

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

Dell Engineering
July 2016

Revisions

Date	Description	Authors
July 2016	Version 4 Update wording, section 8	DNSE: EAB,
June 2016	Version 3 Updated to include new VMware software and M820 to M830 server hardware	DNSE: EAB, DS, EC, SC
April 2014	Version 2 Updated to include new components and revisions.	DNSE: EAB, CB, MM
June 2013	Initial release	DNSE: PP

Copyright © 2013- 2016 Dell Inc. or its subsidiaries. All Rights Reserved.

Except as stated below, no part of this document may be reproduced, distributed or transmitted in any form or by any means, without express permission of Dell.

You may distribute this document within your company or organization only, without alteration of its contents.

THIS DOCUMENT IS PROVIDED “AS-IS”, AND WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED. IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE SPECIFICALLY DISCLAIMED. PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/vn/terms-of-sale-commercial-and-public-sector-warranties>

Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell’s recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text: Dell™, the Dell logo, Dell Boomi™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of Dell. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of QLogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.

Table of contents

Revisions.....	2
1 Introduction.....	4
2 Overview.....	6
2.1 Solution Capabilities and Use Cases	6
2.2 Solution Components	8
2.2.1 VMware vSphere	10
2.2.2 Dell PowerEdge VRTX	11
2.2.3 Dell Networking N3048 Switch	15
3 Design Principles.....	17
4 Reference Architecture.....	18
5 Network Architecture	19
5.1 Physical Network Architecture	19
5.2 Virtual Network Architecture	21
6 Storage Architecture.....	25
6.1 Shared Storage for Virtualization.....	25
6.2 Storage for Management and Infrastructure Services.....	25
6.2.1 Virtual Disks for Workloads	26
7 Virtualization Cluster	27
8 Chassis Management.....	28
8.1 Single Chassis Infrastructure Management Interface	28
8.2 Management and Infrastructure Services Integration	30
8.2.1 Dell OpenManage Essentials (OME)	33
8.2.2 Dell Management Plug-in for VMware vCenter (DMPVV).....	34
8.2.3 VMware vCloud Connector.....	34
8.2.4 Quest vRanger from Dell	35
8.3 Backup and Recovery	37
8.4 Replication and Disaster Recovery	38
9 Summary	41
A Terminology.....	42
B Additional Resources	43
About Dell EMC	43

1 Introduction

This document describes a Dell EMC-designed and validated reference architecture for virtualization. Dell EMC based the architecture on the state of the art Dell PowerEdge VRTX system, Dell PowerEdge M830 blade servers, Dell Networking N3000 switches and VMware vSphere Hypervisor. This architecture targets remote/branch offices and small businesses.

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 document provides reference architecture for virtualization that meets 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 Gb switches with two 10Gb uplinks. In this paper, the 10Gb uplinks connect the Dell PowerEdge VRTX to the LAN and the Gb ports are reserved for access ports for branch office workstations, printers, and so on.

This document also provides recommended settings for integrating infrastructure lifecycle-management components, cloud connectivity and data protection. These components include the following:

- VMware vCenter Server
- Dell Management Plug-in for VMware vCenter
- Dell OpenManage Essentials
- VMware vCloud Connector
- Quest vRanger from Dell

The document 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 deploy this architecture into production environments quickly and confidently. This helps eliminate costly and time consuming trial-and-error work often encountered during complex deployments. Before purchasing components, this document aids customers in preparation tasks, such as the following:

- Sizing
- Selecting appropriate license levels
- Use-case planning
- Deployment preparation

After purchasing components, this document aids with setup, deployment and infrastructure configuration.

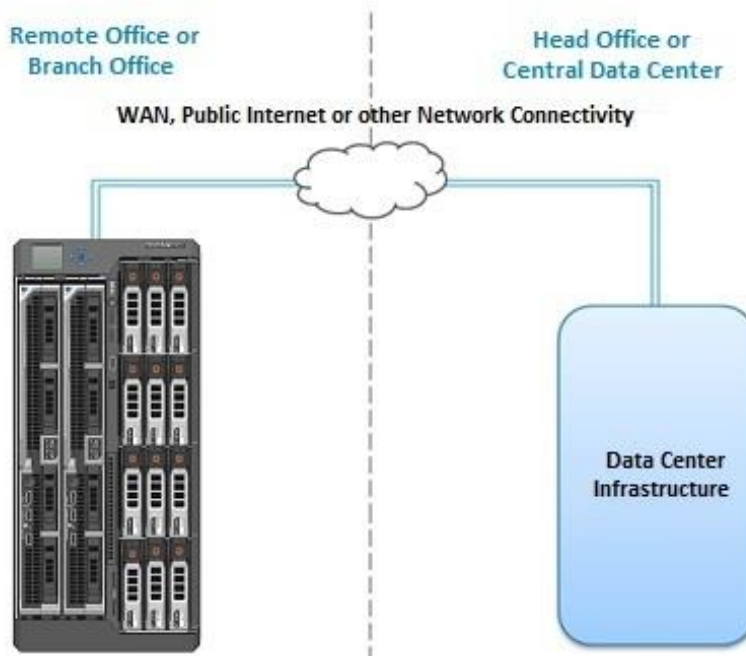
Audience

Dell EMC intends information in this document for IT administrators and managers who purchase, or plan to purchase virtualization infrastructure. The document specifically describes systems for remote offices, branch offices, small offices or small and medium businesses. Concerned parties in these locations 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 provides an enterprise-class, cost-effective, virtualization infrastructure for remote offices, branch offices and small businesses. The solution integrates simple, highly efficient lifecycle management, cloud connectivity, data protection services and other necessary infrastructure services like domain services, databases, and so on. Dell EMC envisions the solution 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 starts with a Dell PowerEdge VRTX platform using Dell PowerEdge blade servers. This platform includes best-in-class embedded server-management capabilities, shared storage and network I/O ports.
 - **Dense compute:** The solution uses two 13th generation Dell PowerEdge M830 blade servers. Each blade server supports up to 3TB of RAM (48 x 64 GB DIMMs).
 - **10GbE networking:** The solution uses QLogic Gigabit network interface cards in the M830 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 enables cost savings by consolidating power and cooling. It 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 enables a VMware vSphere cluster comprising of two server nodes and allows 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 feature for remote or branch offices, which often must move workloads between their locations and a head office or a central data center. The following is a partial list of the advantages of virtualization:

- Resource consolidation
- Increased energy efficiency
- Better business continuity
- HA
- Cloud enablement

These advantages can be critical for remote and branch offices, which often rely on limited physical compute capacity and limited real estate for infrastructure.

- **Integrated management:** The solution leverages the Dell Chassis Management Controller (CMC) 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.

For remote or branch offices, some of these components may also be installed at the head office or central data center. Having the components available locally allows the remote/branch office to continue operation after a loss of network connectivity to the head office or data center. The single interface provided by the Power Edge VRTX's Chassis Management Controller enables local, remote or branch office administrators to manage the complete solution. This can include servers, shared storage, network modules, PCIe devices and so on.

- **Data protection:** This reference architecture supports and integrates Quest vRanger within the PowerEdge VRTX-based virtualization infrastructure. It enables backup of workload VMs, replication of VMs, failover capabilities and data recovery from the head office. It also includes recommendations for configuring these components at the remote/branch office, and at a head office or central data center.

