Remote and Branch Office Reference Architecture for VMware vSphere with Dell PowerEdge VRTX

Dell Engineering
April 2014

A Dell Reference Architecture

Revisions

Date	Description	Authors
June 2013	Initial release	Pranav Parekh
April 2014	Version 2 Updated to include new components and revisions.	Ed Blazek, Curtis Bunch, Mike Matthews

Table of contents

1 Introduction

This white paper describes a reference architecture for virtualization, based on the state of the art Dell PowerEdge VRTX system, Dell PowerEdge M820 blade servers, Dell Networking N3000 switches and VMware vSphere Hypervisor. While the architecture defined in this document is targeted at remote offices, branch offices or small businesses. This architecture has been designed and validated by Dell Network Engineering.

The remote and branch offices of today require infrastructure with enterprise class features to support virtualization and to run high-end applications. At the same time, simplified and efficient infrastructure management capabilities that suit the skill-sets and resources available at these remote and branch locations are essential. This white paper provides reference architecture for virtualization that will meet the needs of remote, branch and small businesses. It also describes validated use-cases for management services integration and data protection to enable the customer to leverage the full potential of virtualization.

This reference architecture is designed to provide a virtualization infrastructure based on VMware vSphere. Dell PowerEdge VRTX provides enterprise class computing, integrated shared storage for virtualization clusters, flexible network interfaces and a single management interface. Dell Networking N3000 series switches are used in the architecture documented in this paper. The Dell Networking 3000 series switches are 1GB switches with two 10Gb uplinks. In this paper the 10Gb uplinks connect the Dell PowerEdge VRTX to the LAN and the 1Gb ports are reserved to be used for access ports for branch office workstations, printers etc.

The white paper provides recommended settings to integrate within the infrastructure recommended lifecycle management components for infrastructure management, cloud connectivity, and, data protection. These components include VMware vCenter Server, Dell Management Plug-in for VMware vCenter, Dell OpenManage Essentials, VMware vCloud Connector, and Quest vRanger from Dell. The white paper goes on to discuss backup and disaster recovery scenarios – essential for any remote or branch office.

The extensive design and engineering work put into this solution allows customers to quickly and confidently deploy this architecture into production environments, thereby helping to eliminate costly and time consuming trial-and-error work often encountered during complex deployments. Before purchasing components, the information provided in this white paper will aid customers in sizing their solution, selecting the appropriate license levels, planning for appropriate use-cases and preparing for the deployment. After purchasing components, the white paper aids with setup, deployment and configuration of the infrastructure.

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 small and medium business can use this document to understand the design elements, hardware and software components, and the overall architecture of the solution.

2 Overview

This reference architecture is designed to provide an enterprise class, cost effective, virtualization infrastructure for remote offices, branch offices and small businesses. The solution is designed to integrate simple, yet highly efficient lifecycle management, cloud connectivity, data protection services and other necessary infrastructure services like domain services, databases, etc. The solution is envisioned to be a complete infrastructure allowing efficient deployment and migration workloads.
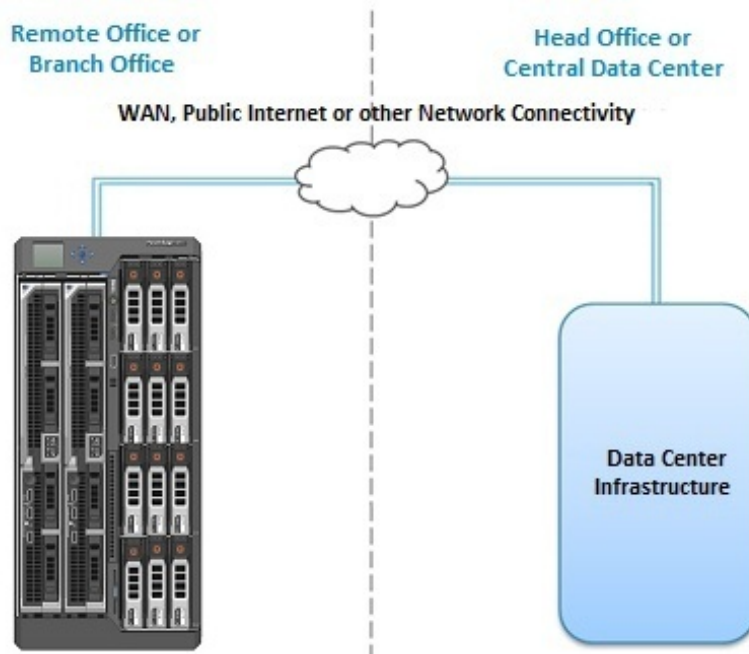
Figure 1   Remote Office / Branch Office

2.1 Solution Capabilities and Use Cases

The reference architecture is designed to enable the following capabilities and use cases:

- **Enterprise class infrastructure:** The reference architecture is based on the Dell PowerEdge VRTX platform, with Dell PowerEdge blade servers, which include best-in-class embedded server management capabilities, shared storage and network I/O ports.
    - **Dense compute:** The solution uses two 12th generation Dell PowerEdge M820 blade servers. Each blade server supports up to 1440 GB of RAM (48 x 32 GB DIMMs).
    - **10 GbE networking:** The solution uses Broadcom Gigabit network interface cards in the M820 Blade Servers, Dell Networking 10GB I/O Module and external N3000 switches to provide networking for the LAN traffic. The design provides sufficient bandwidth for workload and management traffic, with room to expand the bandwidth for future growth.

- **Shared storage:** The integrated shared storage within PowerEdge VRTX enables virtualization with enterprise class features. Having integrated shared storage allows a cost effective, simple and fast way of creating shared data stores without the cost and complexity of a storage area network.

- **Embedded systems management:** The solution includes Dell Chassis Management Controller (CMC), Dell iDRAC (Integrated Dell Remote Access Controller), and Dell Lifecycle Controller to provide enterprise class out-of-band management capabilities and deep-level hardware monitoring functionalities.

The PowerEdge VRTX combines servers, shared storage and networking interfaces in a small form-factor. This shared infrastructure design in a small physical form-factor not only enables cost savings by consolidating power and cooling but also enables rapid deployment and configuration capabilities that provide fast time-to-value. These unique capabilities and the small physical form-factor makes the PowerEdge VRTX perfectly suited for remote and branch offices.

- **Virtualization with VMware vSphere:** The solution is designed to enable a VMware vSphere cluster comprising of two server nodes, and will allow rapid provisioning and migration of customer-workloads in virtual machines (VMs). The design also supports VMware vMotion, vSphere High Availability (HA), vSphere Distributed Resource Scheduler (DRS) and other features with data center-like capabilities.

  Virtualization can be a critical key feature for any remote or branch office. The advantages of virtualization include resource consolidation, increased energy efficiency, better business continuity, high availability, cloud enablement, etc. These advantages are critical for remote and branch offices that rely on limited physical compute capacity, limited real estate for the infrastructure, and often require moving workloads between the remote or branch offices and a head office or a central data center.

- **Integrated management:** Dell Chassis Management Controller (CMC) is leveraged to enable a single management interface for chassis infrastructure management. It also integrates all necessary management services for infrastructure management, virtualization management, cloud management and data protection.

  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 also 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, PCIe devices, etc.

- **Data protection:** This reference architecture is designed to support and integrate Quest vRanger within the PowerEdge VRTX based virtualization infrastructure. It enables backup of workload VMs, replication of VMs, failover capabilities and data recovery from head office. Recommendations for configuring these components at the remote/branch office, and at a head-office or central data center are provided.

Data protection is a critical requirement for any remote or branch office. A solution that provides the ability to backup entire virtual machines and to recover data in the case of a failure is necessary for business continuity. Organizations with multiple remote and branch offices often require replicating critical VMs at multiple sites enabling a quick recovery from a failure at a remote or branch site.

- **Cloud connectivity:** The design supports the necessary components and services to enable VMware vCloud connectivity. This capability facilitates rapid enablement of connectivity to a private or public cloud.

  A remote or branch office often requires migrating workloads to another remote or branch office of the same organization, to a head office, or to a central data center. This capability enables a remote or branch office to implement strategies for optimizing the geographical locality of different workloads for network bandwidth, hardware availability, performance, etc. This reference architecture is designed to support this critical capability.

2.2 Solution Components

This section provides a high-level overview of the major building blocks, i.e., hardware components, software components, data protection and cloud connectivity components. As described in the previous section, the solution is based on Dell PowerEdge VRTX system with Dell PowerEdge M820 blade servers, Dell Networking N3000 switches and VMware vSphere for virtualization.

The following sections provide an overview of the critical components (Figure 2) of this reference architecture.

VMware vSphere
- ESXi 5.x
- High Availability Cluster
- Management with VMware vCenter Server

Dell PowerEdge VRTX
- State-of-the-art chassis infrastructure
- 12th generation Dell PowerEdge M820 blade servers
- Shared storage within the chassis
- 1GbE Ethernet switch, 1GbE Ethernet pass-through or 10GBE Ethernet switch
- Support for up to eight PCIe devices
- Flex Addressing
- CMC and KVM for enclosure management

Dell Networking N3000 Switches
- Layer 3 switching at Gigabit Ethernet bandwidth
- Two 10Gb expansion ports
- Two integrated 10GbE SFP+ ports
- Two GbE combo media ports for copper or fiber flexibility
- Enterprise class management and security
- 12 unit stacking support

Quest vRanger from Dell
- Backup
- Replication
- Disaster recovery

VMware vCloud Connector
- Cloud connectivity
- Public or hybrid cloud

Management Components
- VMware vCenter Server
- Dell Management Plug-in for VMware vCenter
- Dell OpenManage Essentials

Figure 2   Reference Architecture Components Overview

The key components used in the architecture are described in Table 1.

Table 1   Solution Components

Component	Details
Virtualization	
Hypervisor	VMware vSphere 5.x
Dell PowerEdge VRTX Chassis	
Hypervisor Host Servers	2x Dell PowerEdge M820 Each server with: 2x Intel Xeon E5-4650 2.7Ghz Memory 230GB DDR-3 (This CPU and memory configurations are provided for guidance. Other supported CPU and memory configurations can also be used.) 2x Broadcom 57810-k Network Daughter Cards
Networking	Dell PowerEdge VRTX 10Gb R1-2210 switch module
PCIe Devices	2x Broadcom 57810 dual port PCIe NICs
Storage	Shared storage within PowerEdge VRTX Up to 25x 2.5" HDDs/SSDs or up to 12x 3.5' HDDs/SSDs
External Network Switches	
Switches	2x Dell Networking N3048 switches
Management Services	
Management components	VMware vCenter Server Dell Management Plug-in for VMware vCenter Dell OpenManage Essentials VMware vCloud Connector Quest vRanger from Dell

2.2.1   VMware vSphere

VMware vSphere 5.x includes the ESXi hypervisor as well as vCenter Server that is used to configure and manage VMware hosts. Key capabilities for the ESXi Enterprise Plus license level include:

- **VMware vMotion:** VMware vMotion technology provides real-time migration of running virtual machines (VM) from one host to another with no disruption or downtime.

