

# Dell EMC Ready Stack Design Guide for VMware vSphere on Unity

April 2018

## Revisions

| Date       | Description     |
|------------|-----------------|
| April 2018 | Initial Release |

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA April 2018.

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

## TABLE OF CONTENTS

|       |   |    |
|-------|---|----|
| 1     | Introduction.....   | 5  |
| 1.1   | Executive Summary.....  | 5  |
| 1.2   | Description.....  | 5  |
| 1.3   | Audience and Scope .....  | 5  |
| 2     | Solution Overview .....   | 6  |
| 2.1   | Design Principles .....   | 6  |
| 2.2   | Design Overview.....  | 6  |
| 2.3   | Compute .....   | 7  |
| 2.4   | Networking.....   | 8  |
| 2.5   | Storage .....   | 9  |
| 2.6   | Supported Platforms.....  | 11 |
| 2.7   | Design Configurations .....   | 12 |
| 3     | Solution Architecture .....   | 15 |
| 3.1   | Solution Requirements .....   | 15 |
| 3.2   | Compute Design .....  | 15 |
| 3.2.1 | Compute Server.....   | 15 |
| 3.3   | Network Design .....  | 16 |
| 3.3.1 | Dell EMC Networking S5048-ON .....  | 16 |
| 3.3.2 | PowerEdge Rack Servers Network Configuration.....                               | 17 |
| 3.3.3 | VMware vSphere Distributed Switch (vDS) LAN Traffic Network Configuration ..... | 17 |
| 3.4   | Storage Design .....  | 18 |
| 3.4.1 | Dell EMC Unity All-Flash Models.....  | 18 |
| 3.4.2 | Storage Fabric Configuration.....   | 19 |
| 3.4.3 | Storage Connectivity for Compute and Management Servers .....                   | 20 |
| 3.4.4 | Dell EMC Unity x50F enclosures.....   | 20 |
| 3.4.5 | Storage Configuration.....  | 21 |
| 3.5   | Management Design .....   | 22 |
| 3.5.1 | Management Infrastructure .....   | 22 |
| 3.5.2 | Management Server .....   | 22 |
| 3.5.3 | VMware vCenter Server Appliance .....   | 23 |
| 3.5.4 | Dell EMC Virtual Storage Integrator .....                                       | 24 |
| 3.5.5 | Dell EMC OpenManage Integration for VMware vCenter .....                        | 24 |
| 3.5.6 | Dell EMC Avamar Virtual Edition.....  | 25 |
| 3.5.7 | Dell EMC Data Domain .....  | 26 |
| 3.5.8 | Management Software Resources .....   | 27 |
| 4     | Component Specifications.....   | 29 |

|         |  |    |
|---------|--|----|
| 5       | Scaling Dell EMC Ready Stack .....               | 35 |
| 6       | Dell EMC Ready Stack Validation .....            | 37 |
| 6.1     | Validation Strategy .....                        | 37 |
| 6.2     | Validation Methodology Overview .....            | 38 |
| 6.2.1   | Leveraging existing certification programs ..... | 38 |
| 6.2.2   | Versions Matrix .....                            | 38 |
| 6.2.3   | Deployment Validation.....                       | 39 |
| 6.2.4   | Simulate Workload.....                           | 39 |
| 6.2.5   | Design Validation.....                           | 40 |
| 6.2.5.1 | Host Layer.....                                  | 40 |
| 6.2.5.2 | Network Layer.....                               | 40 |
| 6.2.5.3 | Storage Layer .....                              | 41 |
| 6.2.5.4 | Stress Testing.....                              | 41 |
| 7       | Solution Summary .....                           | 42 |
| 8       | References .....                                 | 43 |

# 1 Introduction

With more choice and complexity, it is becoming increasingly difficult to manage IT infrastructure in a data center. Data center infrastructure has evolved over time and the impact of making changes has a deeper impact on stability, performance, costs and easier upgradability.

## 1.1 Executive Summary

This design guide for Dell EMC Ready Stack provides businesses with a reliable, industry-leading Dell EMC converged infrastructure (CI) solution to manage their virtualized workloads. Built around Dell EMC's latest 14G servers, Unity storage, S-Series switches and data protection products, this document provides design principles, best practices, architectural guidance and validated configurations for compute, management, networking, and storage.

This design guide also includes a framework for deploying VMware vSphere, a virtualization platform built for performance, availability and efficiency. Dell EMC Ready Stack is proven, tested, and optimized to help customers meet long-term data center needs for a variety of mixed workloads.

## 1.2 Description

Ready Stack provides the simplicity of a complete yet flexible validated CI stack based on the industry's leading building blocks. This design guide is based on the following Dell EMC components:

- Unity Storage
- PowerEdge Servers
- S-Series Switches
- VMware vSphere
- Data Domain & Avamar Data Protection

Dell EMC Ready Stack provides the following benefits:

- One trusted vendor for all infrastructure components (Compute, Storage, Networking and Data Protection) of the converged infrastructure stack
- Validation, interoperability testing and best practices incorporated into design and deployment guidance
- Virtualization planning and configuration burden reduced by providing pre-set configurations for easy quoting and ordering
- Design guidance built for performance, scale, flexibility and high availability
- Physical topology diagrams and cable connectivity guidelines incorporated in the reference architecture
- Unified management and system monitoring through VMware vCenter

## 1.3 Audience and Scope

The target audience for this design guide includes, but is not limited to channel partners, system integrators, service providers, IT managers, sales engineers and Do It Yourself (DIY) customers.

A working knowledge of virtualization technologies including servers, storage, networking and data protection is recommended.

## 2 Solution Overview

Installing, configuring, and running a production-ready, converged infrastructure involves multiple considerations, including:

- The appropriate operating system and virtualization software distributions
- Monitoring and management software
- Allocation of cluster services and data storage to physical nodes and arrays
- Selection of appropriate server hardware
- Design of the network fabric
- Sizing and scalability
- Performance

These considerations are complicated by the need to understand the types of workloads that will be running on the cluster, the fast-moving pace of the industry, and the challenges of managing a system designed to accommodate multiple general-purpose, virtualized workloads.