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

- **Cloud connectivity:** This 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. Workload migration enables a remote or branch office to optimize the geographical locality of different workloads for network bandwidth, hardware availability, performance and so on. This reference architecture supports this critical capability.

2.2 Solution Components

This section provides a high-level overview of the major building blocks, that is, hardware components, software components, data protection and cloud connectivity components. As described previously, this solution is based on a Dell PowerEdge VRTX system with Dell PowerEdge M830 blade servers, Dell Networking N3000 switches and VMware vSphere for virtualization.

The following sections provide an overview of the critical components of this reference architecture, as in Figure 2

Figure 2 Reference Architecture Components Overview



VMware vSphere

- ESXi 6.X
- High Availability Cluster
- Management with VMware vCenter Server

Dell Power Edge VRTX

- State-of-the-art chassis infrastructure
- 13th generation Dell Power Edge M830 Blade Servers
- Shared storage within the chassis
- 1 GbE Ethernet switch, 1 GbE Ethernet pass-through or 10 GbE 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 integrated 10GbE SFP+ ports
- Two GbE combo media ports for copper or fiber flexibility
- Enterprise class management and security



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 Open Manage Essentials

Table 1 describes the key components the architecture uses:

Table 1 Solution Components

Component	Details
Virtualization	
Hypervisor	VMware vSphere 6.x
Dell PowerEdge VRTX Chassis	
Hypervisor Host Servers	Two Dell PowerEdge M830 Each server with: Four Intel Xeon E5-4650 2.1Ghz Memory 64GB DDR 4 (This CPU and memory configurations are provided for guidance. Other supported CPU and memory configurations can also be used.) Two QLogic BCM57810S-k Network Daughter Cards (NDCs)
Networking	Dell PowerEdge VRTX 10Gb R1-2210 switch module
PCIe Devices	Two QLogic BCM57810 dual port PCIe NICs
Storage	Shared storage within PowerEdge VRTX Up to twenty-five 2.5" HDDs/SSDs or up to twelve 3.5" HDDs/SSDs
External Network Switches	
Switches	Two 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 6.x includes the ESXi hypervisor as well as vCenter Server, which administrators use to configure and manage VMware hosts. Key capabilities for the ESXi Enterprise Plus license level include the following items:

- **VMware vMotion:** VMware vMotion technology provides real-time migration of running VMs from one host to another with no disruption or downtime.
- **VMware High Availability:** VMware HA provides high availability at the VM level. Upon host failure, VMware HA automatically re-starts VMs on other physical hosts running ESXi. VMware vSphere 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 achieve automatic load balancing based on resource requirements. When VMs in a DRS cluster need fewer resources, such as nights and weekends, DPM consolidates workloads onto fewer hosts and powers-off idle hosts 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.

The architecture referenced in this document uses the VMware vCenter Server Virtual Appliance (vCSA). The vCSA provides an alternative option for organizations that chose not to run the Windows vCenter Server but still require centralized management of VMware vSphere deployments in the enterprise. The configuration of vCSA 6.0 is contingent on size requirements, shown in Figure 3.

Figure 3 vCSA Configuration Details

Size	vCPU	vRAM (GB)	Hosts (Max)	VMs(MAX)
Tiny	2	8	20	400
Small	4	16	150	3000
Medium	8	24	300	6000
Large	16	32	1000	10,000

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 within a chassis. The chassis also provides the consolidated power and cooling infrastructure for the servers, shared storage and the other components. Table 2 provides a summary of these characteristics.

Table 2 PowerEdge VRTX System Overview

Feature	Description
Server Nodes	Up to four server nodes
Supported Servers	Dell PowerEdge M520, M620, M630, M820 and M830
I/O Modules	One GbE switch, one GbE pass-through or one 10GbE switch
External PCIe Slots	Five low-profile Gen2 PCIe slots and three full-height Gen2 PCIe slots
Shared Storage	Up to twelve 3.5" SAS HDDs/SSDs or up to twenty-five 2.5" SAS HDDs/SSDs
Raid Controllers	One or two (redundant) Dell PowerEdge RAID Controllers
Management	One or two (redundant) Dell CMC
Power Supplies	Up to four 1100W PSUs
Form Factor / Placement	Standalone "Tower" (vertical position) or 5U Rack unit (horizontal)

Figure 4 PowerEdge VRTX System Overview



The PowerEdge VRTX also includes the following features:

- 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 4 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 share storage. 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 appears 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 directly 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 QLogic BCM57810 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 into the back of the system. The I/O module provides connectivity on Fabric A. The internal ports on the I/O modules connect to the network daughter card (NDC) in each PowerEdge M520, M620, M630, M820 or M830 blade server. The PowerEdge VRTX system supports the R1-2210 10Gb switch module, R1-2401 Gb switch module and R1-PT 1 Gb pass-through module in the I/O module slot.

PCIe Devices: The PowerEdge VRTX has eight PCIe slots on the back of the chassis. Out of the eight slots, five support low-profile form factor cards, while the remaining three slots support full-height form factor cards. Each PCIe slot is associated to either Fabric B or C in the chassis. The user can map the PCIe slots to servers via the CMC. A PCIe slot in Fabric B connects to the Fabric B of the server to which it is mapped. Similarly, 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, that is, each slot is dedicated to one server in the system. Table 3 provides a summary of the PowerEdge VRTX PCIe slot characteristics.

Table 3 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 2.5-inch disks (up to 25) or 3.5-inch disks (up to 12). 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 the creation of 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. This technology creates four virtual adapters (VA) on the PERC8 card. The PCIe switching technology within the PowerEdge VRTX chassis assigns each VA to one of the server slots to provide the server in that slot access to the shared storage. Each PERC driver within each server corresponds to the VA that maps to the associated 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 an 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.