- **VMware High Availability (HA):** VMware HA provides high availability at the virtual machine (VM) level. Upon host failure, VMware HA automatically re-starts VMs on other physical hosts running ESXi. VMware vSphere 5.x uses Fault Domain Manager (FDM) for High Availability.

- **VMware Distributed Resource Scheduler (DRS) and VMware Distributed Power Management (DPM):** VMware DRS technology enables vMotion to automatically achieve load balancing according to resource requirements. When VMs in a DRS cluster need fewer resources, such as during nights and weekends, DPM consolidates workloads onto fewer hosts and powers off the rest to reduce power consumption.

- **VMware vCenter Update Manager:** VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.

- **VMware Storage vMotion:** VMware Storage vMotion enables real-time migration of running VM disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.

- **Host Profiles:** Host Profiles standardize and simplify the deployment and management of VMware ESXi host configurations. They capture and store validated configuration information, including host compliance, networking, storage and security settings.

In the upcoming release of ESXi 5.5, VMware vSphere server is replaced with the vCenter Appliance (vCSA). The vCSA is a preconfigured Linux-based virtual machine optimized for running vCenter Server and associated services. vCVA 5.5 eliminates the requirement of a dedicated Windows Server license. It also includes the software needed to provide the following functionality:

1. vCenter Single Sign-On
2. vSphere Web Client
3. vCenter Inventory Service
4. vCenter Server

In the architecture referenced in this document, vCSA 5.5 was utilized with the default configuration settings shown in (Table 2).

Table 2   vCSA Default Configuration

Component	Default Configuration
CPU	2
Memory	8 GB
Hard Disk 1	25 GB (Thin Provisioned)
Hard Disk 2	100 GB (Thin Provisioned)
SCSI Controller 0	LSI Logic Parallel
Network Adapter 1 type	VMXNET 3

Note: At the time this document is being updated Dell's version of ESXi 5.5 has not been released; the information contained in the preceding section is based on a pre-release version of 5.5. Some changes made in the final version of the software may affect the information contained in this document.

For more information on VMware vSphere, refer to: www.vmware.com/products/vsphere.

2.2.2   Dell PowerEdge VRTX

This section provides a high-level overview of the relevant capabilities and features of PowerEdge VRTX. For detailed information on PowerEdge VRTX platform architecture, features, and capabilities, please refer to the PowerEdge VRTX page on Dell.com and the PowerEdge VRTX manuals.

Chassis Enclosure: Dell PowerEdge VRTX integrates servers, external network ports, external PCIe slots, and shared storage integrated within a chassis. The chassis also provides the consolidated power and cooling infrastructure for the servers, shared storage and the other components. Table 3 provides a summary of the characteristics of Dell PowerEdge VRTX.

Table 3   PowerEdge VRTX System Overview

Feature	Description
Server Nodes	Up to four server nodes
Supported Servers	Dell PowerEdge M520, M620 and M820
I/O Modules	1 GbE switch, 1 GbE pass-through or 10 GbE switch
External PCIe Slots	5 x low-profile Gen2 PCIe slots and 3 x full-height Gen2 PCIe slots
Shared Storage	Up to 12 x 3.5" SAS HDDs/SSDs or up to 25 x 2.5" SAS HDDs/SSDs
Raid Controllers	1 or 2 (redundant) Dell PowerEdge RAID Controllers
Management	1 or 2 (redundant) Dell Chassis Management Controllers (CMC)
Power Supplies	Up to 4 x 1100W PSUs
Form Factor / Placement	Standalone "Tower" (vertical position) or 5U Rack unit (horizontal position)

Figure 3   PowerEdge VRTX System Overview

PowerEdge VRTX also includes a front LCD panel for convenient chassis management, a Keyboard-Video-Mouse (KVM) port that can be assigned to a server, a DVD drive that can be assigned to a server, redundant power supply units and fans. Figure 3 provides an overview of the chassis form-factor and the main components.

PowerEdge VRTX uses a PCIe switch domain within the chassis to assign the external PCIe slots to one or more servers and to use the storage in a shared fashion. PowerEdge VRTX allows up to eight PCIe devices and the flexibility to assign them to any server within the chassis. This design also allows sharing of the storage across multiple servers. When a PCIe slot is assigned to a server, the server will appear to have the PCIe device in that slot as a local PCIe device. Similarly, when the storage is shared with a server, the server will appear to have direct attached storage.

Fabrics: The PowerEdge VRTX chassis has three separate fabrics referred to as A, B and C. Fabric A refers to the Ethernet connectivity provided to the server modules by the I/O module installed in the VRTX enclosure. Fabrics B and C refer to the PCIe connections between the Broadcom 5810 dual port PCIe cards installed in the server modules and the VRTX enclosure.

I/O Module: PowerEdge VRTX supports up to eight external Ethernet connections through a single I/O module plugged in the back of the system. The I/O module provides connectivity on Fabric A of the

system. The internal ports on the I/O modules connect to the network daughter card (NDC) in each PowerEdge M520, M620 or M820 blade server. PowerEdge VRTX system supports R1-2210 10Gb switch module, R1-2401 1Gb switch module, and R1-PT 1 Gb pass-through module in the I/O module slot.

PCIe Devices: PowerEdge VRTX has eight PCIe slots on the back of the chassis. Out of the eight slots, five slots support low-profile form factor, while the remaining three slots support full-height form factor. Each PCIe slot is associated to either Fabric B or C in the chassis. The user can map the PCIe slots to s via the CMC. A PCIe slot in Fabric B connects to the Fabric B of the server to which it is mapped, while a PCIe slot in Fabric C connects to the Fabric C of the server to which it is mapped. The PCIe slots are non-shared, i.e. are dedicated to one server in the system. Table 4 provides a summary of the PowerEdge VRTX PCIe slot characteristics.

Table 4   PowerEdge VRTX PCIe Slots Overview

PCIe Expansion Slot #	Associated Fabric (B or C)	Link Width	Supported Card Size
1	C	x8	Full Height, Full Length
2	C	x8	Full Height, Full Length
3	B	x8	Full Height, Full Length
4	C	x8	Low-Profile
5	C	x8	Low-Profile
6	B	x8	Low-Profile
7	B	x8	Low-Profile
8	B	x8	Low-Profile

Shared Storage: The PowerEdge VRTX chassis houses a storage enclosure that provides shared storage capability to the servers. The storage enclosure within the PowerEdge VRTX chassis provides the option of using either up to 25 x 2.5" disks or up to 12 x 3.5" disks. The PowerEdge VRTX chassis uses a Dell PERC8 controller, shared among the servers within the PowerEdge VRTX chassis, providing SAS connectivity to the disks. This solution enables creating storage volumes that can be selectively shared across some or all servers in the PowerEdge VRTX chassis.

The PERC8 controller in PowerEdge VRTX leverages SR-IOV technology to provide multiple channels of virtual I/O. Using this technology, four virtual adapters (VA) are created on the PERC8 card. The PCIe switching technology within the PowerEdge VRTX chassis allows each VA to be assigned to one of the server slots to provide the server in that slot access to the shared storage. Within each server, the PERC drivers are assigned to virtual adapters that map to the corresponding virtual drives in the storage array.

Chassis Management: Dell PowerEdge VRTX has integrated management through a redundant Chassis Management Controller (CMC) module for enclosure management and integrated KVM port. Through the CMC, the enclosure supports FlexAddress technology, which enables the enclosure to lock the Media Access Control (MAC) addresses of the Ethernet controllers to specific server slots. This enables seamless swapping or upgrading of servers without affecting the LAN configuration.

The Chassis Management Controller (CMC) for PowerEdge VRTX provides a single interface to configure, manage and monitor the entire PowerEdge VRTX system. It can be used to control the power-up and power-down operations for the chassis, servers and the I/O module. It provides an interface to configure and manage the servers, PCIe resources, the I/O module, mapping PCIe slots to servers, physical disks, RAID controller, virtual disks, mapping virtual disks, chassis power budget, etc. The CMC also provides an interface for logging faults, sending alerts, controlling the front panel, and for many other general functions related to management, configuration and monitoring of the PowerEdge VRTX chassis.

Dell PowerEdge M820 Blade servers: The PowerEdge M820 is an Intel processor-based 4-socket, full-height blade server offering:

- Intel E5-4600 series processors offer the greatest per node performance of any blade.

- Each PowerEdge M820 can deploy up to 1.5TB of RAM per server.

- Each server supports two dual-port network daughter cards (NDCs) and two PCIe mezzanine cards Broadcom 5810 dual port PCIe cards (Figure 4).

- 'Agent Free' management with the iDRAC7 with Lifecycle Controller allows customers to deploy, update, maintain, and monitor their systems throughout the system lifecycle without a software management agent, regardless of the operating system.

