connected across switches.

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 7 helps to understand the network design and its configuration.

Figure 7   Cluster Network Design
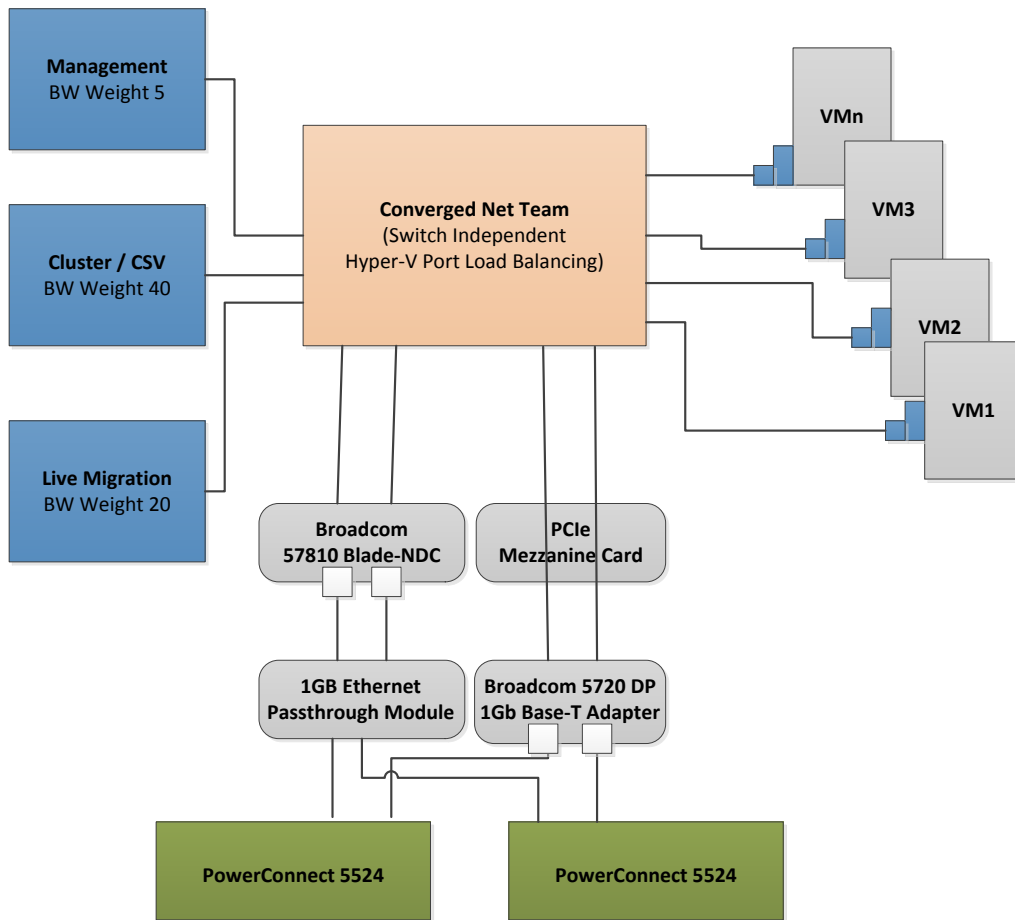
The below tables 2 describes in detail about the usage of the different traffic types associated with the Hyper-V cluster in the network architecture

Table 5   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.
Tenant 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 6 below.

Table 6   Network VLAN configuration

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

Hyper-V QoS in Windows Server 2012 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 will be 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.

The below table 7 provides a sample QoS Bandwidth allocation for the different hyper-V cluster virtual network adapters.

Table 7   Network QoS bandwidth allocation

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

6.4   Storage architecture

The PowerEdge VRTX Chassis offers an integrated, shareable storage with a new I/O Virtualization technology which 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.

It supports either 12 x 3.5" hot-swap drives or 25 x 2.5" 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 to 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 and also the ability to map a Virtual Disk to a Virtual Adapter. The VRTX Chassis does have a default assignment for the Virtual Adapters as mentioned in the below Table 8.