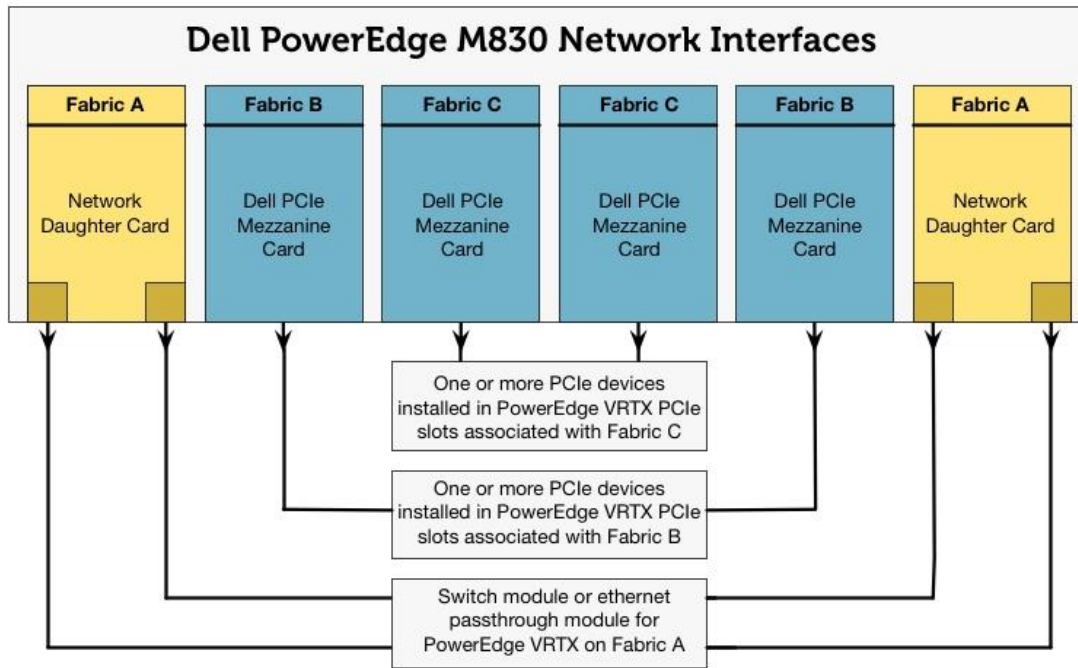
The CMC for PowerEdge VRTX provides a single interface for configuring, managing and monitoring the entire PowerEdge VRTX system. You can use it to control the power-up and power-down operations for the chassis, servers and the I/O module. It provides an interface to configure and manage the servers, PCIe resources and I/O module. It maps PCIe slots to servers, physical disks, RAID controller, virtual disks, mapping virtual disks, chassis power budget and so on. The CMC also provides an interface for general functions including, but not limited to, the following:

- Logging faults
- Sending alerts
- Controlling the front panel
- Chassis management
- Chassis configuration
- Chassis monitoring

Dell PowerEdge M830 Blade servers: The PowerEdge M830 is an Intel processor-based, 4-socket, full-height blade server that offers the following advantages:

- Intel E5-4600 series processors offer the greatest per node performance of any blade.
- Each PowerEdge M830 can deploy up to 3TB of RAM per server.
- Each server supports two dual-port network daughter cards (NDCs) and four PCIe mezzanine cards, as shown in Figure 5.
- Agent Free' management using the iDRAC with Lifecycle Controller allows customers to deploy, update, maintain and monitor their systems throughout the system lifecycle. Lifelong 'Agent Free' management using the iDRAC with Lifecycle Controller also does not require a software management agent, regardless of the operating system.
- Internal, dual SD modules to provide failover capability for embedded hypervisors.

Figure 5 Network Interfaces of PowerEdge M830 in PowerEdge VRTX



Embedded Management with Dell's Lifecycle Controller: The Lifecycle Controller is the engine for advanced embedded management and is part of iDRAC Enterprise in 13th-generation Dell PowerEdge blade servers. It includes 1GB of managed and persistent storage that embeds systems management features directly on the server. This eliminates the media-based delivery of system management tools and utilities previously needed for systems management. Embedded management includes the following features:

- Unified Server Configurator (USC) supports local, one-to-one deployment via graphical user interface (GUI) for OS install, update, configuration and 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 remotely located servers. 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 failed-parts replacement, thereby reducing maintenance downtime.

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 forty-eight 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. It offers simple management and scalability using an 84Gbps, high-availability stacking architecture that allows management of up to 12 switches from a single IP address. It also enables sharing of the dual SFP+ across

the stack for uplinks to your network's next layer. 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 and port mirroring. Additional features include broadcast storm control, loop-free redundancy without spanning tree using multi-chassis link aggregation (MLAG) and more. It also supports private VLAN extensions and private VLAN edge support.

Manageability features of the Dell Networking N3048 switch include QoS, multicast support and link aggregation. You can choose management using 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. It 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 design meets the requirements for remote offices, branch offices and small businesses. Some of its highlighted capabilities are as follows:

- Small form-factor for efficient use of space in an office environment while still providing an 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
- Shared power and cooling infrastructure in PowerEdge VRTX for energy efficiency
- Remote monitoring capability through Dell OpenManage Essentials (OME) for single-map view of the PowerEdge VRTX infrastructure

Virtualization: Each server configuration includes the processor, memory and network adapters required for virtualization. The solution includes an integrated VMware ESXi hypervisor, shared storage and two nodes to create a highly available cluster.

High availability: The design supports VMware vSphere HA clusters. Additionally, other aspects of the design, such as networking, power and VRTX cooling, support high availability.

Manageability: Dell's CMC provides a single management interface for configuring and managing server nodes, PCIe NICs, I/O Modules, shared storage and other PowerEdge VRTX components. The design also integrates enterprise management components.

Data protection integration: The design integrates Quest vRanger for data protection and provides specific recommendations for configuring vRanger for backup, replication and disaster recovery. The design supports multi-site, VM-level replication and remote-site disaster recovery, which are critical requirements for a remote or branch office.

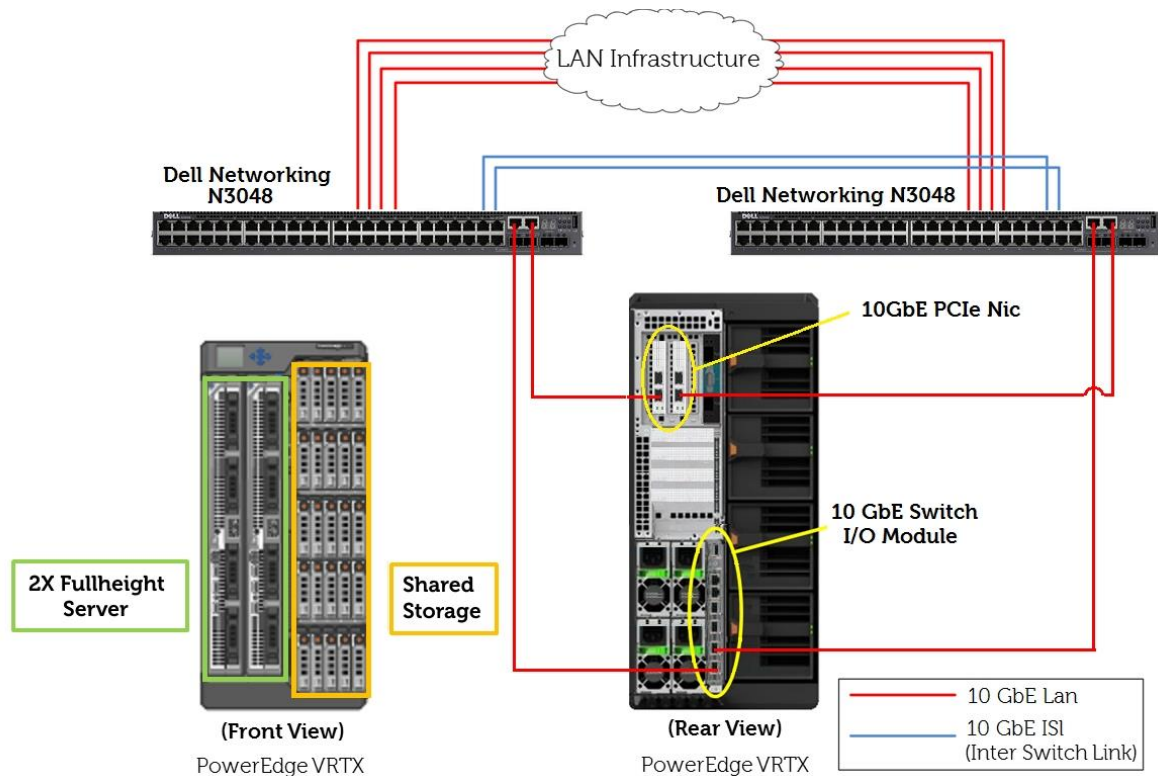
Cloud enablement: The design includes VMware vCloud Connector for cloud enablement and cloud connectivity.