- Internal dual SD module to provide failover capability for embedded hypervisors.
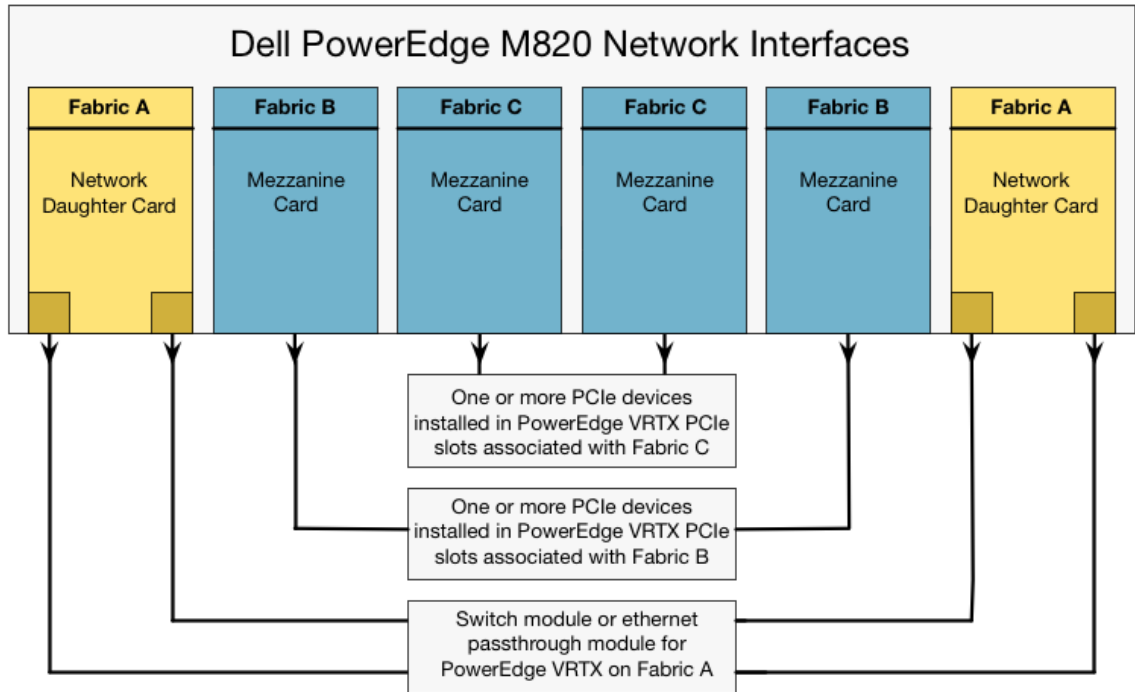
Figure 4   Network Interfaces of PowerEdge M820 in PowerEdge VRTX

Embedded Management with Dell's Lifecycle Controller: The Lifecycle Controller is the engine for advanced embedded management and is delivered as part of iDRAC Enterprise in 12th-generation Dell PowerEdge blade servers. It includes 1GB of managed and persistent storage that embeds systems management features directly on the server, thus eliminating the media-based delivery of system management tools and utilities previously needed for systems management. Embedded management includes:

- Unified Server Configurator (USC) supports local one-to-one deployment via a graphical user interface (GUI) for operating system install, updates, configuration, and for performing diagnostics on single, local servers. This eliminates the need for multiple option-ROMs for hardware configuration.

- Remote Services are standards-based interfaces that enable consoles to integrate, for example, bare-metal provisioning and one-to-many OS deployments, for servers located remotely. Dell's Lifecycle Controller takes advantage of the capabilities of both USC and Remote Services to deliver significant advancement and simplification of server deployment.

- Lifecycle Controller serviceability simplifies server re-provisioning and/or replacing failed parts, and thus reduces maintenance downtime.

For more information on Dell Lifecycle Controller, see http://content.dell.com/us/en/enterprise/dcsm-embedded-management.

2.2.3   Dell Networking N3048 Switch

The Dell Networking N3048 switch offers a secure, fixed-port, Gigabit Ethernet switching solution that delivers full wire-speed switching performance. It has 48 x 10/100/1000 Mb auto-sensing ports, two integrated 10GbE SFP+ ports, two GbE combo media ports for copper or fiber flexibility and two dedicated rear-stacking mini-SAS ports. The switch supports high throughput with 1Gbps bandwidth and 10Gbps fiber uplinks, and offers simple management and scalability via an 84Gbps high-availability stacking architecture that allows managing up to 12 switches from a single IP address, and share the dual SFP+ across the stack for uplinks to the next layer in your network. These capabilities make the Dell Networking N3048 switch a great solution for remote offices, branch offices and businesses.

The Dell Networking N3048 switch provides wire-rate performance with features like auto speed negotiation, flow control, port mirroring, broadcast storm control, loop-free redundancy without spanning tree using multi-chassis link aggregation (MLAG), etc. It also supports private VLAN extensions and private VLAN edge support.

The Dell Networking N3048 switch includes manageability features like, QoS, multicast support and link aggregation; and offers management via an intuitive and familiar CLI, embedded web server (GUI), SNMP, telnet and serial connections. It also supports LLDP (Link Layer Discovery Protocol) which allows for troubleshooting and enhanced network management over multi-vendor environments, and includes a USB port to allow auto configuration of switches through a USB drive without the need of TFTP to transfer configuration files. The Dell Networking N3048 switch offers Energy Efficient Ethernet (IEEE 802.3az) helping reduce standby power consumption and disabling the port if no cables are connected.

For more information on, visit [Dell Networking N3048 page on Dell.com](#).

3 Design Principles

The following principles are central to the design and architecture of this solution.

Remote offices, branch offices and small businesses: This reference architecture has been designed to meet the requirements of remote offices, branch offices, and small businesses. Here are some of the highlighted capabilities.

- Small form-factor for efficient use of real estate in an office environment and still provide enterprise class four-node hypervisor cluster with integrated shared storage and Gigabit Ethernet based LAN infrastructure

- A single chassis infrastructure management interface for simple yet efficient management

- Energy efficiency achieved through shared power and cooling infrastructure in PowerEdge VRTX

- Capability to monitor the PowerEdge VRTX infrastructure at different remote locations from a single map-view through Dell OME.

Virtualization: Each server is configured with appropriate processor, memory and network adapters, as required for virtualization. The solution includes integrated VMware ESXi hypervisor, shared storage, and two nodes to create a highly available cluster.

High availability: The design supports VMware vSphere HA cluster. Additionally, high availability has been incorporated in other aspects of the design including networking, power, and cooling the VRTX.

Manageability: The Dell Chassis Management Controller provides a single management interface for configuring and managing server nodes, PCIe NICs, I/O Modules, shared storage and other components of the PowerEdge VRTX chassis. The integration of enterprise management components is also part of the design.

Data protection integration: Quest vRanger is integrated as the data protection solution. Specific recommendations are provided for configuring vRanger for backup, replication and disaster recovery. Multi-site VM level replication and disaster recovery at a remote location are critical requirements for a remote or branch office. This design supports these scenarios.

Cloud enablement: The design includes VMware vCloud Connector for cloud enablement and cloud connectivity.

Flexible configurations: The solution supports additional options for server processors, server memory and disk and RAID configurations for shared storage. This flexibility is intended to enable the customer to optimize the compute and storage for any specific workload requirements, while taking advantage of this reference architecture solution.

4 Reference Architecture

This solution consists of a Dell PowerEdge VRTX chassis populated with Dell PowerEdge M820 blade servers running VMware ESXi hypervisor, shared storage and networking. Figure 5 provides the major components as well as a high-level overview of the network connectivity in the reference architecture.
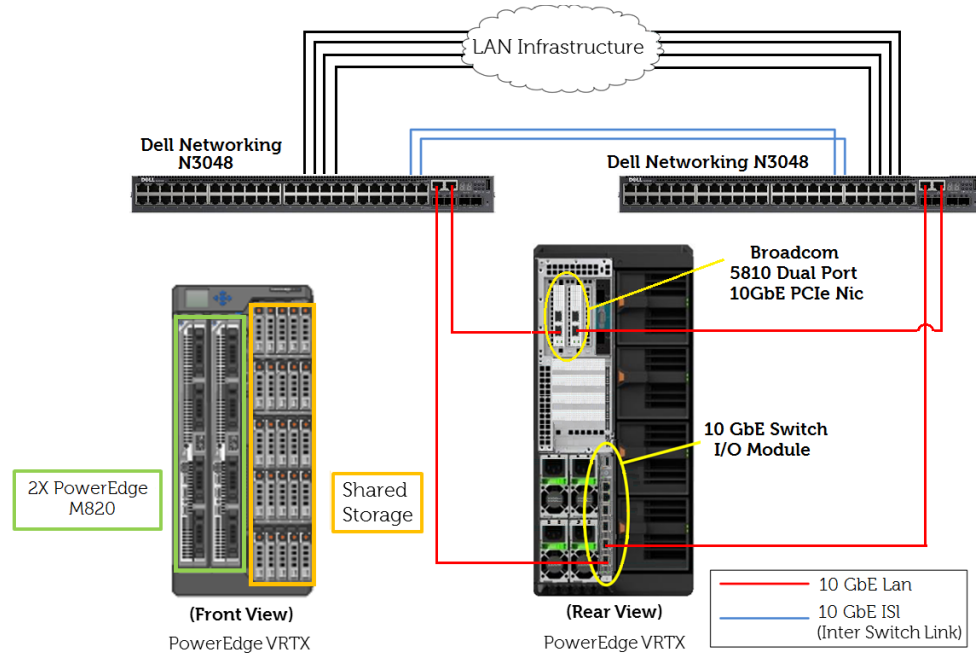
Reference Architecture Overview

The subsequent sections of this document provide additional information on network connectivity and configuration of the major components.

5 Network Architecture

This reference architecture is designed to equip each PowerEdge M820 blade server with two 10Gb Ethernet interfaces for LAN traffic. This section provides details of the design and configuration of the different physical subsystems for networking and the hypervisor configuration to setup the corresponding virtual networking.

5.1 Physical Network Architecture

This section provides details on the components of the networking subsystem and their configuration. Moreover, information on virtual networking configuration is also provided.

Server network ports: Each PowerEdge M820 blade server is configured with two Broadcom 5810 dual port network daughter cards (NDC) in Fabric A and one Broadcom 5810 dual port PCIe card in Fabric C. The solution also includes a Dell PowerEdge VRTX 10Gb R1-2210 switch module on Fabric A. This switch I/O module has 16 internal ports on Fabric A and six external ports (four 10GbE and two 1GbE) for external connectivity. All four available ports spilt across both Broadcom NDCs are connected to the corresponding internal ports of the R1-2210 10Gb switch module within the Dell PowerEdge VRTX chassis. Two Broadcom 5810 dual port PCIe NICs are installed in two full-profile PCIe slots in the PowerEdge VRTX chassis. Using the CMC, each Broadcom 5810 dual port NIC is mapped to one of the PowerEdge M820 blade servers on Fabric C. Table 5 provides the configuration of the PCIe slots in the PowerEdge VRTX chassis and their mapping to the server slots.