Table 8   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.

6.4.1   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, one for cluster Quorum configuration, the second one 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 the 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" disk PowerEdge VRTX chassis the cluster Quorum and the Management infrastructure CSV is configured on the first 2 disks with RAID 10 to provide greater protection with the configuration settings mentioned in the below table 6 . 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" PowerEdge VRTX chassis the cluster Quorum and the Management infrastructure CSV is configured on the first 4 disks with RAID 10 along with the respective settings mentioned in the table 6.

Below Figure 8 and 9 helps to understand in detail on how the storage is configured in this reference architecture.

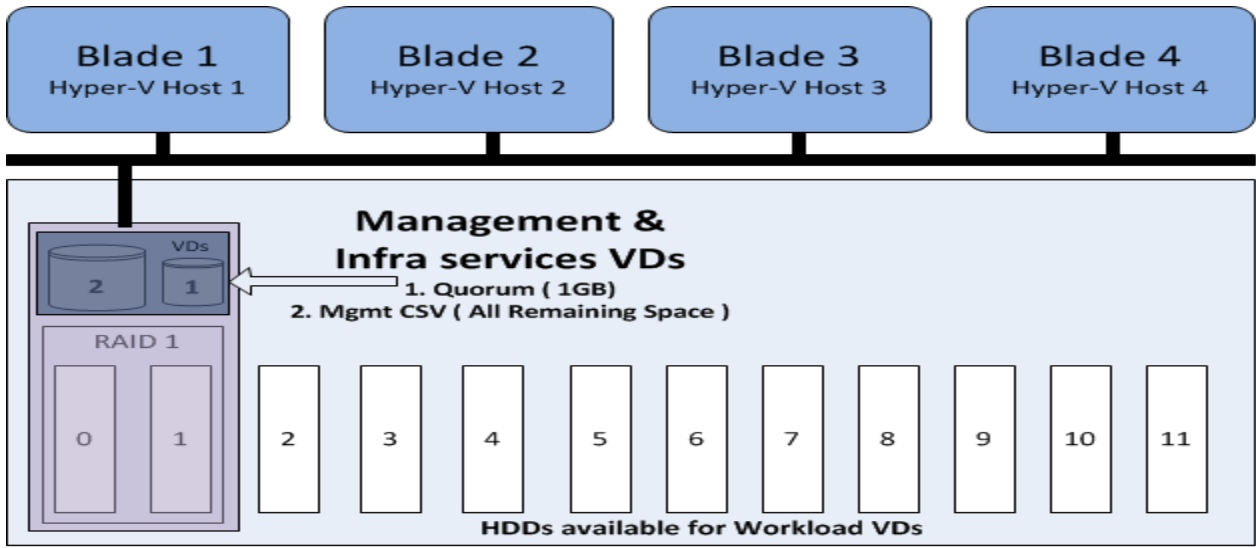Figure 8   Virtual disk configuration on 3.5" PowerEdge VRTX chassis
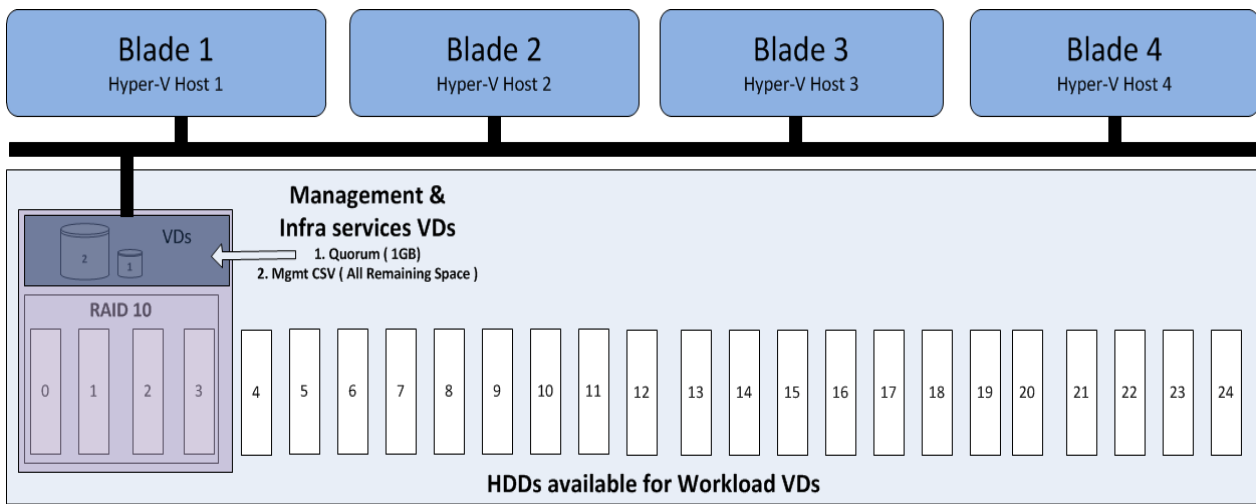