Flexible configurations: The solution supports additional options for server processors, server memory and Disk RAID configurations for shared storage. This flexibility enables the customer to optimize compute and storage resources for specific workload requirements.

4 Reference Architecture

This solution consists of a Dell PowerEdge VRTX chassis populated with Dell PowerEdge M830 blade servers running VMware ESXi hypervisor, shared storage and networking. Figure 6 shows the major components and a high-level overview of the network connectivity in the reference architecture.

Figure 6 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 design equips each PowerEdge M830 blade server with two 10Gb Ethernet interfaces for LAN traffic. This section provides design and configuration details for the different physical networking subsystems 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. This section also provides information on virtual networking configuration.

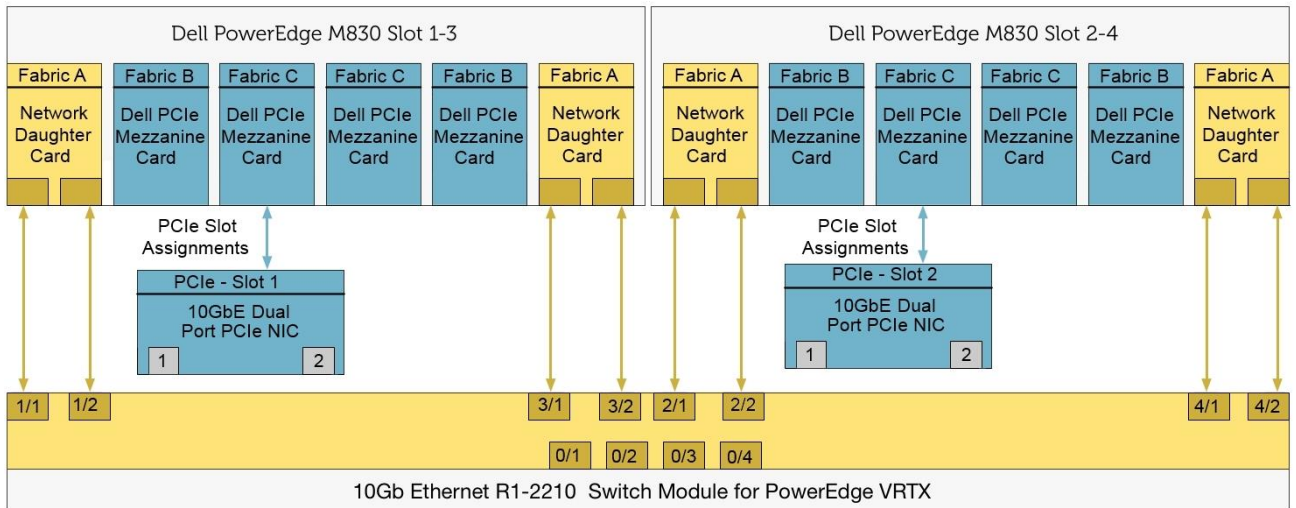
Server network ports: Each PowerEdge M830 blade server's configuration includes two QLogic BCM57810 dual-port NDC in Fabric A and one QLogic BCM57810 dual-port PCIe card in Fabric C. The solution also includes a Dell PowerEdge VRTX 10Gb R1-2210 switch module on Fabric A. This switch I/O module has 16 internal ports on Fabric A and six ports (four 10GbE and two GbE) for external connectivity. All available ports spilt across both QLogic NDCs connect to the corresponding internal ports of the R1-2210 10Gb switch module within the Dell PowerEdge VRTX chassis. Two QLogic BCM57810 dual port PCIe NICs reside in two full-profile PCIe slots in the PowerEdge VRTX chassis. Each QLogic BCM57810 dual port NIC maps to one of the PowerEdge M830 blade servers on Fabric C using the CMC. Table 4 provides the configuration of the PCIe slots in the PowerEdge VRTX chassis and their mapping to the server slots.

Table 4 PowerEdge VRTX PCIe Slots Configuration

PCIe Expansion Slot #	Fabric (B or C)	PCIe Device	Server Mapping
1	C	QLogic BCM57810 Dual Port 10GB	Server-1
2	C	QLogic BCM57810 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

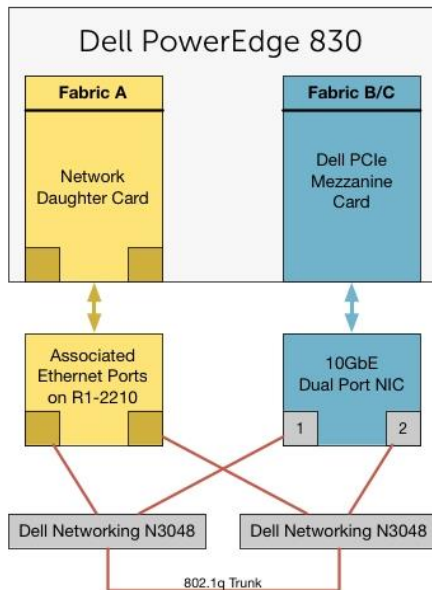
This physical networking configuration provides each M830 blade server with four 10Gb Ethernet ports on Fabric A through the R1-2210 10Gb switch module. It also provides two 10Gb Ethernet ports on Fabric C through a PCIe QLogic BCM57810 dual-port NIC mapped to each server. Because of this configuration, each M830 blade server has six 10Gb Ethernet ports for LAN traffic. Figure 7 summarizes this configuration.

Figure 7 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 uplinks to one of the two Dell Networking N3000 switches. This reference architecture configures all links as trunk ports carrying the appropriate VLAN IDs. Figure 8 provides an overview of how the network interfaces on each PowerEdge M830 blade server connect to the Dell Networking N3000 network switches.


Figure 8 Network Edge Connectivity



Uplink: There are several uplink options from the Dell Networking N3000 switches to the customer LAN infrastructure. The uplink option selected 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. Then create uplink LAGs from the Dell Networking N3000 switches to the core network. For additional 10 GbE SFP+ ports, the N3000 switches have an expansion module on the switch's rear for installation of a SFP+ Dual-Port 10GbE Module. Figure 9 provides an overview of the inter-switch-link and network uplink.