Table 5   PowerEdge VRTX PCIe Slots Configuration

PCIe Expansion Slot #	Fabric (B or C)	PCIe Device	Server Mapping
1	C	Broadcom 5810 Dual Port 10GB	Server-1
2	C	Broadcom 5810 Dual Port 10GB	Server-2
3	B	Empty	None
4	C	Empty	None
5	C	Empty	None
6	B	Empty	None
7	B	Empty	None
8	B	Empty	None

This physical networking configuration provides each M820 blade server with four 10Gb Ethernet ports on Fabric A through the R1-2210 10Gb switch module, as well as two 10Gb Ethernet ports on Fabric C through a PCIe Broadcom 5810 dual port NIC mapped to each server. Because of this configuration, each M820 blade server has six 10Gb Ethernet ports for LAN traffic. Figure 6 summarizes this configuration.
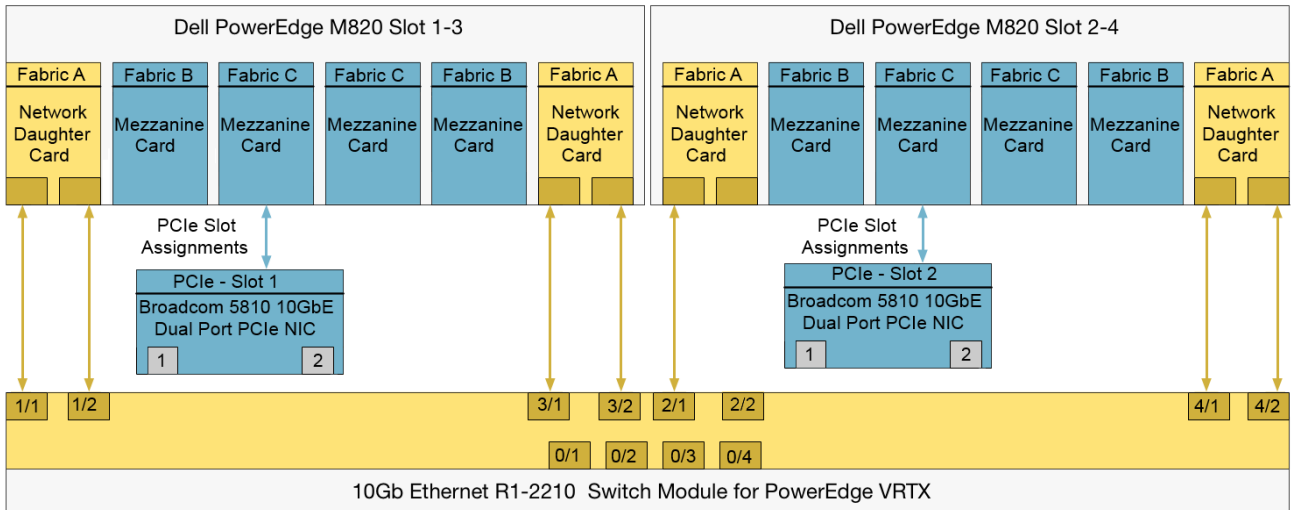
Figure 5   PowerEdge VRTX Network Interface Configuration Overview

Network Edge Connectivity: The reference architecture includes two Dell Networking N3000 switches. Each external port on the PowerEdge VRTX chassis is uplinked to one of the two Dell Networking N3000 switches. All links in this reference architecture are configured as trunk ports carrying the appropriate VLAN IDs. Figure 7 provides an overview of how the network interfaces on each PowerEdge M820 blade server is connected to the Dell Networking N3000 network switches.
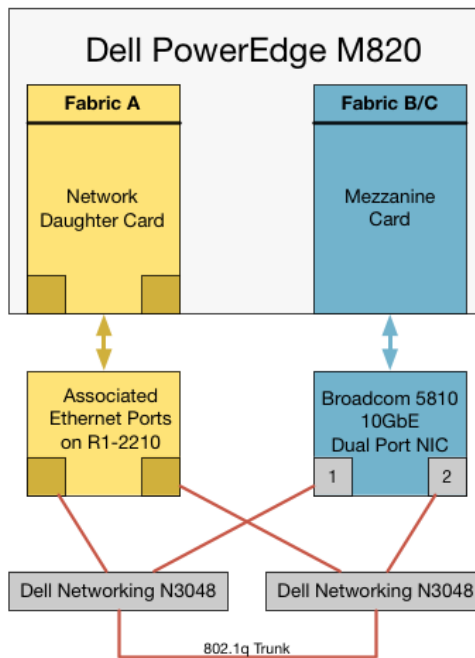
Figure 6   Network Edge Connectivity

Uplink: There are several options to uplink the Dell Networking N3000 switches to the customer LAN infrastructure. Selecting the uplink option depends on the customer network and customer requirements. One simple option is to create multiple uplinks on each switch and connect them to the customer LAN network switches. Uplink LAGs can then be created from the Dell Networking N3000 switches to the core network. If additional 10 GbE SFP+ ports are required, the N3000 switches have an expansion module on the rear of the switch allowing a SFP+ Dual-Port 10GbE Module to be installed. Figure 8 provides an overview of the inter-switch-link and network uplink.

Figure 7   Network Uplink Overview

Redundancy: The physical network is configured to have redundancy for all critical subsystems, and to avoid any single point-of-failure. For network connectivity, each PowerEdge M820 server has a Broadcom 5810 dual port NDC in Fabric A, and a Broadcom 5810 dual port NIC in Fabric B or C. As a result, the solution provides redundant network interface devices in each server. The design also includes two Dell Networking N3000 switches, connected to each other through an inter-switch-link. This design allows each PowerEdge M820 blade server to maintain network connectivity in the event of an individual network component failure.

5.2   Virtual Network Architecture

Using the VMware vSphere Client, one standard virtual switch is created on each ESXi host. Both network ports associated with the PowerEdge M820 blade servers are connected as uplinks to the virtual switch. This allows creation of a team of two network ports, enabling NIC failover and load balancing for the vSwitch. The virtual switch is configured for the LAN traffic types associated with the servers. The LAN traffic in this solution is categorized into four traffic types: vSphere management traffic, vMotion traffic, workload VM traffic, and management and Infrastructure services VM traffic. The solution also includes one more type of LAN traffic: Out-of-band management traffic. This traffic is associated with CMC, and is not handled by the vSwitch.

The virtual switch in each ESXi host has the following VMkernel connections and port groups: vSphere management, vMotion, workload VM port group, and management & infrastructure services VM port group. Figure 9 provides the view of this vSwitch from vCenter Server.
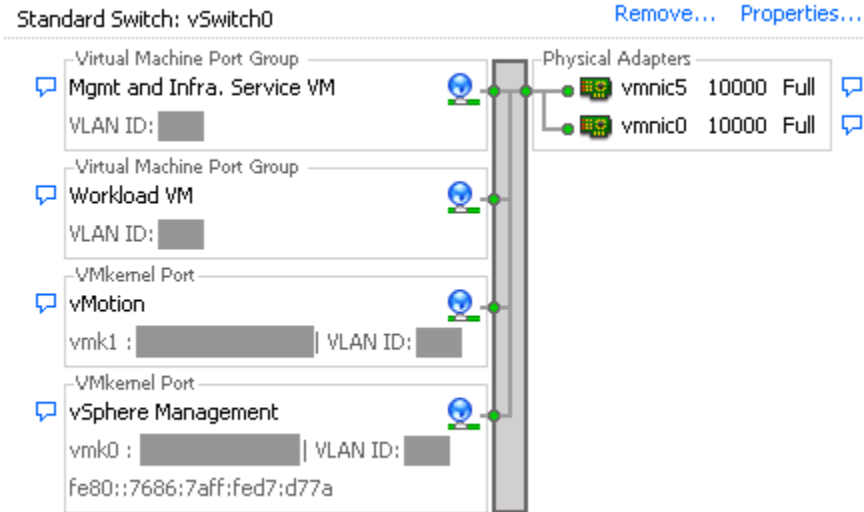
Figure 8   VMware vSphere vSwitch Overview

Congestion management and failover capability: Each VMkernel connection or VM port group uses one or more vmnics in active or standby fashion to provide congestion management and failover capabilities. vSphere management traffic has vmnic5 as active and all other vmnics as standby. vMotion traffic has vmnic0 as active and all other vmnic as standby. The virtual machine port groups have all vmnics as active. This design allows separation of vSphere management and vMotion traffic on separate physical ports during normal operating conditions. In addition, this configuration allows load balancing of all virtual machine traffic across all four vmnics. This overall design brings high availability by providing failover capability for each VMkernel connection and each VM port group during any failure on a physical port, NDC, PCIe mezzanine card, PCIe NIC, Ethernet pass-through or a Dell Networking N3000 switch. Table 6 summarizes the configuration of each connection to the vSwitch.

Table 6   Virtual Switch Configuration Details

vSwitch	Connection Type	Name	Active	Standby	Unused	Failback
vSwitch0	VMkernel	vSphere Management	Vmnic5	Vmnic0	None	Yes
	VMkernel	vMotion	Vmnic0	Vmnic5	None	Yes
	Virtual Machine Port Group	Workload VM	Vmnic5 Vmnic0	None	None	Yes
	Virtual Machine Port Group	Management and Infrastructure Services VM	Vmnic5 Vmnic0	None	None	Yes

Load Balancing and Failover: *Route based on the originating virtual switch port ID* configuration is used at the vSwitch for load balancing the LAN traffic. Any given virtual network adapter will use only one physical adapter port at any given time. In other words, if a VM has only one virtual NIC, it will use only one physical

adapter port at any given time. The reason for choosing this option is that it is easy to configure and provides load balancing across VMs, especially in the case of a large number of VMs.