Dell EMC's customer-centered approach is to create rapidly-deployable and highly-optimized, end-to-end converged infrastructure solutions running on hyperscale hardware. Dell EMC listened to its customers and designed a solution that is unique in the marketplace, combining optimized hardware, software, and services to streamline deployment and improve the customer experience.

Dell EMC Ready Stack embodies all the hardware, software, resources and services needed to run a scalable, highly-available converged infrastructure environment. This end-to-end solution approach means that you can be operational in a shorter time than is typically possible with homegrown solutions.


### 2.1 Design Principles

The following principles are central to the design and architecture of the Dell EMC Ready Stack, which is built and validated using these design principles.

- **No single point-of-failure:** Redundancy is incorporated in the critical aspects of the solution, including server high availability features, redundant networking, and multi-path storage.
- **Integrated Management:** Provide integrated management of the Ready Stack using OpenManage Integration for VMware vCenter and Virtual Storage Integrator.
- **Hardware configuration for virtualization:** Dell EMC Ready Stack is designed for general-purpose virtualized workloads. Each server is equipped with appropriate processor, memory, Fibre channel, and Ethernet network adapters as required for virtualization.
- **Best practices adherence:** Storage, networking and vSphere best practices of the corresponding components are incorporated into the design to ensure availability, serviceability and optimal performance.
- **Flexible configurations:** The Dell EMC Ready Stack can be configured to suit most customer needs for a virtualized infrastructure. The solution supports flexibility in form of various options, such as rack server model, number of compute servers, server processor model, server memory capacity, type of Fibre channel storage model, and Fibre channel storage capacity based on customer needs.

### 2.2 Design Overview

This section provides an overview of the Dell EMC Ready Stack architecture applicable to 14G PowerEdge servers and Unity storage. The following figure provides a high-level overview of the architecture, including compute servers (showing flexible compute nodes), management servers, LAN switches, SAN switches, storage arrays and out-of-band switches:



**Figure 1. Dell EMC Ready Stack – Design Overview**

## 2.3 Compute

A complete portfolio of Dell EMC PowerEdge rack servers is designed to optimize application performance and ensure a stable, worry free environment with intuitive tools that simplify and automate throughout the entire server lifecycle. Every datacenter and virtual environment has unique requirements, and PowerEdge provides the flexibility needed to build scalable infrastructure that is tightly integrated with VMware vSphere.

The PowerEdge R440 delivers a combination of affordability, performance and density with a feature set right-sized for scale-out infrastructure environments. It is built on a scalable architecture that provides the choice and flexibility to optimize performance and density. The PowerEdge R440 is a 1U, 2 socket server platform, which supports up to 16 DIMM slots, 10 2.5” drives and 2 PCIe cards.

The PowerEdge R640 is an ideal dual-socket, 1U platform for dense scale-out data center computing. The R640 combines density, performance and scalability to optimize application performance and data center density. The scalable business architecture of the R640 is designed to maximize application performance and provide the flexibility to optimize configurations based on the application and use case. The PowerEdge R640 is a 1U, 2 socket server platform, which supports up to 24 DIMM slots, 10 2.5” drives, and 3 PCIe cards.

The PowerEdge R740/740xd delivers a balance between storage scalability and performance. The 2U two-socket platform is ideal for software defined storage, service providers or as virtual desktop infrastructure. The PowerEdge R740 was designed to accelerate application performance leveraging accelerator cards and storage scalability. The R740xd versatility is highlighted with the ability to mix any drive type to create the optimum configuration of NVMe, SSD and HDD for either performance, capacity or both. The PowerEdge R740/R740xd is a 2U, 2 socket server platform, which supports up to 24 DIMM slots, 32 2.5” drives, and 8 PCIe cards.

The integrated Dell EMC Remote Access Controller 9 (iDRAC9) with Lifecycle Controller is the next generation of remote server administration. Embedded in every PowerEdge server, iDRAC 9 provides secure and remote server access for a multitude of common management functions. The iDRAC with Lifecycle Controller operates regardless of operating system state or the presence of a hypervisor and

offers a complete set of server management features including configuration, OS deployment, firmware updates, health monitoring and maintenance.

The iDRAC 9 provides a variety remote connectivity interfaces and protocols to allow administrators to securely configure, deploy, manage, monitor and update the server. IPMI, Redfish, SMASH-CLP, and WS-Man are just a few of the common standard management interfaces that are supported. Additionally, the iDRAC 9 HTML5 Web GUI allows connectivity securely through HTTPS. With an iDRAC9 Enterprise license, Group Manager is a new 14G iDRAC 9 feature that creates a one-to-many console experience, allowing users to view the details of a set of servers by permitting more powerful management than by inspecting servers visually for faults and other manual methods.




Figure 2. Dell EMC iDRAC 9 GUI

## 2.4 Networking

For Ethernet network traffic, Dell EMC Ready Solutions uses the Dell EMC Networking S-Series S5048F-ON switch. The S5048-ON is a high-performance, multi-function, 10/25/40/50/100 GbE TOR switch purpose-built for applications in high-performance data center, cloud and computing environments. The compact S5048F-ON model design provides industry-leading density with up to 72 ports of 25GbE or up to 48 ports of 25GbE and 6 ports of 100GbE in a 1RU form factor. Using industry-leading hardware and a choice of Dell's OS9 or select 3<sup>rd</sup> party network operating systems and tools, the S5048F-ON delivers non-blocking performance for workloads sensitive to packet loss. 3.6 Tbps (full-duplex) of non-blocking, store and forward switching fabric delivers line-rate performance under full load. Key features include scalable L2 and L3 Ethernet switching with QoS and a full complement of standards-based IPv4 and IPv6

features, including OSPF and BGP routing support. L2 multipath is supported via Virtual Link Trunking (VLT) and multiple VLT (mVLT) multi-chassis link aggregation technology.

Dell Networking OS 9 has been tested and hardened to meet stringent requirements for reliability, scalability and serviceability. OS 9 supports the full portfolio of Dell Networking data center switch products and enables you to build cost-effective, end-to-end networks while reducing operational complexity. A fully compliant, industry-standard CLI allows certified engineers to be productive from day one. It provides the primary method to configure, monitor and administer Dell OS 9 applications and Dell Networking switches.