Figure 9 Network Uplink Overview



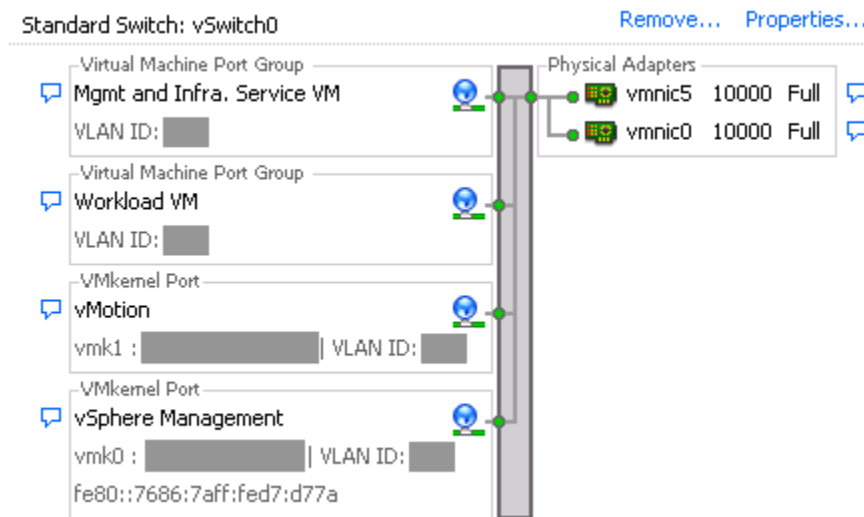
Redundancy: The physical network includes redundant configurations for all critical subsystems. These configurations avoid any single point-of-failure. For network connectivity, the PowerEdge M830 servers have a QLogic BCM57810s dual-port NDC in Fabric A, and QLogic BCM57810 dual-port NICs in Fabric 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 allows each PowerEdge M830 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 M830 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 the 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 10 provides the view of this vSwitch from vCenter Server.

Figure 10 VMware vSphere vSwitch Overview



Congestion management and failover capability: Each VMkernel connection or VM port group uses one or more active or standby VM NIC (VMNIC) 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 VM port groups have all VMNICs as active. This places vSphere management and vMotion traffic on separate physical ports during normal operating conditions. It also provides for load balancing of all VM traffic across all four VMNICs. This overall design brings high availability by providing failover capability for each VMkernel connection and each VM port group. The design provides this protection during any failure on a physical port, NDC, PCIe mezzanine card, PCIe NIC, Ethernet pass-through or Dell Networking N3000 switch. Table 5 summarizes the configuration of each connection to the vSwitch.

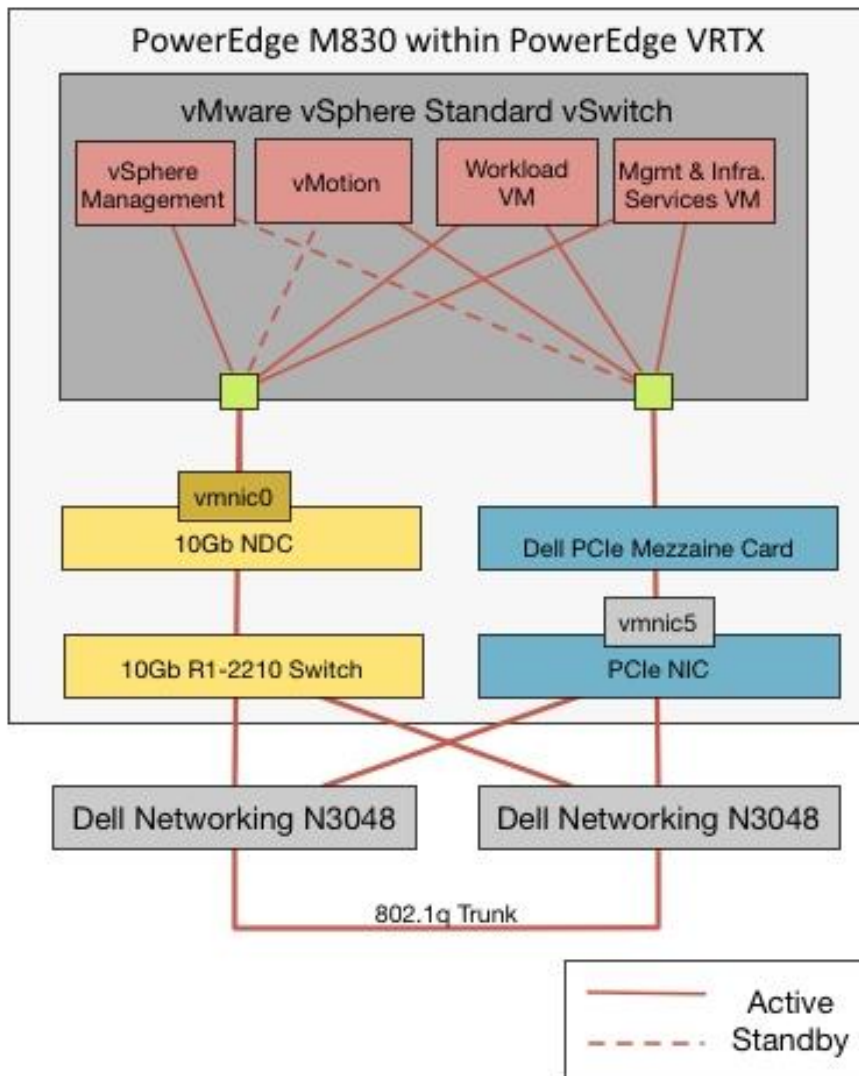
Table 5 Virtual Switch Configuration Details

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 at the vSwitch provides for load balancing the LAN traffic. Any given virtual network adapter uses only one physical adapter port at any given time. In other words, if a VM has only one virtual NIC, it uses only one physical adapter port at any given time. This option provides for easy configuration as well as load balancing across VMs, which is especially helpful with a large number of VMs.

Each host uses this vSwitch configuration, which Figure 11 describes in detail.

Figure 11 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 6

Table 6 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 includes a single VMware vSphere cluster spanned across both PowerEdge M830 blade servers, as described in section 7. This design allows all VMs to benefit from VMware vSphere High Availability and DRS capabilities. This reference architecture solution uses the storage within the PowerEdge VRTX chassis shared across both servers to enable this virtualization infrastructure.

You can create one or more shared virtual disks according to customer workload requirements. This solution also includes a shared Virtual Disk (VD) to host all management and infrastructure services, as described in section 6.2. Dell EMC recommends configuring a hard drive as the global hot spare in a shared storage configuration to provide data protection beyond the RAID configuration. To enable shared storage, assign each virtual adapter of the shared PERC8 controller to its corresponding server slot. This provides shared storage access to the PowerEdge VRTX in that slot. Table 7 provides the details of the virtual adapter assignment.