Figure 9   Virtual disk configuration on 2.5" PowerEdge VRTX chassis

The Table 9 provides detailed information on the virtual disk configuration used in this reference architecture:

Table 9   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	2 x 2TB 7.2k NL–SAS Disks (on 12 disk VRTX chassis) 4 x 900GB 10k SAS Disks (on 25 disk VRTX chassis)	RAID 1 (12 disk chassis) 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)	9 Disks * (12 disk chassis) 20 Disks * (25 disk chassis)	RAID 5 * (12 disk chassis) RAID 10* (25disk Chassis)	Depending upon the workload requirement	Adaptive Read Ahead*	Write Back*	64KB*
Global Hotspare		1 Disk					

*   Configurable as per the application best practices.

6.5   Hyper-V cluster Design

The 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. In a small infrastructure deployments where the compute environment is limited the fabric management virtual machines could be hosted on the fabric itself to better utilization and management.

In this solution a Hyper-V Failover cluster was designed for the purpose of virtualizing the common infrastructure to increases availability and enables access to applications during planned or unplanned downtime and also it is designed for better resource utilization. There are separate highly available virtual network adapters configured for the cluster heartbeat, live migration and the management as discussed earlier in this paper.

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 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 node 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 where used for the Management VMs and the Workload VMs to store the VHDs. The below figure 10 explains how the Hyper-V cluster is designed in the VRTX solution and also how the VMs are placed on the same infrastructure.