For Fibre Channel storage traffic, Dell EMC Ready Stack uses the Dell EMC Connectrix DS6500B series switches. DS6500B switches deliver up to 16 Gbps Fibre channel performance and scale from 12 to 96 ports. This document focuses on, but is not limited to, the DS6505B and DS6510B models. The DS6505B comes standard with twelve 8/16Gbps ports enabled; however, the Ports on Demand (POD) license can enable the remaining twelve ports (twenty-four ports in total). The DS6510B comes standard with twenty-four 8/16Gbps ports enabled; however, the Ports on Demand (POD) license can enable the remaining twenty-four ports (forty-eight ports in total).

Brocade Web Tools is an embedded graphical user interface (GUI) on the Connectrix DS6500 switches that enables administrators to monitor and manage single or small fabrics, switches, and ports. Web Tools is launched directly from a web browser, or from the Brocade Network Advisor.




Figure 3. Brocade Web Tools GUI

## 2.5 Storage

The Dell EMC Unity All-Flash product line sets new standards for storage with compelling simplicity, modern design, flexible deployments and affordable prices. Dell EMC Unity All-Flash storage systems implement an integrated architecture for block, file, and VMware VVols with concurrent support for native

NAS, iSCSI, and Fibre Channel protocols. Each system leverages dual storage processors, full 12 Gb SAS back end connectivity and Dell EMC's patented multicore architected operating environment to deliver performance & efficiency. Additional storage capacity is added via Disk Array Enclosures (DAEs) and for additional performance, online controller upgrades are available. Dell EMC Unity All-Flash supports file and block environments, point-in-time snapshots and thin clones, synchronous and asynchronous replication, built-in encryption, tiering to the cloud and deep ecosystem integration with VMware, Microsoft and OpenStack.

Dell EMC Data Reduction includes deduplication and compression logic within the space savings algorithm. Dell EMC Unity Data Reduction features are available for block LUNs and VMFS datastores in an All-Flash pool. Dell EMC Unity Data Reduction can be enabled at the resource-level, so that storage resources with and without data reduction enabled can exist in the same pool on a system. Also, data reduction savings can be seen at multiple levels including system-level, pool-level, and resource-level thereby providing varying levels of granularity in reporting for management purposes and helps show the business value of utilizing the advanced data feature.

Dynamic Pools increases the flexibility of configuration options within a Dell EMC Unity system with an entirely redesigned Pool structure. Dynamic Pools replace the existing Pool technology as the default Pool type created within Unisphere for Dell EMC Unity All Flash systems. Dynamic Pools, as with Traditional Pools, can be created, expanded, and deleted, but include other improvements. Expanding a Dynamic Pool can be as granular as a single drive to the Pool to increase its capacity. This provides a completely flexible deployment model, which improves the planning and provisioning process and reduces the total cost of ownership.

Host I/O Limits is the ability to limit the amount of I/O activity that is serviced by the Dell EMC Unity system. Host I/O Limits can be applied on LUNs, VMware VMFS Datastores, Thin Clones, and their associated attached snapshots. Use Host I/O Limits to limit incoming host activity on the basis of IOPS, Bandwidth, or both. Consider configuring Host I/O Limits on LUNs or datastores that are consuming a large portion of the system's resources and reducing the performance of other resources on the system. Another feature in Dell EMC Unity is burst policies for Host I/O Limits. This option is available for any Host I/O Limit policy and allows host activity/traffic to exceed the base policy limit by a user-specified percentage. The user specifies parameters of amount (percentage), duration, and frequency of the I/O limit burst allowance.

VMware vStorage APIs for Array Integration enable the Dell EMC Unity system to handle offloaded VMware operations, improving performance and reducing network overhead. VMware Aware Integration (VAI) and VMware vStorage APIs for Storage Awareness (VASA) enable inter-communication between the storage and virtualization layers so that administration of one layer can be performed from the other. This can enable more efficient workflows for those familiar with one interface. Furthermore, the VASA protocol forms the basis on which VMware Virtual Volumes are constructed. VMware vStorage APIs for Array Integration (VAAI) improves the utilization rates of ESXi hosts in a VMware environment by offloading tasks to the Dell EMC Unity system. For both block and file datastores as well as Virtual Volumes, various storage operations are processed by the Dell EMC Unity system. This includes clone and snapshot operations, as well as hardware-assisted moves. Using VAI, Unisphere discovery and registration of a VMware environment is simple and efficient. The information imported provides relevant detail for the user, who is no longer required to switch between VMware vSphere and Unisphere to obtain a clear picture of their environment.

Dell EMC Unisphere, the HTML5 graphical management interface for Dell EMC Unity, can be used to present storage to all host servers. For vSphere hosts, Unisphere utilizes integration points to reduce required steps, perform several host-side functions and the presentation of VMware VMFS, NFS, or VVOL Datastore storage. System status and performance information can also be displayed in Unisphere, with a graphical representation of the Dell EMC Unity system provided and highlighted for areas of interest such as drive faults, network link failures, etc.




Figure 4. Dell EMC Unisphere GUI

## 2.6 Supported Platforms

The Dell EMC Ready Stack utilizes Dell EMC's latest enterprise-level rack servers, LAN and SAN switches, and All-Flash storage arrays; as well as state-of-the-art management software, carefully selected and tested to provide outstanding performance and value.

Table 1 provides an overview of the supported platforms.