Table 7 PERC8 Virtual Adapter Mapping

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

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

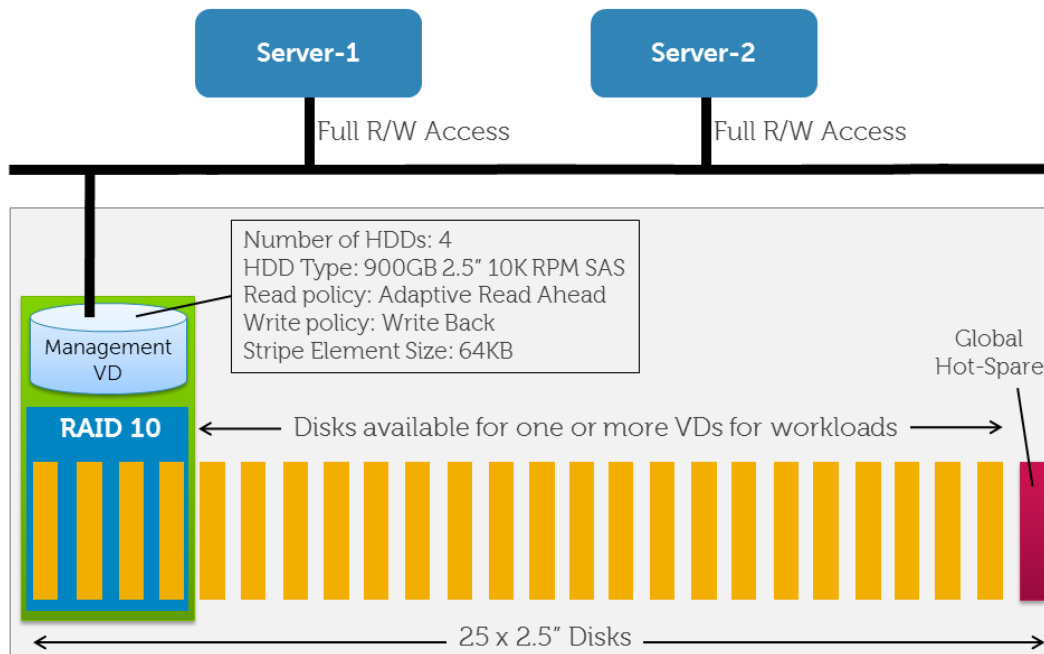
This reference architecture solution includes a shared VD configuration for management and infrastructure services, and additional storage customizable to meet the customer workload requirements. You can install the management and infrastructure services components in VMs and host them on the PowerEdge VRTX system along with customer workloads. See Chassis Management for further details on the solution's management and infrastructure services components, their sizing and deployment. Create a volume shared across all servers to host storage for all management and infrastructure services VMs. To enable the running of management and infrastructure VMs on the same VMware vSphere clusters as the workload VMs, allow access to all servers. This also enables these VMs to benefit from HA and DRS functionalities. The rest of this document refers to the shared VD for management and infrastructure services 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-inch HDDs. The shared storage configuration with twenty-five 2.5-inch HDDs provides the dense storage with capacity and IOPS suitable for a virtualization cluster. Create the

management VD on four 900GB, 10K RPM SAS HDDs. Dell EMC recommends configuring these disks with RAID 10 due to the criticality of the management and infrastructure services.

Provide full read/write access for the management VD to both PowerEdge M830 blade servers through the CMC user-interface. See Figure 12 for the details of the shared storage configuration and the management VD.

Figure 12 Management VD Configuration



6.2.1 Virtual Disks for Workloads

After creating the management VD, populate the remaining disk slots with supported HDDs or SSDs. Create one or more VDs on these HDDs or SSDs to support the different customer workloads. Use different disk size, speed, RAID levels and caching policies for different VDs, each to support workloads with different capacity, performance and data protection requirements

7 Virtualization Cluster

VMware ESXi resides on a pair of redundant SD cards in each M830 server, with the configured ESXi hosts on a vCenter Server for management. This solution considers a single VMware vSphere cluster spanned across both ESXi host servers. Enable VMware features, including vMotion, VMware High Availability (HA), and Distributed Resource Scheduling (DRS) if licensing allows.

Create a single VMware vSphere cluster for the following advantages:

- **High availability:** VMware vSphere High Availability (HA) provides high availability for applications running in VMs. In the event of physical server failure, affected VMs automatically restart. Creation of a single HA cluster across more than one server allows the cluster to survive failure of one or more servers.
- **Simplified management:** A single cluster allows convenient configuration and management of all hosts from vCenter server.

8 Chassis 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 VMs, 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 and 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 QLogic BCM57810 dual port NICs. As described in the reference architecture, the PCIe slots 1 and 2 are mapped to server slots 1 and 2 respectively.

Figure 13 PCIe Slot Assignment from CMC

PCIe Slot				Server Slot		Action
Slot	Name	Fabric	Power Status	Name	Slot	
1	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet [14e4:1006]	C	On		1	<input type="text" value="No Action"/>
2	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet [14e4:1006]	C	On		2	<input type="text" value="No Action"/>
3	Empty	B	N/A		N/A	<input type="text" value="No Action"/>
4	Empty	C	N/A		N/A	<input type="text" value="No Action"/>
5	Empty	C	N/A		N/A	<input type="text" value="No Action"/>
6	Empty	B	N/A		N/A	<input type="text" value="No Action"/>
7	Empty	B	N/A		N/A	<input type="text" value="No Action"/>
8	Empty	B	N/A		N/A	<input type="text" value="No Action"/>

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 and select **Chassis Overview > Storage** and click the Setup tab. As in Figure 14, VA 1 and VA 2 map to server 1 and server 2, respectively:


Figure 14 PERC8 Virtual Adapter Assignment from CMC

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

Virtual disk assignment Mode: CMC provides the capability to configure the virtual disk assignment mode. Multiple servers can access one virtual disk only after selecting the “Multiple Assignment” mode. To configure the virtual disk assignment mode, open the CMC and select **Chassis Overview > Storage**. Then click the Setup tab to display the CMC virtual disk Assignment Mode dialog, as in Figure 15:

Figure 15 CMC virtual disk Assignment Mode dialog

Assignment Mode	Description
<input type="radio"/> Single Assignment	This mode allows a virtual disk to be assigned to a single virtual adapter at a time
<input checked="" type="radio"/> 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

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** and select **Chassis Overview > Storage > Virtual Disk > Create** to display the CMC Virtual Disk creation dialog, as in Figure 16:

Note: The figure shows the properties for creating the management VD.

Figure 16 CMC Virtual Disk creation dialog

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

See below for capacity limits

Assign virtual disk access: CMC provides the capability to manage and configure access to each virtual disk for different servers. To assign virtual disk access, open the **CMC** and select **Chassis Overview > Storage > Virtual Disk > Assign**. Figure 17 shows the access to the management VD for each server. Dell EMC recommends granting full read/write access for each server to any new VD that hosts VMs.

Figure 17 Assigning Virtual Disk Access to Servers from CMC

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 ▼
				<input type="button" value="Cancel"/> <input type="button" value="Apply"/>

The details above help configuration of 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 CMC 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, the Dell management plug-in for vCenter, Dell OpenManage Essentials, Quest vRanger, VMware vCloud Connector, and so on. The infrastructure services may include Active Directory (AD), domain name service (DNS), network time protocol (NTP) server, and so on. This section provides recommendations for the specific management components and the deployment strategy for them as well as the necessary infrastructure services components.