Each host is configured with this vSwitch configuration. This configuration is described in detail in Figure 10.
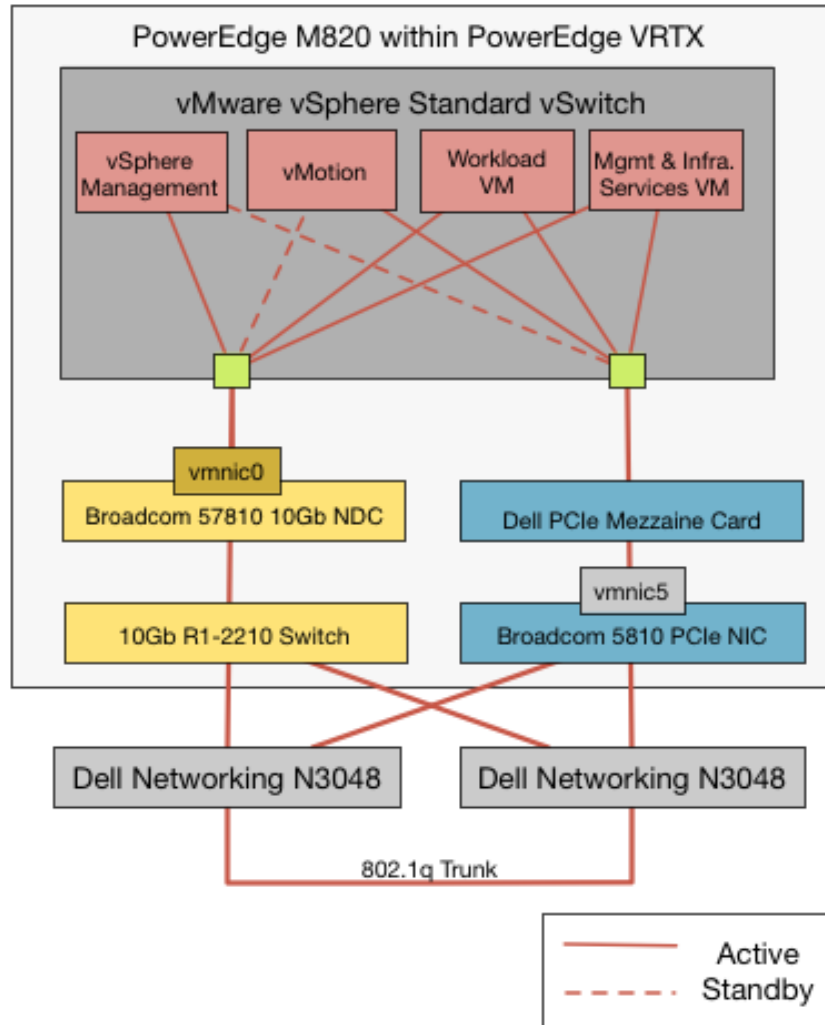
Figure 9   VMware vSphere vSwitch Overview

Traffic isolation using VLANs: This solution considers VLAN segregation of the LAN traffic types. The LAN traffic can be separated into three unique VLANs, as described in Table 7.

Table 7   VLAN Configuration:

VLAN	Associated Traffic Types
Management	vSphere Management Management & infrastructure Services VM traffic Out-of-band management traffic
vMotion	vMotion Traffic
Workload	Workload VM traffic

The network traffic should be tagged with the respective VLAN ID for each traffic type in the virtual switch. Additionally, the port channels configured on Dell Networking N3000 for inter-switch-links should be configured to trunk all VLANs. This VLAN segregation is not essential but it is recommended for traffic isolation. A different VLAN segregation strategy can be used based on the customer workload requirements. For example, more than one port group for workload VMs, each on a separate VLAN, can be created in the virtual switch.

6   Storage Architecture

This section provides details on the shared storage configuration required to support the reference architecture. This section also provides details on the configuration of a shared volume to host the management and infrastructure services within the PowerEdge VRTX chassis.

6.1   Shared Storage for Virtualization

This reference architecture solution has a single VMware vSphere cluster spanned across both PowerEdge M820 blade servers, as described in section 7. This design allows all virtual machines to benefit from VMware vSphere High Availability and DRS capabilities. To enable this virtualization infrastructure, this reference architecture solution uses the storage within the PowerEdge VRTX chassis shared across both servers.

One or more shared virtual disks can be created as per the customer workload requirements. As described in section 6.2 of this document, this solution also includes a shared Virtual Disk (VD) to host all management and infrastructure services. For shared storage configuration, it is recommended that one of the hard drives be configured to be the global hot spare to provide data protection in addition to the RAID configuration. In order to enable the shared storage, each virtual adapters of the shared PERC8 controller should be assigned to a corresponding server-slot, in order to provide the PowerEdge VRTX in that slot access to the shares storage. Table 8 provides the details of the virtual adapter assignment.

Table 8   PERC8 Virtual Adapter Mapping

Virtual Adapter	Server-Slot
VA1	Server Slot-1
VA2	Server Slot-2
VA3	Server Slot-3
VA4	Server Slot-4

In order to create VDs that all servers can access, enable Multiple Assignment mode for the shared storage from the CMC user interface.

6.2   Storage for Management and Infrastructure Services

This reference architecture solution includes a shared VD configuration for management and infrastructure services, and additional storage that can be customized to meet the customer workload requirements. The management and infrastructure services components can be installed in virtual machines, and can be hosted on the PowerEdge VRTX system, along with the customer workloads. Section 8 of this document provides further details on the management and infrastructure services components included in this solution, their sizing, and how they can be deployed. From the storage point-of-view, a volume that is shared across all servers should be created to host all management and infrastructure services virtual machines. Allowing access to all servers will enable the management and infrastructure VMs to be run on the same VMware vSphere cluster as the workload VMs; and will enable

these VMs to benefit from HA and DRS functionality. In the rest of this document, this shared VD for management and infrastructure services VMs will be referred to as the "management VD".

This reference architecture solution provides the design and configuration details for the management VD for a PowerEdge VRTX system with 2.5" HDDs. The shared storage configuration with 25x 2.5" HDDs provide the dense storage with capacity and IOPS suitable for a virtualization cluster. The management VD should be created on four 900 GB 10K RPM SAS HDDs. Considering the criticality of the management and infrastructure services, it is recommended that these disks are configured with RAID 10.

Full read/write access for the management VD should be provided to both PowerEdge M820 blade servers through the CMC user-interface. Figure 11 shows the details of the shared storage configuration and the management VD.
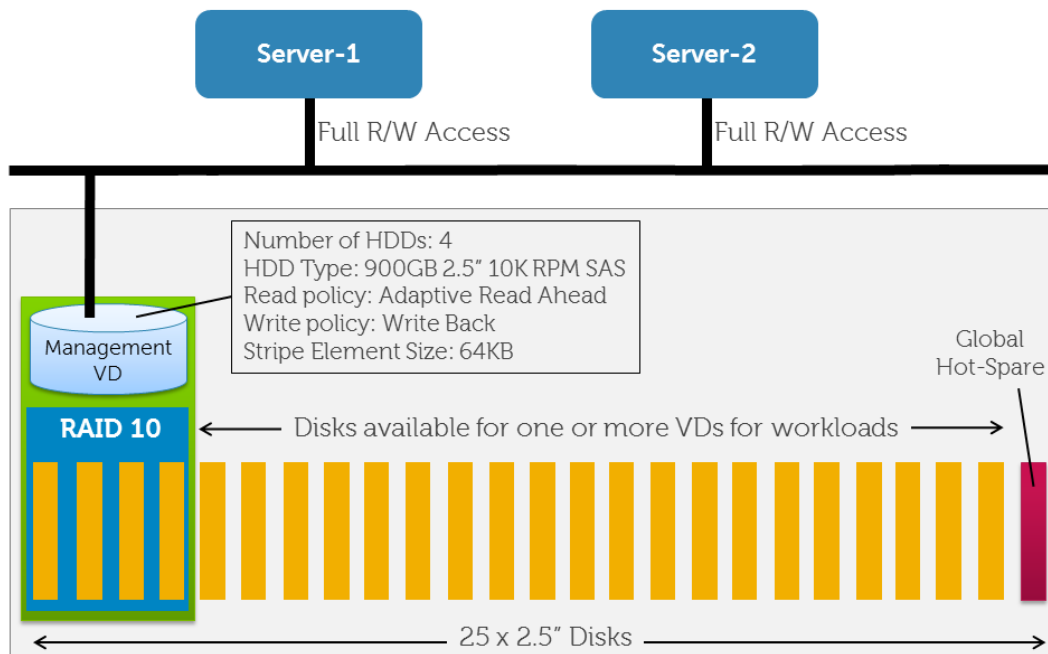
Figure 10   Management VD Configuration

6.2.1   Virtual Disks for Workloads

After creating the management VD, the remaining disk slots can be populated with supported HDDs or SSDs to be used for different workloads. One or more VDs can be created on these HDDs or SSDs to support different the customer workloads. Different disk size and speed, different RAID levels, and different caching policies can be used for different VDs to support different workloads with different capacity, performance and data protection requirements.

7 Virtualization Cluster

VMware ESXi is installed on a pair of redundant SD cards in each M820 server. The configured ESXi hosts are added to a vCenter Server for management. This solution considers a single VMware vSphere cluster spanned across both ESXi host servers. VMware features, including vMotion, VMware High Availability (HA), and Distributed Resource Scheduling (DRS) should be enabled if allowed by licensing.

Creating a single VMware vSphere cluster provides the following advantages:

- **High availability:** VMware vSphere High Availability (HA) provides high availability for applications running in virtual machines. In the event of physical server failure, affected virtual machines are automatically restarted. Creating a single HA cluster across more than one server allows the cluster to sustain failure on one or more servers.

- **Simplified management:** A single cluster allows convenient configuration and management of all hosts from vCenter server.

8 Converged Management

This reference architecture includes Dell Chassis Management Controller (CMC), which provides the capability to manage the complete infrastructure within the PowerEdge VRTX chassis through a single interface. Additionally, the solution is designed to integrate certain recommended management services, hosted as virtual machines, to configure virtualization infrastructure, cloud connectivity, and data protection services. The necessary infrastructure services, like AD, DNS, DHCP server, etc., can also be hosted within the PowerEdge VRTX chassis, if desired.

8.1 Single Chassis Infrastructure Management Interface

The CMC provides a single interface to configure, manage and monitor the entire PowerEdge VRTX system. It can be used to control the power-up and power-down operations for the chassis, servers and the I/O module. It provides an interface to configure and manage the servers, PCIe resources, the I/O module, mapping of the PCIe slots to servers, physical disks, RAID controller, virtual disks, virtual disks mappings, chassis power budget, etc. The CMC also provides an interface for logging faults, sending alerts, controlling the front panel, and for many other general functions related to management, configuration.

PCIe slot assignment: CMC provides the capability to assign the PCIe slots to server slots. When a PCIe slot is assigned or unassigned, the server in the corresponding server slot must be turned off. To assign a PCIe slot, open the **CMC** select **Chassis Overview** > **PCIe Overview** > **Setup Tab** > **Mapping**. As shown in Figure 13, the PCIe slots 1 and 2 have Broadcom 5810 dual port NICs. As described in the reference architecture, the PCIe slots 1 and 2 are mapped to server slots 1 and 2 respectively.

PCIe Slot				Server Slot		
Slot	Name	Fabric	Power Status	Name	Slot	Action
1	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet [14e4:1006]	C	On		1	No Action
2	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet [14e4:1006]	C	On		2	No Action
3	Empty	B	N/A		N/A	No Action
4	Empty	C	N/A		N/A	No Action
5	Empty	C	N/A		N/A	No Action
6	Empty	B	N/A		N/A	No Action
7	Empty	B	N/A		N/A	No Action
8	Empty	B	N/A		N/A	No Action
						Cancel Apply

Figure 11   PCIe Slot Assignment from CMC

PERC8 Virtual Adapter assignment: CMC provides the capability to assign the four virtual adapters of PERC8 to server slots. To assign the virtual adapters, open the **CMC** select **Chassis Overview** > **Storage** >

Setup Tab. As shown in Figure 14, the VAs 1 and 2 are mapped to the servers 1 and 2 respectively – as per this reference architecture.