Table 1 Supported Platforms

| Component        | Details  |
|------------------|--|
| Server Platforms | Dell EMC PowerEdge R440  |
|                  | Dell EMC PowerEdge R640  |
|                  | Dell EMC PowerEdge R740/R740xd   |
| LAN Connectivity | Dell EMC Networking S5048-ON 25/40/50/100GbE Switches                      |
| SAN Connectivity | Dell EMC Connectrix DS6500 (6505B or 6510B) 16 Gbps Fibre Channel Switches |

| Component                      | Details  |
|--------------------------------|--|
| Out-of-band (OOB) Connectivity | Dell EMC Networking S3048-ON 10/100/1000Mb Switch                    |
| Storage Array                  | Dell EMC Unity 350F, 450F, 550F, or 650F All-Flash Storage Platforms |
| Management Server Platforms    | Dell EMC PowerEdge R640  |
|                                | Dell EMC PowerEdge R440  |
| Management Software Components | VMware vCenter Server Appliance (VCSA)                               |
|                                | Dell EMC OpenManage Integration for VMware vCenter (OMIVV)           |
|                                | Dell EMC Virtual Storage Integrator (VSI)                            |
|                                | Dell EMC Avamar Virtual Edition (AVE)                                |
|                                | Dell EMC Data Domain Virtual Edition (DD VE)                         |

## 2.7 Design Configurations

Dell EMC Ready Stack offers flexibility to address various virtualization use-cases and workloads. Sizing a physical infrastructure for all of customer's virtualized workloads requires knowledge of all the individual workloads and their compute, memory, storage and network performance requirements. Generally, this can be estimated by creating a virtual machine profile that is an average of all the different individual workloads that will run in the environment. Based upon this VM profile, the total number of VMs needed in the environment, and the amount of future growth expected, the physical infrastructure can be approximated. It's also important to factor server and storage maintenance requirements and virtualization features, such as high-availability (HA), when finalizing on sizing.



**Figure 5. Dell EMC Ready Stack Enterprise Large Configuration**

For the purposes of this document, we have proposed these small, medium and large sample configurations based upon following Balanced VM profile.

**Balanced VM Profile: 2 vCPU, 6 GB Memory, 60 GB Virtual Disk, 50 IOPS**

- Enterprise Small (250 VMs, 500 vCPU, 1.5TB vRAM, 15TB vDisk, 12,500 IOPS totals)
- Enterprise Medium (500 VMs, 1000 vCPU, 3TB vRAM, 30TB vDisk, 25,000 IOPS totals)
- Enterprise Large (1000 VMs, 2000 vCPU, 6TB vRAM, 60TB vDisk, 50,000 IOPS total)

We chose to size our configurations based upon a vCPU to CPU over-subscription of 4:1. Based upon these assumptions, the configurations listed below provide the physical infrastructure resources necessary to accommodate the total number of virtual machines indicated.

Table 2 Dell EMC Ready Stack Sample Configurations

|                | Enterprise Small | Enterprise Medium | Enterprise Large |
|----------------|------------------|-------------------|------------------|
| Storage Fabric | Fibre channel    | Fibre channel     | Fibre channel    |

|   | <b>Enterprise Small</b> | <b>Enterprise Medium</b> | <b>Enterprise Large</b> |
|---|-------------------------|--------------------------|-------------------------|
| Compute Server<br>(R440, R640, or R740)             | (4) PowerEdge R440      | (7) PowerEdge R740       | (13) PowerEdge R640     |
| Storage<br>(350F, 450F, 550F, or 650F)              | Unity 350F All-Flash    | Unity 450F All-Flash     | Unity 550F All-Flash    |
| LAN Switches  | (2) S5048F-ON           | (2) S5048F-ON            | (2) S5048F-ON           |
| SAN Switches<br>(DS6505B or DS6510B)                | Connectrix DS6505B      | Connectrix DS6505B       | Connectrix DS6510B      |
| OOB Switch  | None                    | S3048                    | S3048                   |
| Management Servers<br>(None, Two R440, or Two R640) | None                    | (2) R640                 | (2) R640                |
| Hypervisor  | ESXi 6.5                | ESXi 6.5                 | ESXi 6.5                |

## 3 Solution Architecture

This chapter summarizes all the main components of a production Ready Stack solution, including the physical compute server hardware, network fabric, storage infrastructure, software layers, and ongoing management.

### 3.1 Solution Requirements

Dell EMC Ready Stack depends on certain infrastructure elements within the existing data center implementation. The following list summarizes the infrastructure elements that are required.

1. Domain Name Server (DNS) must be available on the management network.
2. Network Time Protocol (NTP) Server must be available on the management network.
3. An existing Ethernet infrastructure to integrate with is required. 10/25 GbE and/or 40/100 Gb Ethernet infrastructure is recommended. Additional components, such as Dell Networking cables and transceivers, are needed to uplink the solution to the customer network. The necessary components depend upon customer networking and uplink requirements.
4. Sufficient power and cooling to support the solution must be present.

### 3.2 Compute Design

The latest Intel® SP Skylake generation processors power Dell EMC Ready Stack. With up to 28 cores per CPU and clock speeds of up to 3.6 GHz in the Dell EMC PowerEdge rack servers you can reduce CPU socket-based licensing costs and achieve greater VM density. Scale compute resources with Intel® Xeon® Scalable processors, delivering a 27% increase in processing cores and 50% increase in bandwidth over previous generation of Xeon processors.

The Dell EMC PowerEdge rack server platforms support the RDIMM and LRDIMM memory types. Load Reduced DIMM (LRDIMM) uses an iMB buffer to isolate electrical loading from the host memory controller. This buffer and isolation allows for the use of quad ranked DIMM to increase overall memory capacity. For general-purpose virtualization solutions, 2666 MT/s RDIMMs are recommended. Memory can be configured in various modes from within the BIOS. Optimizer Mode is the default mode and is recommended for most virtualization use cases to provide the optimized memory performance. For improved reliability and resiliency, other modes such as mirror mode and spare mode are available.

The Mellanox CX4 LX 10/25 Gb dual-port Ethernet adapter is the recommend adapter for Dell EMC Ready Stack for virtual workload and host communication. The QLogic 2692 Dual Port 16 Gb Fibre Channel Adapter is the recommended adapter for Fibre channel SAN storage connectivity.

Dell EMC PowerEdge servers support various BIOS configuration profiles that control the processor, memory, and other configuration options:

- Enable the Performance Optimized Virtualization profile
- Disable processor C-States and C1E to ensure the highest performance in a virtualized environment
- Enable Intel Hyperthreading and Virtualization features

#### 3.2.1 Compute Server

The Dell EMC Ready Stack offers customers a selection of rack servers for their compute infrastructure. Table 3 details the supported compute servers.