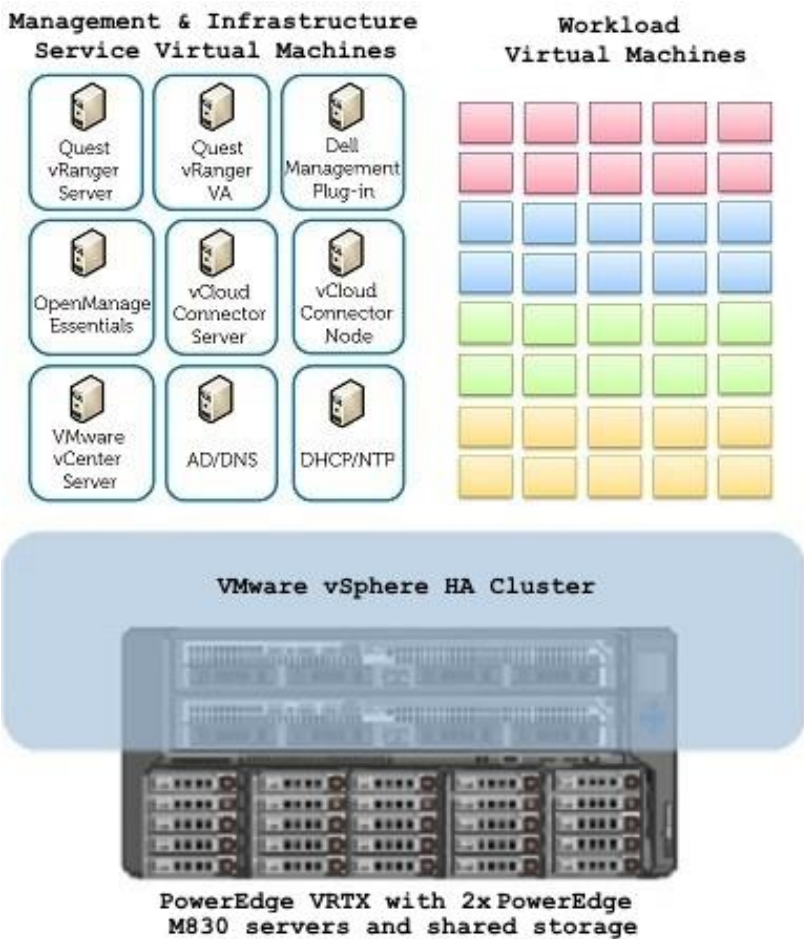
The necessary management and infrastructure services run as VMs on the same vSphere HA cluster as the workload VMs. Virtual switch configuration on each ESXi host handles the management and infrastructure services VM traffic on a dedicated port group (and dedicated VLAN, if desired). Section 5.2 describes virtual switch configuration on ESXi hosts. The PowerEdge VRTX shared storage provides the capability to create a separate VD for management and infrastructure services components, as described in [Storage Architecture](#). This design allows the customer to ensure availability of certain storage capacity and IOPS (and a specific RAID level) for management and infrastructure services.

The reference architecture solution includes the following management components:

- 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, infrastructure services components like Active Directory, Domain Controller, DHCP server, NTP server and so on, can run on the same infrastructure as the management components. As Figure 18 illustrates these components, are installed as VMs in the management infrastructure:

Figure 18 Management and Infrastructure Services Deployment



The product best practices and the requirements of the solution determine the sizing of each VM. Table 8 provides the guidelines for virtual processors (vCPU), Memory (vRAM) and disk configuration.

Most of these management components support a variety of operating systems. Resource calculations for management components, as Table 8 shows, consider Windows Server 2012 as the guest operating system, where 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 8 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 addition to these management components, the following items require further resources:

- **Infrastructure services components:** The components include some combination of Active Directory, Domain Controller or Read-Only Domain Controller, DNS, DHCP server, NTP server and so on. If the workload requires a SQL server, these components can include a separate installation of SQL server. The number of VMs and their sizing to support these infrastructure services depends upon the specific customer requirements and usage model.
- **VMware vCloud Connector sizing optimization:** Size vCloud Connector Node storage appropriately to hold the largest VM that customers plan to 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. Deployment of a virtual appliance allows scaling back up, restore and replication activities across multiple hosts or clusters. Virtual appliance deployment also allows simultaneous maintenance of 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.

The recommended sizing for the management VD, as described in [section 6.2](#), accommodates the management components that Table 8 describes, 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.

Take the resource requirements for the management and infrastructure services components into consideration while deploying the workload VMs to ensure these components always get the needed compute

and storage resources. The customer can use the vSphere Resource Pool capability to create a resource pool for the management and infrastructure services VMs, if supported by the chosen vSphere license. Use the information in Table 8 to determine the resource reservations for the resource pool. Using a resource pool helps 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 VM. The VM's size and configuration enables an installation of OME to monitor the components of this reference architecture solution. OME's configuration utilizes a local SQL Express database. For fullest functionality, Dell EMC recommends direct internet access, or internet access through a proxy.

OME is utilized for discovery, inventory and hardware level monitoring of PowerEdge VRTX chassis, PowerEdge M830 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 19 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 19 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). To provide ease-of-use of the appliance, firmware updates and warranty information, it is recommended 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:

- Retrieve in-depth details 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' VM and two 'node' VMs. 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:

- Expanded view across hybrid clouds. Use a single management interface that seamlessly spans your private vSphere and public vCloud environment.
- Extended 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 VM. The VM 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 VMs
- Quickly restores the entire VM or just specific files
- Manages disaster recovery strategies and protects critical data in virtual environments
- Provides the ability to replicate VMs 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 VM. 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 VM eliminates the need for dedicated hardware while maintaining high performance. In addition, vRanger has to be installed in a VM 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 VM.

- **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 two vCPUs 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 for a user-friendly inventory of all vCenter servers, ESXi hosts and VMs, and manages all backup, recovery and replication jobs.

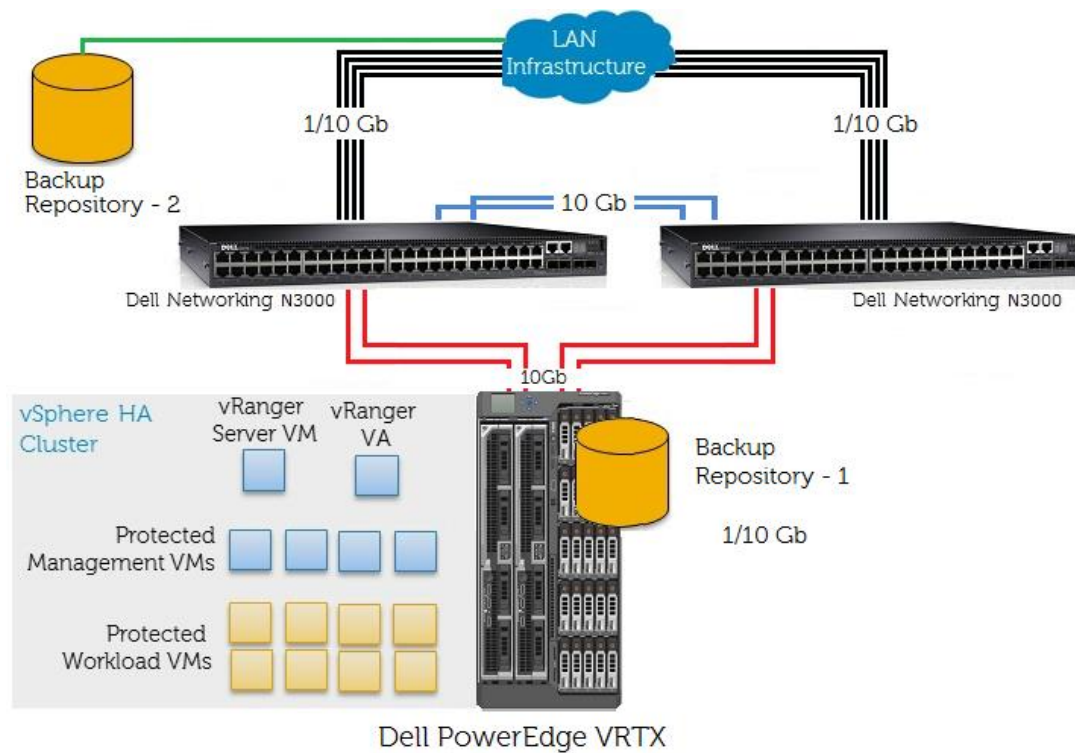
8.3 Backup and Recovery

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 VM. This VM can be configured on the management VD 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 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 20 below shows both of the options for creating a backup repository.