Virtual Adapter	Server Slot Mapping	Action
Virtual Adapter 1	SLOT-01	No Action
Virtual Adapter 2	SLOT-02	No Action
Virtual Adapter 3	Unmapped	No Action
Virtual Adapter 4	Unmapped	No Action

Restore Default   Cancel   Apply

Figure 12   PERC8 Virtual Adapter Assignment from CMC

Virtual disk assignment Mode: CMC provides the capability to configure the virtual disk assignment mode. One virtual disk can be accessed by multiple servers only when the "Multiple Assignment" mode is selected. To configure the virtual disk assignment mode, open the **CMC** select **Chassis Overview** > **Storage** > **Setup Tab** (Figure 15).

Assignment Mode	Description
◯ Single Assignment	This mode allows a virtual disk to be assigned to a single virtual adapter at a time
◉ Multiple Assignment	This mode allows a Virtual Disk to be assigned to multiple Virtual Adapters at a time ⚠ Do not use this mode unless the servers have Cluster Services installed on them. Use of this mode without Cluster Services may lead to corrupted or lost data.

Cancel   Apply

Figure 13   Virtual Disk Assignment Mode Selection from CMC

Create virtual disk: CMC provides the capability to create a new virtual disk with the desired properties and characteristics. To create a new virtual disk, open the **CMC** select **Chassis Overview** > **Storage** > **Virtual Disk** > **Create** (Figure 16). The figure shows the properties selected to create the management VD.

Name	Management VD	Read Policy	Adaptive Read Ahead
Controller	Shared PERC8	Write Policy	Write Back
RAID Level	RAID 1	Disk Cache Policy	Default
Media Type	HDD		
Stripe Element Size	64KB		
Capacity	GB *See below for capacity limits*	Number of Spans	1

Figure 14   Creating Virtual Disk from CMC

Assign virtual disk access: CMC provides the capability to manage and configure the access to each virtual disk for different servers. To assign virtual disk access, open the **CMC** select **Chassis Overview** > **Storage** > **Virtual Disk** > **Assign**. Figure 17 shows the access to the management VD for each server. Full read/write access to the management VD is granted for each server. It is recommended that full read/write access be granted for each server to any new VD that will host VMs.

Virtual Disk Name	[Virtual Adapter 1] Server Slot SLOT-01	[Virtual Adapter 2] Server Slot SLOT-02	[Virtual Adapter 3] Server Slot Unmapped	[Virtual Adapter 4] Server Slot Unmapped
	To setup virtual adapter mapping or a virtual disk assignment mode, go to page: Setup Storage *Current Assignment Mode:* Multiple Assignment			
VMDK5-R10-6H2S	Full Access ▾	Full Access ▾	No Access ▾	Full Access ▾
			Cancel	Apply

Figure 15   Assigning Virtual Disk Access to Servers from CMC

The details above will help configure a PowerEdge VRTX system as per this reference architecture. The Dell CMC also enables many other configuration, management and deployment capabilities for all the components within the PowerEdge VRTX system. For comprehensive information on the Chassis Management Controller for Dell PowerEdge VRTX, see the PowerEdge VRTX Manuals.

8.2   Management and Infrastructure Services Integration

The reference architecture is designed to include management and infrastructure services within the solution infrastructure. It includes VMware vCenter server, Dell management plug-in for vCenter, Dell OpenManage Essentials, Quest vRanger, VMware vCloud Connector, etc. The infrastructure services may include Active Directory (AD), domain name service (DNS), network time protocol (NTP) server, etc. This section provides recommendations for the specific management components and the deployment strategy for them as well as the necessary infrastructure services components.

The necessary management and infrastructure services will be run as virtual machines on the same vSphere HA cluster as the workload virtual machines. As described in the section 5.2, the virtual switch on each ESXi host is configured to handle the management and infrastructure services VM traffic on a dedicated port group (and a dedicated VLAN, if desired). The PowerEdge VRTX shared storage provides the capability to create a separate VD for this management and infrastructure services components, as described in the *Storage Architecture* section. This design allows the customer to ensure that certain storage capacity and IOPS, and a specific RAID level are available for the management and infrastructure services.

The following management components are included in the reference architecture solution.

- Dell OpenManage Plugin for vCenter
- Dell OpenManage Essentials
- Quest vRanger Server
- Quest vRanger Virtual Appliance
- VMware vCenter Server
- VMware vCloud Connector Server
- VMware vCloud Connector Node

Additionally, the infrastructure services components, like Active Directory, Domain Controller, DHCP server, NTP server, etc., can be run on the same infrastructure along with the management components.

As illustrated in Figure 18, these components are installed as virtual machines in the management infrastructure:
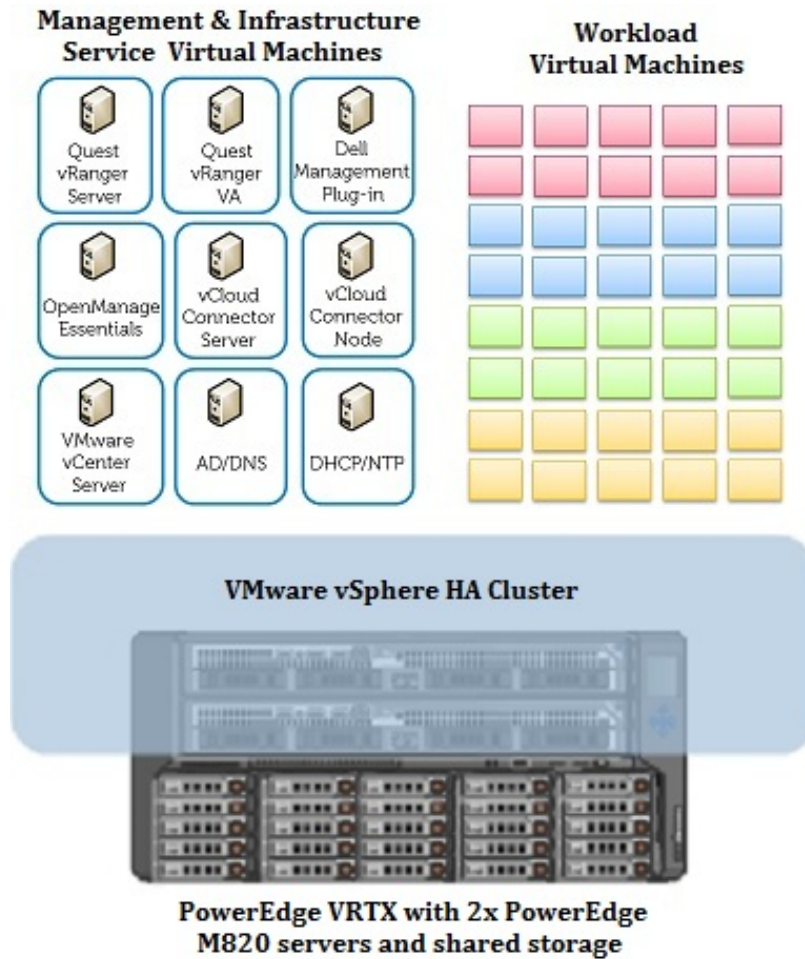
Figure 16   Management and Infrastructure Services Deployment

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 in Table 9.

Most of these management components support a variety of operating systems. The resource calculations for the management components, as shown below, consider Windows Server 2012 as the guest operating system, where it is supported by the particular management component. If a management component does not support Windows Server 2012, the highest version of the operating system supported by that particular component has been considered.

Table 9   Management Component Sizing

Management Components as Virtual Machines	vCPU	vRAM	Disk
VMware vCenter Server Appliance (Guest OS: Suse Enterprise Linux)	2	12 GB	120 GB
Dell OpenManage Essentials (Guest OS: Windows Server 2012)	2	4 GB	50 GB
Dell Management Plug-in for VMware vCenter (Virtual appliance)	2	3.5 GB	40 GB
VMware vCloud Connector Server (Virtual appliance)	2	3 GB	13 GB
VMware vCloud Connector Node (Virtual appliance)	2	2 GB	53 GB
Quest vRanger Server (Guest OS: Windows Server 2008 R2)	4	4 GB	50 GB
Quest vRanger Virtual Appliance (Virtual appliance)	2	1 GB	14 GB

In additional to these management components, additional resources are needed for the following:

- **Infrastructure services components:** The components will include some combination of Active Directory, Domain Controller or Read-Only Domain Controller, DNS, DHCP server, NTP server, etc. If the workload requires a SQL server, a separate installation of SQL server can also be included. The number of virtual machines and their sizing to support these infrastructure services will depend upon the specific customer requirements and usage model.

- **VMware vCloud Connector sizing optimization:** vCloud Connector Node storage should be appropriately sized to hold the largest VM that customers will migrate multiplied by the expected number of concurrent migrations. Consult vCloud Connector documentation on vmware.com for further details.

- **Quest vRanger deployment optimization:** Quest vRanger architecture allows using virtual appliances in addition to a single vRanger server. The vRanger virtual appliance can process backup and restore tasks in addition to replication tasks. Deploying a virtual appliance allows scaling back up, restore and replication activities across multiple hosts or clusters, while maintaining central scheduling and reporting control from a single vRanger server. While one vRanger virtual appliance can be deployed per vSphere cluster, a VM dense environment may benefit from deploying one vRanger virtual appliance per ESXi host. In addition, scaling up the size of the virtual appliance may also be helpful in some cases.

The recommended sizing for the management VD, as described in section 6.2, accommodates the management components (Table 9), and the additional variable requirements of the infrastructure services components, vCloud Connector, and vRanger.

This model of running management and infrastructure services in the same vSphere cluster as the workload VMs provides several benefits. First, this architecture brings the benefit of vSphere high availability cluster to the management and infrastructure services components, along with the workload VMs.

While deploying the workload virtual machines, the resource requirements for the management and infrastructure services components should be taken into consideration to ensure these components always get the needed compute and storage resources. The customer can use vSphere Resource Pool capability to create a resource pool for the management and infrastructure services VMs, if supported by the chosen vSphere license. The information in Table 9 can be used to determine the resource reservations for the resource pool. Using a resource pool will help ensure the availability of required resources for the management and infrastructure services components.