Table 3 Compute Server Configurations

| Components                       | Details  |
|----------------------------------|--|
| Server Platform Models           | PowerEdge R440, R640, R740, and/or R740xd                                      |
| Processors                       | (1 or 2) Intel® Xeon® SP Bronze, Silver, Gold, and Platinum Skylake processors |
| Memory (depends on server model) | (1 – 24) 2666 MT/s RDIMMs, up to 768 GB of RAM                                 |
|                                  | (1 – 24) 2666 MT/s LRDIMMs, up to 3 TB of RAM                                  |
| Network Adapter                  | Mellanox CX4 LX 10/25Gb Ethernet Network Adapter                               |
| Host Bus Adapter                 | QLogic® 2692 Dual Port Fibre Channel Adapter                                   |
| Boot Device                      | Internal Dual SD Module (IDSDM)  |
|                                  | PERC H730p RAID Controller   |
| Out-of-band Server Management    | iDRAC9 Enterprise  |
| Hypervisor                       | VMware ESXi 6.5  |


### 3.3 Network Design

This section provides an overview of the network architecture including compute and management server connectivity. Details around the Top-of-Rack (TOR) and virtual switch configuration are provided.

#### 3.3.1 Dell EMC Networking S5048-ON

The network architecture employs Virtual Link Trunking (VLT) connection between the two Top-of-Rack (TOR) switches. The inherent redundancy of a non-VLT environment requires standby equipment, which drives up infrastructure costs and increases risks. In a VLT environment, all paths are active, adding immediate value and throughput while still protecting against hardware failures. VLT technology allows a server or bridge to uplink a physical trunk into more than one Dell EMC Networking S5048-ON switch by treating the uplink as one logical trunk. A VLT connected pair of switches acts as a single switch to a connecting bridge or server. Both links from the bridge network can actively forward and receive traffic. VLT provides a replacement for Spanning Tree Protocol (STP) based networks by providing both redundancy and full bandwidth utilization using multiple active paths. Major benefits of VLT technology are:

- Dual control plane for highly available resilient network services
- Full utilization of the active LAG interfaces
- Active / Active design for seamless operations during maintenance events




(2) Dell Networking S5048

**Figure 6. Dell EMC Networking S5048-ON Virtual Link Trunk Interconnect (VLTi) Configuration**

The Dell EMC Networking S5048-ON switches each provide six 40/100 GbE uplink ports. The Virtual Link Trunk Interconnect (VLTi) configuration in this architecture uses two 40/100 GbE ports from each Top-of-Rack (TOR) switch to provide a 200 Gb data path between the switches. The remaining four 40/100 Gb ports allow for high speed connectivity to spine switches or directly to the data center core network infrastructure. They can also be used to extend connectivity to other racks.

### 3.3.2 PowerEdge Rack Servers Network Configuration

The compute cluster consists of Dell EMC PowerEdge rack servers. This section describes the network connectivity if rack servers are used for compute servers, and the management servers. The following image is an example of the connectivity between the compute and management Dell EMC PowerEdge rack servers and Dell EMC Networking S5048-ON switches. The compute and management rack servers have two 10/25 GbE connections to S5048-ON switches through one Mellanox ConnectX-4 LX dual port 10/25 GbE network card.



**Figure 7. PowerEdge Rack Server Connectivity**

### 3.3.3 VMware vSphere Distributed Switch (vDS) LAN Traffic Network Configuration

Customers can achieve bandwidth prioritization for different traffic classes such as host management, vMotion, and VM network using VMware Distributed Virtual Switches. The VMware vSphere Distributed Switch (vDS) can be configured, managed, and monitored from a central interface and provides:

- Simplified virtual machine network configuration
- Enhanced network monitoring and troubleshooting capabilities
- Support for network bandwidth partitioning when NPAR is not available

The following Figure 8 shows the virtual distributed switch configuration for the management and compute servers.




Figure 8. Distributed Virtual Switch for Dual Port Configuration

## 3.4 Storage Design

The Dell EMC Ready Stack uses the Dell EMC Unity x50F All-Flash Storage Platform with vSphere integration connected to a 16 Gb Fibre channel storage network. The Dell EMC Unity x50F All-Flash Storage Platforms is a flexible storage product, purposely designed for All-Flash. With all-inclusive software, these All-Flash systems deliver consistent performance with low response times and is a perfect fit for mixed virtual workload requirements. Two Connectrix DS6505B or 6510B series switches make up the Fibre Channel fabrics.

### 3.4.1 Dell EMC Unity All-Flash Models

Table 4 below shows a comparison of the Dell EMC Unity x50F All-Flash storage arrays that are supported for the Dell EMC Ready Stack solution. The Disk Processor Enclosure contains up to twenty-five SSD drives and includes two Storage Processors. Additional drive enclosures can be used to expand overall storage capacity. Additional Information can be found in the Unity Specification Sheet.


Table 4 Dell EMC Unity All-Flash Systems Comparison

| Components             | Unity 350F  | Unity 450F                     | Unity 550F                               | Unity 650F                               |
|------------------------|---|--------------------------------|--|--|
| Min/Max Drives         | 6/150   | 6/250                          | 6/500                                    | 6/1000                                   |
| Array Enclosure        | A 2U Disk Processor Enclosure (DPE) with twenty-five 2.5" drives                  |                                |  |  |
| Drive Enclosure        | All models support 2U twenty-five drive and 3U eighty drive trays for 2.5" drives |                                |  |  |
| RAID Options           | 1/0, 5, 6   |                                |  |  |
| SAS IO Ports per Array | 4 Embedded x 4-lane 12Gb/s SAS  | 4 Embedded x 4-lane 12Gb/s SAS | 4 Embedded x 4-lane 12Gb/s SAS           | 4 Embedded x 4-lane 12Gb/s SAS           |
|                        |   |                                | 6 x 4-lane, or 2 x 4-lane with IO Module | 6 x 4-lane, or 2 x 4-lane with IO Module |
|                        | 4 Embedded ports: 8/16 Gb FC, 10Gb IP/iSCSI, or 1Gb RJ45                          |                                |  |  |

| Components                              | Unity 350F                       | Unity 450F | Unity 550F | Unity 650F |
|---|----------------------------------|------------|------------|------------|
| CNA Ports per Array                     | 4 Additional ports per IO Module |            |            |            |
| Max Raw Capacity (drive type dependent) | 2.4 PBs                          | 4.0 PBs    | 8.0 PBs    | 16.0 PBs   |
| IOPS (100% Reads, 8K blocks)            | Up to 130K                       | Up to 305K | Up to 395K | Up to 440K |
| IO Modules per Array                    | Up to 4                          | Up to 4    | Up to 4    | Up to 4    |