Figure 10   Hyper-V Cluster architecture
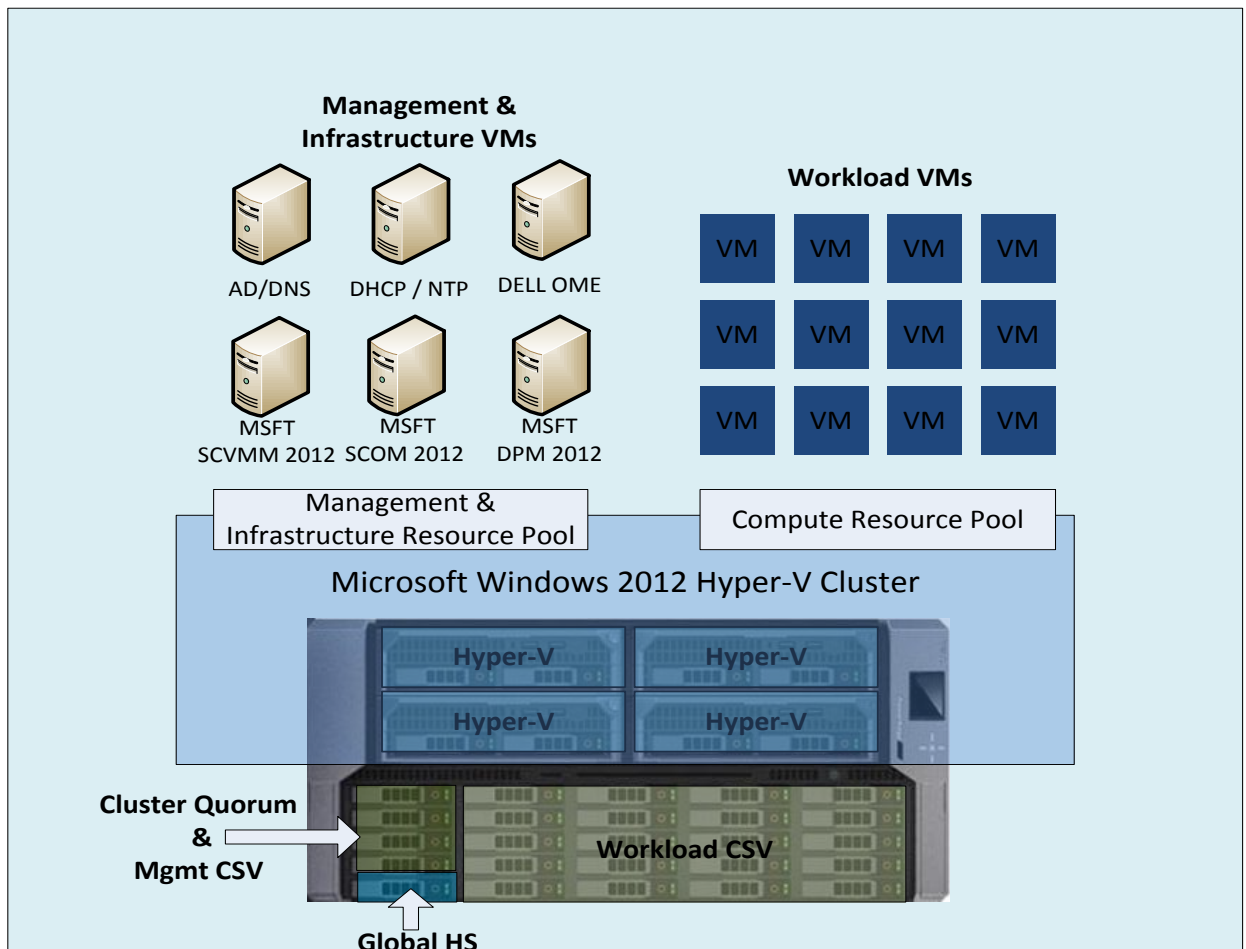
7 Management Infrastructure

The reference architecture solution is designed to include the necessary management and infrastructure services within the solution infrastructure. Each has significant role like infrastructure provisioning, management, monitoring and data protection for the branch office. Additionally, the infrastructure services components, like Active Directory, Domain Controller or Read-Only Domain Controller, DHCP server, NTP server, etc. can be run on the same infrastructure along with the management components

Each of the virtual machines is sized according to the product best practices and the requirements of the solution. The guidelines are provided for virtual processors (vCPU), Memory (vRAM), and Disk configuration

The below figure 11 depicts the Management components used in this remote office/branch office reference architecture.