8.2.1 Dell OpenManage Essentials (OME)

Dell OpenManage Essentials (OME) is installed on an individual virtual machine. The VM is sized and configured to enable an installation of OME to monitor the components of this reference architecture solution. OME is configured to utilize a local SQL Express database. For fullest functionality, direct internet access, or through a proxy, is recommended.

OME is utilized for discovery, inventory and hardware level monitoring of PowerEdge VRTX chassis, PowerEdge M820 blade servers and Dell Networking N3000 network switches. Each of these components are configured to send SNMP traps to the centralized OME console to provide a single monitoring interface for major hardware components. OME provides a comprehensive inventory of solution components through WS-MAN and SNMP inventory calls. As a part of hardware inventory, OME displays information about PCIe slot assignments, mapping for virtual adapters for PERC, and other components. Firmware versions and solution warranty status is reported. OME can be used as the single point of monitoring for all hardware components within the solution. OME can also be used to update the firmware for CMC and the other components within the PowerEdge VRTX chassis.

OME Map View: The Map 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. Figure 18 below shows the Map View feature. The figure shows multiple PowerEdge VRTX systems that are dispersed around the world at the company's offices.

Figure 17   OME Map View Feature

For more information on OpenManage Essentials, see the [Data Center Systems Management](#) page and the [OME page on Dell TechCenter](#).

8.2.2   Dell Management Plug-in for VMware vCenter (DMPVV)

Dell Management Plug-in for VMware vCenter is deployed as a virtual appliance, and is attached to the VMware vCenter Server. DMPVV communicates with the VMware vCenter Server, the hypervisor management interfaces and server out-of-band management interfaces (iDRAC). For ease of using the appliance, firmware updates and warranty information, it is recommend that the DMPVV appliance has access to an internet connect either directly, or through a proxy. Dell Management Plug-in for VMware vCenter enables customers to:

- Get deep-level detail from Dell servers for inventory, monitoring, and alerting — all from within vCenter

- Apply BIOS and Firmware updates to Dell servers from within vCenter

- Automatically perform Dell-recommended vCenter actions based on Dell hardware alerts

- Access Dell hardware warranty information online

- Rapidly deploy new bare metal hosts using Profile features

For more information, see the web page for [Dell Management Plug-in for VMware vCenter](#).

8.2.3   VMware vCloud Connector

VMware vCloud Connector lets you view, operate on and transfer your computing resources across vSphere and vCloud Director in your private cloud environment, as well as a public vCloud. VMware vCloud Connector is deployed with three VMs for the base functionality: A single 'server' virtual machine

and two 'node' virtual machines. The node VMs are responsible for the physical transfer of VM workloads. Within the two-node vSphere HA cluster in PowerEdge VRTX, two of these components, the server and the local node, are installed. The third component, 'remote' node VM, should be installed outside of this reference architecture solution, near the infrastructure to which it provides connectivity.

After deploying the VMware vCloud Connector 'node' VMs, the size of the virtual disk may have to be increased based on the size of the expected VMs to be transferred and the number of concurrent transfers anticipated.

The key capabilities provided by VMware vCloud Connector are:

- Expand your view across hybrid clouds. Use a single management interface that seamlessly spans your private vSphere and public vCloud environment.

- Extend your data center. Move VMs, vApps and templates from private vSphere to a public vCloud to free up your on premise data center resources as needed.

- Consume cloud resources with confidence. Run Development, QA and production workloads using a public vCloud.

For more information, see [Dell vCloud website](#).

8.2.4   Quest vRanger from Dell

Quest vRanger from Dell provides a simple, fast, and scalable data protection solution for VMware vSphere based virtual environments and Microsoft Windows-based physical environments. In the case of this reference architecture solution, it is recommended that Quest vRanger be deployed in a virtual machine. The virtual machine is sized for the reference architecture infrastructure. Additionally, a vRanger virtual appliance (VA) is also deployed to protect the VMware vSphere cluster. The management VD is sized to accommodate four such VAs, one per ESXi host.

Quest vRanger provides a single console for managing backup, replication and recovery. A Quest vRanger-based data protection solution scales with the virtual environment by maximizing resources through distributed processing, while simplifying management with central command and control. Some of the major characteristics and capabilities of Quest vRanger are:

- Performs incremental, differential, or full image backups of virtual machines
- Quickly restores the entire virtual machine or just specific files
- Manages disaster recovery strategies and protects critical data in virtual environments
- Offers the ability to replicate virtual machines to more than one destination. This allows users to accomplish multiple High Availability and Disaster Recovery objectives based on their specific needs.
- Uses an agentless architecture
- Supports backup and recovery of Windows physical servers, files, and folders
- Operates as a low-resource consumption Virtual Appliance (VA) for low-impact scalability
- Performs LAN-free backups using VMware SCSI Hot Add with vRanger installed inside a VM and from the vRanger VA

- Supports VMware vMotion to ensure that the VMs are protected as they move from one host to another, even when backup jobs are running

For more information, see [Quest vRanger website](#).

For backup operations, vRanger captures the complete VM image (including OS, patches and applications) and transfers it to a pre-configured data repository. These image level backups can be restored in a matter of minutes. In essence, replicating a VM is replicating the complete VM image, including the specific changes to these images that reflect user-specified settings for the source VM. vRanger is designed to integrate with VMware vSphere at the API level.

A complete vRanger installation includes the following four components in order to enable backup, replication and disaster recovery capabilities for this reference architecture:

- **vRanger server:** vRanger server can be a physical server or a virtual machine. The benefit of installing vRanger on a physical server is that the resource consumption of backup activity is off-loaded from the virtual environment to the physical server. On the other hand, installing vRanger in a virtual machine eliminates the need for dedicated hardware while maintaining high performance. In addition, vRanger has to be installed in a virtual machine in order to use the VMware SCSI Hot Add functionality on VMware ESXi.

  As discussed in section 7, it is advantageous to use both blade servers in PowerEdge VRTX as hypervisor hosts. In addition, using SCSI Hot Add capability allows LAN-free operation, which significantly reduces the network traffic. Due to these reasons, this reference architecture solution recommends installing vRanger in a virtual machine.

- **vRanger database:** vRanger utilizes a Microsoft SQL database to store application and task configuration data. The database can be either the embedded SQL Express instance (the default option) or a SQL database running on a separate SQL Server or SQL Express instance. If the vRanger cataloging feature is to be used, the SQL instance must be installed on the vRanger server. Due to this fact, and considering the size of this PowerEdge VRTX based infrastructure, this reference architecture solution recommends using the Microsoft SQL Express instance embedded with the vRanger installation, although the reference architecture infrastructure is sufficiently sized to include a separate database server.

- **vRanger virtual appliance(s):** The vRanger virtual appliance can process backup and restore tasks in addition to replication tasks. This allows scaling back up, restore and replication activities across multiple hosts or clusters, while maintaining central scheduling and reporting control from a single vRanger server. One virtual appliance can be deployed on each host, or a single virtual appliance can be shared among the hosts within a cluster. If a virtual appliance is not detected on the host, vRanger will check if the host is part of a cluster, and then if that cluster has a virtual appliance available. This reference architecture considers one virtual appliance for the entire vSphere cluster comprised of all four hosts.

The recommended virtual appliance size is 2x vCPU and 2GB RAM. If there are a large number of VMs on the vSphere cluster, this virtual appliance can be scaled up. Additionally, instead of using one virtual appliance for the entire vSphere cluster, one virtual appliance can be deployed on each ESXi host. Considering that this reference architecture solution only has a four-node cluster, only one virtual appliance is considered, although the management VD and the overall infrastructure is sufficiently sized to deploy one scaled-up VA on each ESXi host. A different clustering strategy or vRanger virtual appliance deployment strategy can also be used based on the workload requirements.

For more information on virtual appliance sizing and deployment strategy, see [vRanger User Guide](#).

- **At least one repository:** A repository is essentially a directory on a supported file system that vRanger uses to store save-points (backup archives). Repositories can be one of these formats: CIFS, NFS, FTP, SFTP and NetVault SmartDisk.

  For more information on NetVault SmartDisk, visit http://www.quest.com/netvault-smartdisk.

Quest vRanger server provides a single interface to provide a user-friendly inventory of all vCenter servers, ESXi hosts and virtual machines; and also to manage all backup, recovery, and replication jobs.

8.3   Backup and Recovery

Quest vRanger allows backup operation for a VM, ESXi host, folder, resource pool, data center, or vCenter and backup all of the VMs located under the nodes in the tree. Incremental, differential or full image backup can be performed based on the requirements of the environment.

The vRanger server is installed in a virtual machine. This virtual machine can be configured on the management VD and to use the Management and Infrastructure Services VM port group for network traffic. Similarly, one vRanger virtual appliance can be installed for the entire vSphere cluster. This virtual appliance can also be configured on the management VD and to use the Management and Infrastructure Services VM port group for network connectivity. It is recommended to use the SAN HotAdd transport option to minimize the network traffic.

The backup repository can be on PowerEdge VRTX shared storage or on another system on the LAN infrastructure where the vRanger server can access it. Although the repository can be configured on the PowerEdge VRTX shared storage, it is recommended to configure it on a system on a separate network and power fault domain outside the PowerEdge VRTX chassis. Figure 19 below shows both of the options for creating a backup repository.
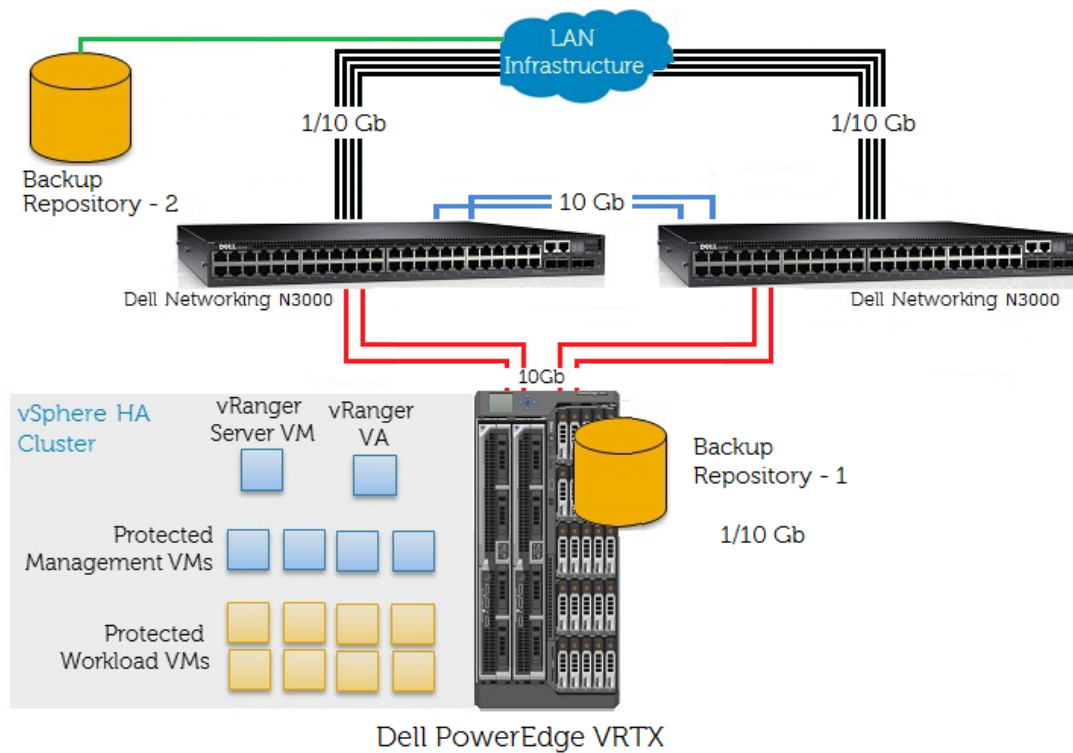
Figure 18   Backup and Recovery with Quest vRanger

8.4   Replication and Disaster Recovery

Quest vRanger supports replication of VMs via the vRanger virtual appliance (VA). A VM on PowerEdge VRTX can be replicated on a compatible ESXi host or a VMware vSphere cluster.