### 3.4.2 Storage Fabric Configuration


This solution is configured with two FC fabrics for high availability. For the Dell EMC Unity x50F arrays, FC port 0 from each controller connects to the FC fabric switch A, while port 1 connects to FC fabric switch B. Dell EMC Unity x50F arrays have expansion slots that can provide additional front end (FC) or back end (mini SAS HD) ports. However, the connectivity shown in the below diagram does not cover the usage of such expansion modules. Please refer to Dell EMC Unity Hardware Information Guide for additional details.



**Figure 9. Dell EMC Unity x50F Storage Fabric Configuration**

### 3.4.3 Storage Connectivity for Compute and Management Servers

Each management and rack compute server is configured with a QLogic® 2692 dual port Fibre Channel adapter for connecting to the storage fabrics, with each port connecting to the Connectrix switch as shown in Figure 10.



**Figure 10. PowerEdge Rack Server SAN Connectivity**

### 3.4.4 Dell EMC Unity x50F enclosures

Dell EMC Unity x50F series storage array is equipped with two back end (BE) buses that utilize mini SAS HD connectivity. When connecting additional enclosures, they should be connected so that the load will be balanced equally between the available buses. Dell EMC Unity x50F arrays have expansion slots that can provide additional front end (FC) or back end (mini SAS HD) ports. However, the connectivity shown in the below diagram does not cover the usage of such expansion modules. Please note that the Disk Processor Enclosure is on Bus 0. Therefore, the first expansion enclosure should be on Bus 1 while the second expansion enclosure should be on Bus 0 and so forth.




Figure 11. Dell EMC Unity x50F Storage Array Backend Connectivity

### 3.4.5 Storage Configuration

For both rack server platforms, each server's HBA port 1 connects to the Brocade switch 1, while HBA port 2 connects to the Brocade switch 2. These ports are then zoned with the Dell EMC Unity array target ports to enable storage access for the hypervisor hosts.

Multiple datastores within the VMware vSphere cluster enable vSphere HA Datastore Heartbeat. This ensures that partitioning, or isolated host networks, do not trigger VM movement within the cluster. By default, the vSphere cluster selects up to five datastores for the Datastore Heartbeat.

Block storage presented to vSphere hosts from Dell EMC Unity has the native Path Selection Policy (PSP) of round robin (RR) applied by default. While RR is the recommended PSP to apply to Dell EMC Unity block storage, the default number of I/Os between switching paths is 1000. By reducing this value, all paths are more efficiently utilized.

VMware currently supports a maximum datastore size of 64 TB and 2048 powered on virtual machines per VMFS datastore. However, in most circumstances and environments, a target of 15-25 virtual machines per 500-750GB sized datastore is the conservative recommendation. LUNs and vSphere

datastores can easily be expanded to address future growth and by maintaining a smaller number of virtual machines per datastore, the potential for I/O contention is greatly reduced. This results in more consistent performance across the environment.

Dell EMC Unity offers thin provisioning as a recommended option when creating block storage and requires thin provisioning to enable compression. Utilizing thin provisioning within VMware on virtual disks does not initially result in additional space efficiency when thin provisioning is enabled on the array. However, the ability to reclaim space from within a compatible guest OS requires thin provisioning be used on both the storage and the virtual disks.

## 3.5 Management Design

This section provides an overview of the Dell EMC Ready Stack's management infrastructure, and the software components that run on virtual machines within the management cluster.

### 3.5.1 Management Infrastructure

The management infrastructure consists of two PowerEdge R440 or R640 servers that form a management cluster. Management components are virtualized to provide high availability. Redundant 10/25 Gb Ethernet uplinks to the network infrastructure, redundant 16Gbps Fibre channel uplinks to the storage array combined with vSphere High Availability ensure that management components stay online. A Dell EMC Networking S3048 switch is used for OOB connectivity. iDRAC ports in each management and compute cluster connect to this switch.

The management software components include:

- VMware vCenter Server Appliance (VCSA)
- Dell EMC OpenManage Integration for VMware vCenter (OMIVV)
- Dell EMC Virtual Storage Integrator (VSI)
- Dell EMC Avamar Virtual Edition (AVE)
- Dell EMC Data Domain Virtual Edition (DD VE)




Figure 12. Management Infrastructure

### 3.5.2 Management Server

The management software components for the Dell EMC Ready Stack require a nominal amount of virtual server resources. There are no strict requirements that these components cannot reside on the compute

server cluster, alongside the compute workload. However, many customers would choose to run these software components on dedicated physical server resources so that there is less resource contention with the compute workloads. Table 5 describes the optional management cluster, consisting of two PowerEdge R440 or R640 servers with a minimal server configuration.

Table 5 Management Server Configuration

| Components                       | Details  |
|----------------------------------|--|
| Server Platform Models           | (2) Dell EMC PowerEdge R440  |
|                                  | (2) Dell EMC PowerEdge R640  |
| Processors                       | (1 or 2) Intel® Xeon® SP Bronze, Silver, Gold, Platinum Skylake processors |
| Memory (depends on server model) | (1 – 24) 2666MT/s RDIMMs/LRDIMMs   |
| Network Adapter                  | Mellanox CX4 LX 10/25Gb Ethernet Network Adapter                           |
| Host Bus Adapter                 | QLogic® 2692 Dual Port Fibre Channel Adapter                               |
| Boot Device                      | Internal Dual SD Module (IDSDM)  |
|                                  | PERC H730p RAID Controller   |
| Out-of-band Server Management    | iDRAC9 Enterprise  |
| Hypervisor                       | VMware ESXi 6.5  |