Figure 11   Management Infrastructure design
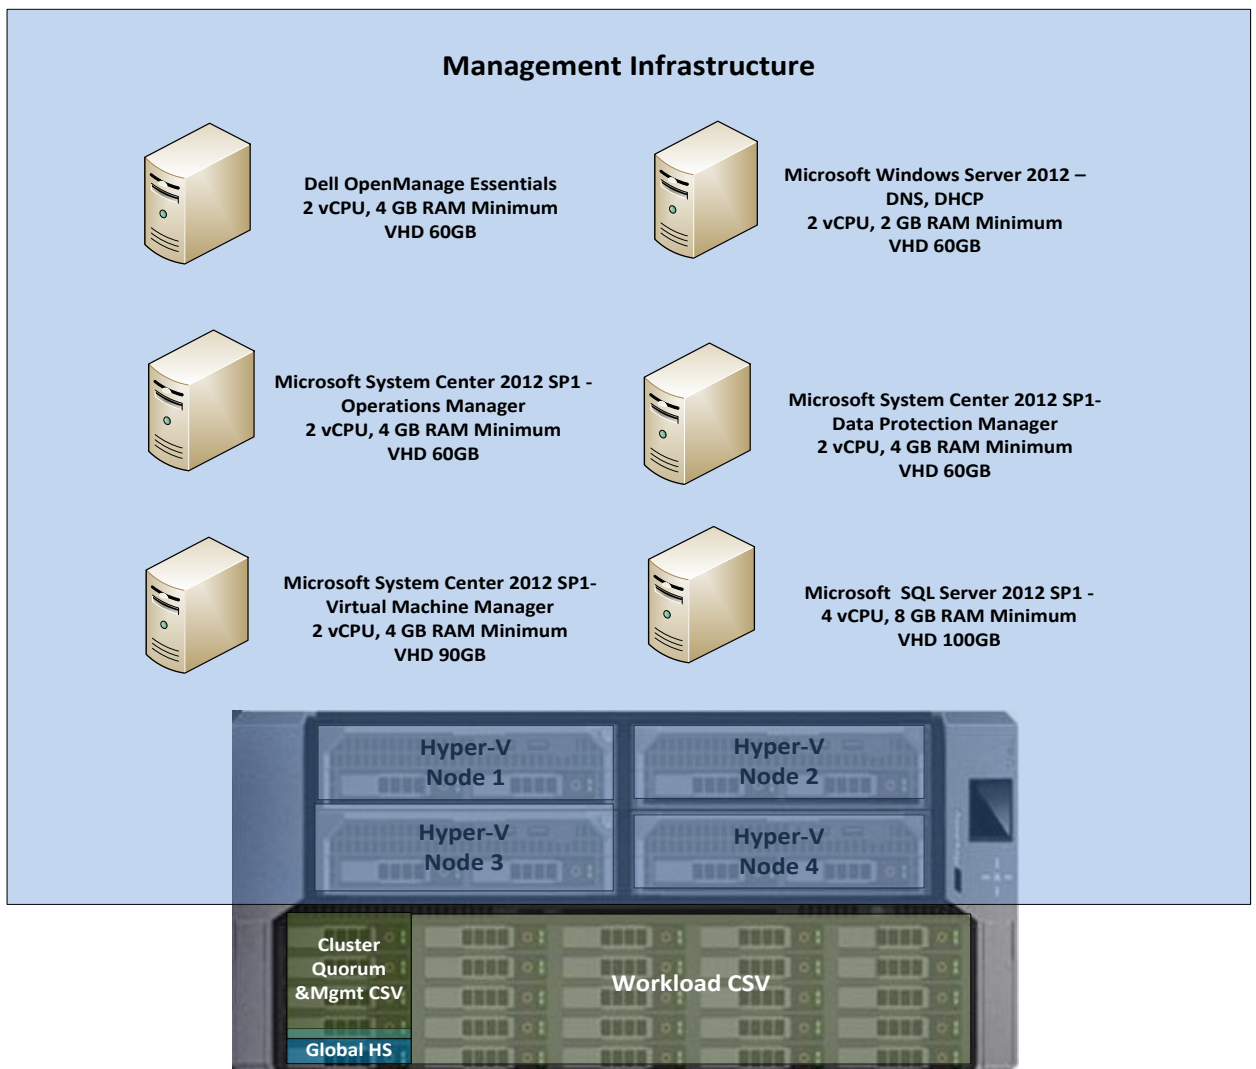
The above mentioned fabric management also shares the same cluster resource infrastructure designed for the workload VMs on the 4 node Hyper-V cluster on the PowerEdge VRTX chassis with a dedicated shared storage volume configured CSV and the redundant network connections. This architecture provides high availability for the management components hosted.

Apart from the above mentioned branch office fabric management software, PowerEdge VRTX Chassis has an embedded management component called CMC to manage and monitor each and every component of the chassis. Detailed information about each of the management components and its capabilities are discussed below.