Figure 20 Backup and Recovery with Quest vRanger



8.4 Replication and Disaster Recovery

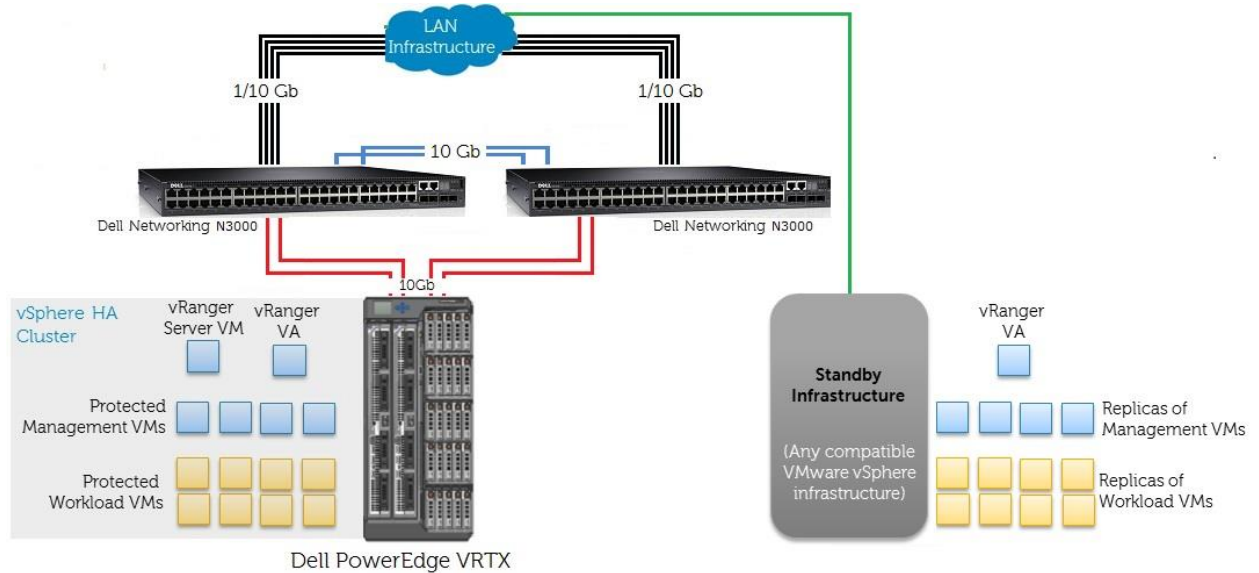
Quest vRanger supports replication of VMs via the vRanger virtual appliance (VA). A VM on PowerEdge VRTX can be replicated on a compatible ESXi host or a VMware vSphere cluster.

The vRanger server is installed in a VM. This VM 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.

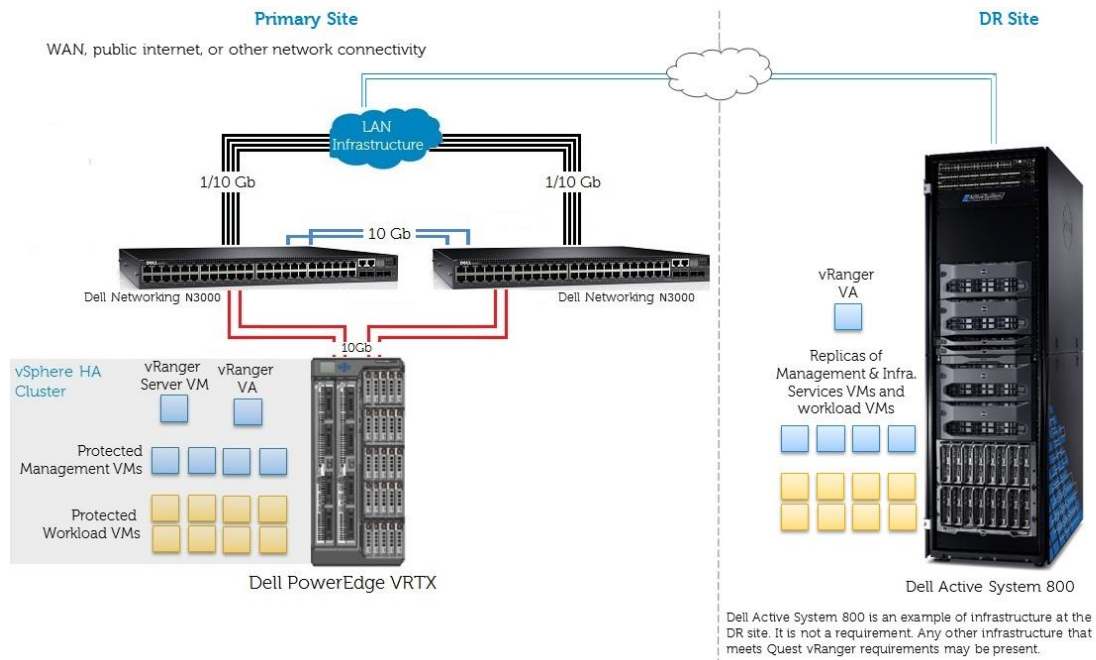
Replication at the local site: The VMs running on a PowerEdge VRTX can be replicated on a destination infrastructure at the local site. This destination infrastructure can be another PowerEdge VRTX system or any other infrastructure compatible with vRanger requirements. At least, one vRanger VA needs to be deployed on this destination infrastructure. Figure 21 below provides a representation of this replication scenario.

Figure 21 Replication with Quest vRanger



Replication at a remote site for disaster recovery: 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 22 below provides a representation of this replication scenario.

Figure 22 Disaster Recovery with Quest vRanger



9 Summary

This document provides a reference architecture for VMware vSphere based virtualization infrastructure for remote offices, branch offices and small businesses. The document 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 document provides configuration details for the PowerEdge M830 blade-server network interfaces, and for the shared PCIe slots, the I/O module slot and the shared storage in PowerEdge VRTX. The document also provides recommendations for virtual and physical network configuration. In addition, the document 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 document 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 EMC Customers and Dell EMC employees for sharing knowledge, best practices and information about Dell EMC 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](#)
- [Broadcom: Qlogic: What has changed with the Broadcom 57xx and 57xxx downloads for PowerEdge Servers](#)

About Dell EMC

Dell EMC is a worldwide leader in data center and campus solutions, which includes the manufacturing and distribution of servers, network switches, storage devices, personal computers, and related hardware and software. For more information on these and other products, please visit the Dell EMC website at <http://www.dell.com>.