### 3.5.3 VMware vCenter Server Appliance

The vCenter Server Appliance 6.5 is a preconfigured Linux virtual machine, which is optimized for running VMware vCenter Server and its associated services. The appliance package contains the following software:

- Project Photon OS 1.0
- PostgreSQL database
- vCenter Server 6.5 and vCenter Server 6.5 components
- Platform Services Controller that contains all the necessary services for running vCenter Server such as vCenter Single Sign-On, License service, and VMware Certificate Authority

VMware vCenter Server provides management of virtualized hosts and virtual machines from a single console. It gives administrators visibility into the configuration of the critical components of a virtual infrastructure. Key capabilities enabled by vCenter Server include VMware vSphere vMotion, VMware vSphere Distributed Resource Scheduler, VMware vSphere High Availability (HA) and VMware vSphere Fault Tolerance. The open plug-in architecture of vCenter Server supports a broad range of additional capabilities from VMware and its partners. The vCenter Server APIs also allow integration of physical and virtual management tools for maximum flexibility.

### 3.5.4 Dell EMC Virtual Storage Integrator

The Dell EMC Virtual Storage Integrator (VSI) for VMware vSphere Web Client is a plug-in for VMware vCenter. It enables administrators to view, manage, and optimize storage for VMware ESX/ESXi servers and hosts, and then map that storage to the hosts. VSI consists of a graphical user interface and the EMC Solutions Integration Service (SIS), which provides communication and access to the storage systems. Using the Solutions Integration Service, a storage administrator can enable virtual machine administrators to perform management tasks on a set of storage pools. For the Dell Unity platforms, tasks that you can perform with VSI include: Datastore provisioning, Space Reclamation and setting multi-pathing policies.

During installation, VSI automatically registers EMC PowerPath/VE and enables you to set multipathing policies for all devices using the VMware Native Multipathing Plug-in (NMP) or EMC PowerPath/VE. After installation, you can provision VMFS-6 datastores in vCenter Inventory Lists in datacenters, folders, clusters, or hosts. Additionally, using the space reclamation feature in VSI, you can reclaim unused storage on datastores, hosts, clusters, folders, and storage folders. Space Reclamation tasks are created per-datastore. Scheduling from the host, cluster, folder, or storage folder level results in one scheduled task per datastore under the target object. Only one scheduled reclaim task can exist for each datastore.

### 3.5.5 Dell EMC OpenManage Integration for VMware vCenter

The OpenManage Integration for VMware vCenter (OMIVV) is designed to streamline the management processes in your data center environment by allowing you to use VMware vCenter to manage your entire server infrastructure - both physical and virtual. From monitoring system level information, bubbling up system alerts for action in vCenter, rolling out cluster level BIOS and firmware updates for an ESXi cluster, to bare metal deployment, the OpenManage Integration will expand and enrich your data center management experience with Dell EMC PowerEdge servers. OpenManage Integration provides deep level details for inventory, monitoring, and alerting of Dell hosts within vCenter and recommends or performs vCenter actions based on Dell hardware events.

Proactive HA is a vCenter feature that works with OMIVV. When you enable Proactive HA, the feature safeguards your workloads by proactively taking measures based on degradation of redundancy health of supported components in a host. When OMIVV detects a change in the redundancy health status of supported components (either through Traps or polling), the health update notification for the component is sent to the vCenter server. Polling runs every hour, and it is available as a fail-safe mechanism to cover the possibility of a Trap loss. After assessing the redundancy health status of the supported host components, the OMIVV appliance updates the health status change to the vCenter server.

Deployment templates contain a system profile, hardware profile, hypervisor profile, a combination of system profile and hypervisor profile, or a combination of the hardware profile and hypervisor profile. The OMIVV Deployment Wizard uses this template to provision server hardware and deploy hosts within vCenter. It is recommended to use the system profile type for 14th generation servers. For Internal Dual SD Module deployment, the IDSDM should be enabled from BIOS before you deploy a hypervisor with OMIVV.

Figure 13 below highlights how OMIVV can create system profiles within vCenter.




Figure 13. Dell EMC OpenManage Integration for VMware vCenter

### 3.5.6 Dell EMC Avamar Virtual Edition

Avamar Virtual Edition (AVE) provides powerful data protection, unified management and hypervisor integration. The hypervisor integration allows administrators to take advantage of self-service data protection while using the native vSphere hypervisor interface. AVE uses Avamar data protection technology to protect VMware virtual environments. AVE brings performance backup features such as: Change Block Tracking (CBT) for backup and recovery, High-speed image-level backup and recovery for bare metal protection of each VM, and Universal proxy load balancing.




Figure 14. Avamar Administrator GUI

### 3.5.7 Dell EMC Data Domain

The Dell EMC Data Domain portfolio consists of an array of software features and backup appliances that transform backup, archive and disaster recovery with protection storage for Small, Midsize and Large Enterprise environments. Data Domain includes features such as Data Domain Boost, Data Domain Encryption, Data Domain Replicator, and Data Domain Management Center. With these advanced options, organizations can benefit from advanced integration with backup and enterprise applications, simple and cost-effective tiering to the public, private, or hybrid cloud for long-term retention, and network-efficient replication. Avamar backup to Data Domain brings value through instant access to a virtual machine by booting the VM directly from Data Domain via an NFS datastore. With the instant access feature, there is no restore operation required.

Table 6 briefly compares the Data Domain portfolio.