7.1   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.
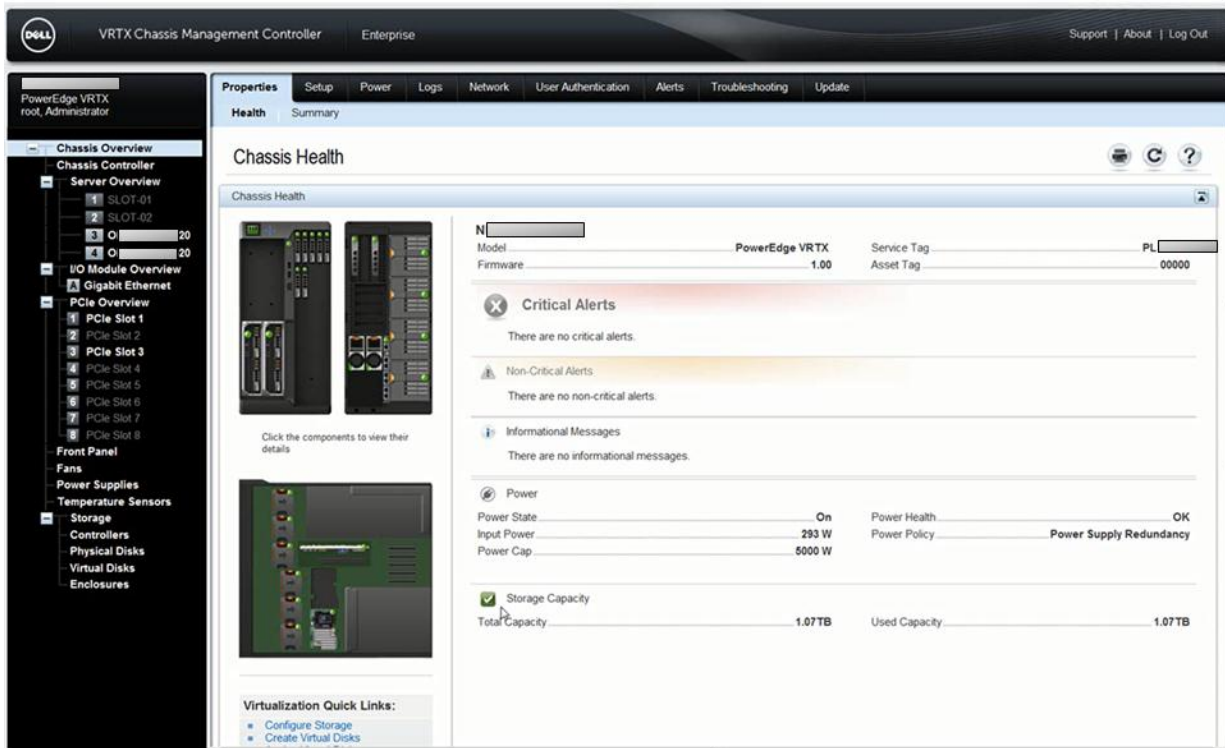
CMC enables an IT administrator to:

- 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.

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.

7.2 Dell OpenManage Essentials

Dell OpenManage Essentials is a hardware management application that provides a comprehensive view of Dell PowerEdge VRTX, devices, and components in the enterprise's network. 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.

It helps to maximize IT performance and uptime with capabilities such as:

- Automated discovery, inventory and monitoring of Dell PowerEdge servers, EqualLogic and PowerVault storage and Dell Networking Switch
- Agent-free server monitoring as well as BIOS, firmware and driver updates for Dell PowerEdge servers, blade systems and internal storage
- Control of PowerEdge servers within Windows, Linux®, VMware® and Hyper-V™ environments

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.

OME GeoView: The GeoView 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.

7.3 Read only domain controller (RODC)

If a Branch office users had to authenticate with a domain controller over a wide area network (WAN), branch offices often cannot provide the adequate physical security that is required for a writable domain

controller. Furthermore, branch offices often have poor network bandwidth when they are connected to a hub site. This can increase the amount of time that is required to log on. It can also hamper access to network resources.