The vRanger server is installed in a virtual machine. This virtual machine can be configured on the management VD and use the Management and Infrastructure Services VM port group for network traffic. Similarly, one vRanger virtual appliance can be installed for the entire vSphere cluster. This virtual appliance can also be configured on the management VD and use the Management and Infrastructure Services VM port group for network connectivity. It is recommended to use the SAN HotAdd transport option to minimize the network traffic. In addition, at least one vRanger VA needs to be deployed on every ESXi server or vSphere cluster that will be used as a replication target.

This reference architecture solution considers two scenarios: Replication at the local site, and replication at a remote site for disaster recovery. Quest vRanger allows replication to multiple destinations, and both of these options can be implemented.

Replication at the local site: The VMs running on a PowerEdge VRTX can be replicated on a destination infrastructure at the local site. This destination infrastructure can be another PowerEdge VRTX system or any other infrastructure compatible with vRanger requirements. At least, one vRanger VA needs to be

deployed on this destination infrastructure. Figure 20 below provides a representation of this replication scenario.
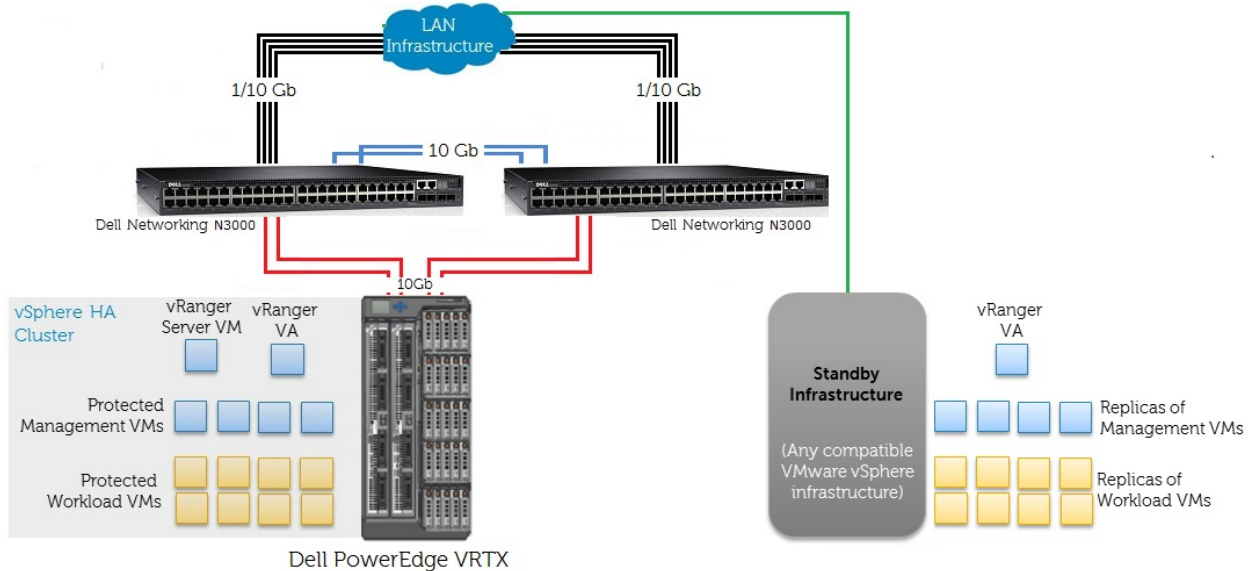

Figure 19   Replication with Quest vRanger

Replication at a remote site for disaster recovery (DR): The VMs running on a PowerEdge VRTX can be replicated at a destination infrastructure at a remote site. This destination infrastructure can be another PowerEdge VRTX system or any other infrastructure compatible with vRanger requirements. As an example, this solution considers Dell Active System 800 as the destination system. In the case of multiple remote and branch offices, each remote or branch office can use this PowerEdge VRTX based reference architecture solution; and a Dell Active System solution can be used at a head-office or at a central data center as the replication target for all remote and branch offices. Figure 21 below provides a representation of this replication scenario.
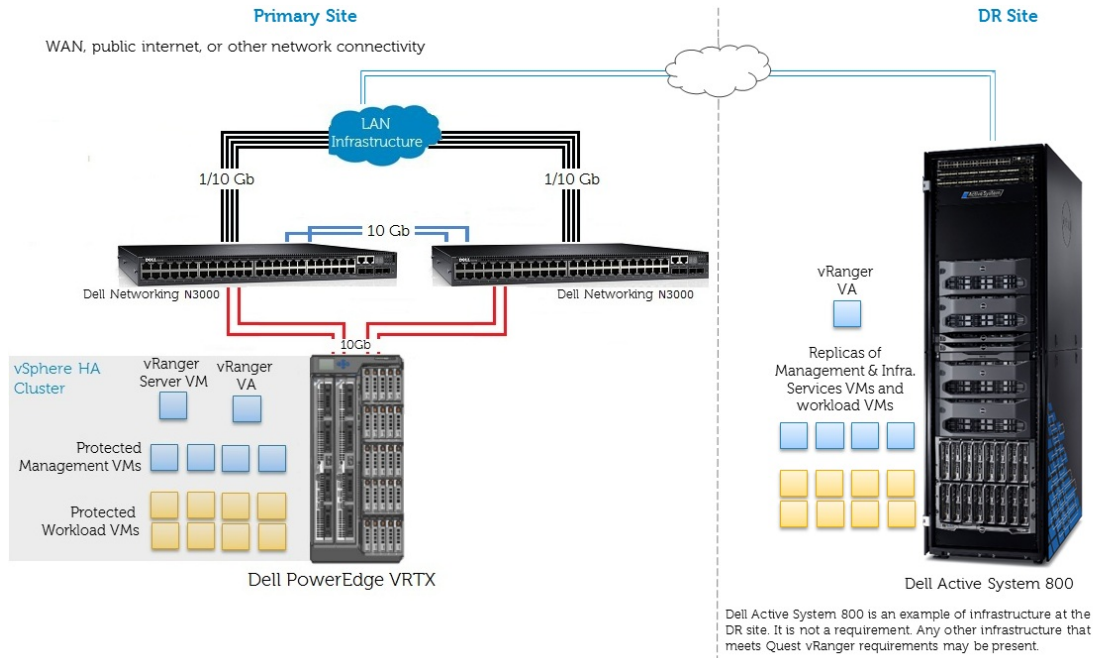
Figure 20   Disaster Recovery with Quest vRanger

9 Summary

This white paper provides a reference architecture for VMware vSphere based virtualization infrastructure for remote offices, branch offices and small businesses. The white paper also provides recommended settings and configuration details to efficiently integrate management, data protection and cloud enablement services within this infrastructure.

The reference architecture enables an enterprise class two-node hypervisor cluster with integrated shared storage and Gigabit Ethernet based LAN infrastructure, all in the small form-factor of a Dell PowerEdge VRTX system for efficient use of real estate at a remote or branch office. The reference architecture uses a redundant pair of Dell Networking N3000 switches for connectivity to the customer LAN infrastructure. The white paper provides configuration details for the PowerEdge M820 blade server network interfaces, and for the shared PCIe slots, the I/O module slot and the shared storage in PowerEdge VRTX. The white paper also provides recommendations for virtual and physical network configuration. In addition, the white paper includes the recommended settings for configuring a virtual disk (VD) on the shared storage in PowerEdge VRTX to host certain management and infrastructure services components.

The white paper describes the recommended settings to integrate within the infrastructure VMware vCenter Server Appliance, Dell Management Plug-in for VMware vCenter, Dell OpenManage Essentials, VMware vCloud Connector and Quest vRanger from Dell.

A   Terminology

API: Application programming Interface

CIFS: Common Internet File System

CLI: Command Line Interface

CMC: Dell Chassis Management Controller for Dell PowerEdge VRTX

DHCP: Domain Host Configuration Protocol

DRS: VMware Distributed Resource Scheduler

FTP: File Transfer Protocol

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

NFS: Network File System

NTP: Network Time Protocol

PERC: Dell PowerEdge RAID Controller

RAID: Redundant Array of Independent Disks

SAS: Serial Attached SCSI

SFTP: SSH File Transfer Protocol

SSD: Solid State Drive

TFTP: Trivial File Transfer Protocol

VA: Virtual Adapter

VD: Virtual Disk

VM: Virtual Machine

B Additional Resources

Support.Dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for sharing knowledge, best practices and information about Dell products and installations.

Below are some helpful resources.

- Dell PowerEdge VRTX

- Dell PowerEdge VRTX manuals

- Dell PowerEdge family manuals

- Manuals for Dell Chassis Management Controller for PowerEdge VRTX, Dell iDRAC, and Dell Lifecycle Controller

- Dell Networking N3048 Systems User's Guide

- Dell Management Plug-In for VMware vCenter references – Solution Brief

- VMware vSphere Documentation

- VMware vSphere License Versions and Kits Comparison

- VMware Virtual Networking Concepts

- Quest vRanger from Dell

- vRanger Installation/Upgrade Guide