This Reference architecture addresses the above problem by deploying a Microsoft Read Only Domain Controller on the branch office. As a result, users in this situation can receive the following benefits:

- Improved security
- Faster logon times
- More efficient access to resources on the network

One way in which the RODC improves the security for the AD environment is by providing a read-only copy of the Active Directory database. This means that no user or administrator can make changes to the Active Directory database from a read-only domain controller.

But the Branch office user can connect to an RODC to read almost all of the same information just like a writeable DC, but will not be able to perform any write operations without connecting to a writeable DC.

Figure 12   Branch office user Authentication using RODC
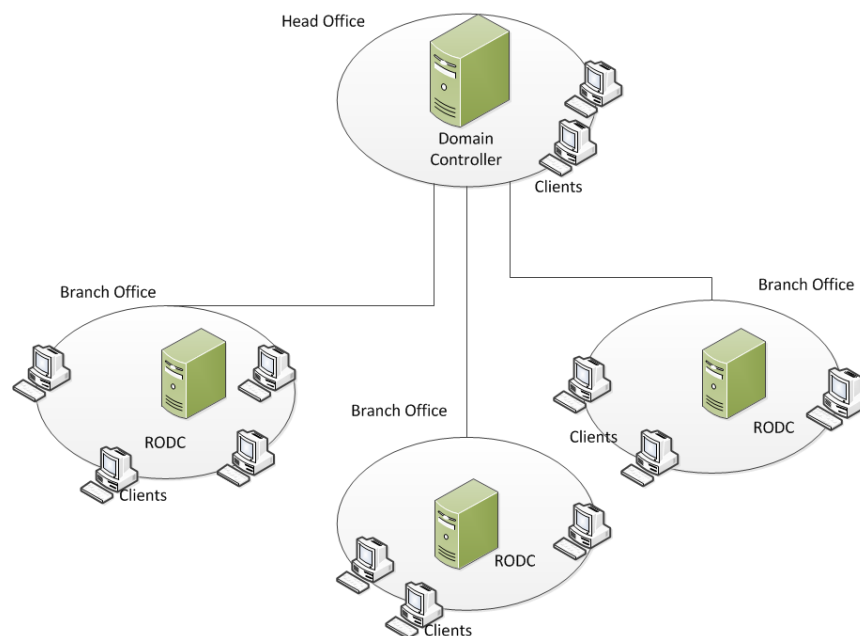
7.4   System Center 2012 virtual machine manager

The Microsoft System Center Virtual Machine Manager 2012 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 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.

7.5   System Center 2012 Operations manager

Operations Manager, a component of System Center 2012, enables to monitor services, devices, and operations for remote office. IT administrators can gain rapid insight into the state of the IT environment and the IT services running across different systems and workloads by using numerous views that show state, health, and performance information, as well as alerts generated for availability, performance, configuration and security situations.

It provides infrastructure monitoring that is flexible and cost-effective, helps ensure the predictable performance and availability of vital applications, and offers comprehensive monitoring for the datacenter. Operations manager can be extended by importing Dell Server Management Pack to integrate with Microsoft System Center Operations Manager 2012 and deliver new features for monitoring and managing Dell PowerEdge servers. Dell Server Management Pack Suite offers agent-free out-of-band monitoring for System Center 2012 Operations Manager leveraging Dell PowerEdge 12th generation embedded server management feature, iDRAC7 with Lifecycle Controller.

7.6   System Center 2012 Data protection manager

Managing the branch office systems and maximizing data protection at Remote Office / Branch Office is a critical portion of the branch office solution. As the branch office infrastructure is limited and it not feasible to have a dedicated backup server on the remote location, deploying the data protection solution on a virtual machine is an ideal solution to isolate the data backup from a given branch office . Microsoft designed its Data Protection Manager (DPM) backup software to protect branch offices over WAN as well as local servers. The DPM agent software is installed on each protected server; to captures the changes made to the file system

It is also a good practice that backup software at the remote office is capable of complete backup, including bare metal backup. Everything except for the bare metal backup is then backed up to the secondary protection configured on a different storage. In this backup solutions for branch office, Microsoft Data Protection Manager (DPM) 2012 provides the flexibility to manage systems and data at remote offices from a central location and also capable of managing the backups locally.

Below is the list of features which Microsoft System Center Data Protection Manager 2012 offers:

- Disk-based data protection and recovery.
- Command-line scripting using Windows PowerShell.
- Enterprise deployment methods for distributing the DPM agent.
- Enterprise monitoring with Operations Manager.
- Tape-based backup and archive solutions.
- Disaster recovery solutions, which provides bare-metal recovery of servers running Windows.

DPM provides protection of the following items:

- File data from volumes, shares, and folders.
- Application data, such as Microsoft Exchange Server storage groups, Microsoft SQL Server databases, Windows SharePoint Services farms, and Microsoft Virtual Server and its virtual machines.
- Data on clustered servers.
- System state for protected file and application servers.

Microsoft DPM provides centralized backup of branch offices, effectively making the company's primary data center its disaster recovery site. It can back up the DPM database to tape, or can use a second DPM server in a geographically separated location to protect the primary DPM server.it can also restore data to protected computers directly from the secondary DPM server. The secondary DPM server can also protect computers until the primary DPM server is brought back online.

It can be time-consuming process to manage multiple Remote office Backup servers since one must move from one server to another to perform various management and maintenance tasks. However, with Data Protection Manager (DPM), it is easy to manage all DPM servers in the network from a single location.

The Central Console management pack on the Operations Manager console can begin monitoring and managing the DPM servers located on various geographic locations.

The important features of the Central Console are:

- Centralized monitoring of DPM servers across different versions of DPM
- Remote administration
- Role-based access control
- Remote recovery
- Remote corrective actions
- Service level agreement (SLA)-based alerting: Alerts are generated only when an SLA is broken
- Alert consolidation
- Support for scripting repetitive DPM jobs

The centralized management solution can monitor all DPM servers from a single location. In the Central Console, it can monitor the health of the various DPM resources like DPM server, protected computers, tape libraries, available disk space, and more. The Central Console also tracks the various tasks in DPM, like whether recovery points are being taken at the scheduled times, or whether a server is still on the network.

For a smaller setup with about five to ten DPM servers, use Remote Administration to manage the DPM server centrally. Remote administration is basically the DPM Administrator Console on the computer to connect and work on any DPM server.

A   Terminology

CMC: Dell Chassis Management Controller for Dell PowerEdge VRTX

ROBO: Remote Office Branch Office

GbE: Gigabit Ethernet

HA: High Availability

HDD: Hard Disk Drive

iDRAC: Integrated Dell Remote Access Controller

KVM: Keyboard, Video, and Mouse

LAN: Local Area Network

OME: Dell OpenManage Essentials.

PERC: Dell PowerEdge RAID Controller

RAID: Redundant Array of Independent Disks

VM: Virtual Machine

RODC: Read Only Domain Controller

SCVMM: System Center Virtual Machine Manager

SCOM: System Center Operations Manager

DPM: Data Protection Manager

SLA: Service level agreement

CSV: Cluster Shared Volumes

LAN: Local Area Network

B   Additional resources

Below are some helpful resources.

Recommended Dell publications:

- Dell PowerEdge VRTX Product Page
- Dell PowerConnect 5524 Manuals and Documents
- Dell OpenManage Essentials

Recommended Microsoft publications:

- Microsoft Windows Server 2012
- Creating a Windows Server 2012 Failover Cluster
- Microsoft System Center 2012 Virtual Machine Manager Deployment Guide
- Microsoft System Center 2012 SP1 Operations Manager
- Microsoft System Center 2012 SP1 Data Protection Manager
- Microsoft Server 2012 Active Directory Read-Only Domain Controller