Table 6 Data Domain Portfolio Comparison

| Compare Data Domain          | DD VE 16TB | DD VE 96TB | DD3300    | DD6300    | DD6800   | DD9300   | DD9800   |
|------------------------------|------------|------------|-----------|-----------|----------|----------|----------|
| Max Throughput               | 2.1 TB/hr  | 4.0 TB/hr  | 4.2 TB/hr | 8.5 TB/hr | 14 TB/hr | 20 TB/hr | 31 TB/hr |
| Max Throughput with DD Boost | 5.6 TB/hr  | 11.2 TB/hr | 7.0 TB/hr | 24 TB/hr  | 32 TB/hr | 41 TB/hr | 68 TB/hr |

| Compare Data Domain                 | DD VE 16TB | DD VE 96TB | DD3300  | DD6300    | DD6800    | DD9300     | DD9800      |
|-------------------------------------|------------|------------|---------|-----------|-----------|------------|-------------|
| Usable Capacity                     | 16 TB      | 96 TB      | 4-32 TB | 36-178 TB | 56-288 TB | 144-720 TB | 200 TB-1 PB |
| Logical Capacity with Deduplication | 800 TB     | 4.8 PB     | 1.6 PB  | 8.9 PB    | 14.4 PB   | 36 PB      | 50 PB       |

Data Domain Virtual Edition (DD VE) is a software-only protection storage appliance - a virtual deduplication appliance that provides data protection for entry, enterprise and service provider environments. Like any Data Domain system, DD VE is always paired with backup software. DD VE runs the Data Domain Operating System (DD OS), and provides the DD OS command line interface (CLI) and the Data Domain System Manager graphical user interface (GUI) for performing all system operations. Data Domain Virtual Edition maintains the core Data Domain features that differentiate it as the industry-leading protection storage. This includes high-speed, variable length deduplication for a 10 - 30x reduction in storage requirements, best in class data integrity to ensure reliable recovery, and seamless integration with leading backup and archiving applications. DD VE also comes with DD Boost, DD Encryption for enhanced security of data, and DD Replicator, which enables network efficient replication for faster time-to-DR readiness.

Data Domain Virtual Edition is a software only solution. It requires that datastore storage has been provisioned to the ESXi host(s) that are running the DD VE instance. Data Domain provides a deployment assessment tool (DAT) which runs on DD VE. It measures the underlying I/O performance and determines the size of the file system. DAT may be used to scan the available physical storage to determine if the storage meets the DD VE requirements.

The Data Domain Boost feature provides integration between Avamar backup and Data Domain systems. With DD Boost, parts of the deduplication process are distributed to the Avamar backup server enabling client-side deduplication so only unique data segments are sent to the Data Domain system. This enables 50% faster backups and reduces network bandwidth requirements by 80 to 98%. DD Boost provides advanced load balancing and failover, which further improves throughput and resiliency.

### 3.5.8 Management Software Resources

The management software components run on virtual machines that reside in the management cluster. Table 7 lists the management components in the bundle and the recommended VM sizing of those components:

Table 7 Management Components Sizing

| Component  | VMs | CPU Cores | RAM (GB) | OS (GB) | NIC |
|--|-----|-----------|----------|---------|-----|
| VMware vCenter Server Appliance                    | 1   | 4         | 16       | 290     | 1   |
| Dell EMC OpenManage Integration for VMware vCenter | 1   | 2         | 8        | 44      | 1   |
| Dell EMC Virtual Storage Integrator                | 1   | 2         | 8        | 11      | 1   |

| Component                            | VMs | CPU Cores | RAM (GB) | OS (GB) | NIC |
|--------------------------------------|-----|-----------|----------|---------|-----|
| Dell EMC Data Domain Virtual Edition | 1   | 8         | 64       | 260     | 1   |
| Dell EMC Avamar Virtual Edition      | 1   | 2         | 16       | 3150    | 1   |
| Dell EMC Avamar Proxy                | 1   | 4         | 4        | 21      | 1   |

## 4 Component Specifications

This section details the hardware and software components used in the solution, which include VMware vSphere 6.5 on the PowerEdge rack server architecture with Dell Networking S5048-ON switches and Dell EMC Unity x50F All-Flash Storage.

The tables below provide a list of components needed to build a virtualized infrastructure for the Dell EMC Ready Stack as described in Solution Architecture. Optional items are indicated as such.

Table 8 Enterprise Small Component Specification

| Component            | Details  |  |
|----------------------|--|--|
| Virtualization Hosts | (4) Dell EMC PowerEdge R440 Server                         |  |
|                      | Processor  | (2) Intel Xeon SP-6152 Family            |
|                      | Memory   | 512 GB: (16) 32 GB DDR4 RDIMMs           |
|                      | Storage Adapter  | PERC H730p                               |
|                      | Local Storage  | (2) 120 GB RI SSD                        |
|                      | OS Volume  | (2) 16 GB microSD in Internal SD Module  |
|                      | Network  | Mellanox CX4 LX Dual Port PCIe NIC       |
|                      | HBA  | QLogic 2692 Dual Port FC PCIe Card       |
|                      | OS   | VMware ESXi 6.5 U1                       |
|                      | OS License   | (2) vSphere Enterprise Plus 1CPU, 3Years |
| Storage Array        | Dell EMC Unity 350F All-Flash<br>(17) 3.84TB RI SSD drives |  |
| Networking           | (2) Dell EMC Networking S5048-ON Ethernet Switches (LAN)   |  |
|                      | (2) Connectrix DS6505B Fibre Channel Switches (SAN)        |  |
| Licenses             | (4) iDRAC9 Enterprise                                      |  |
|                      | (8) vSphere Enterprise Plus 1CPU, 3Years                   |  |
|                      | VMware vCenter Standard License, 3Year                     |  |

| Component          | Details   |
|--------------------|---|
|                    | (4) OpenManage Integration for VMware vCenter, 3Year          |
| Cable Requirements | (8) 25 GbE SFP28 SR DAC Cables                                |
|                    | (12) LC-LC FC Cables  |
|                    | (12) Dell Networking Brocade Transceiver 16Gb SWL SFP         |
|                    | (2) 100 GbE QSFP28 Cables                                     |
|                    | (8) 1000Base-T SFP Transceiver (for Out-of-Band connectivity) |
|                    | (9) Cat5e Ethernet Network Patch Cables                       |

Table 9 Enterprise Medium Component Specification

| Component            | Details                            |   |
|----------------------|------------------------------------|---|
| Virtualization Hosts | (7) Dell EMC PowerEdge R740 Server |   |
|                      | Processor                          | (2) Intel Xeon SP-6152 Family           |
|                      | Memory                             | 512 GB: (16) 32 GB DDR4 RDIMMs          |
|                      | Storage Adapter                    | PERC H730p                              |
|                      | Local Storage                      | (2) 120 GB RI SSD                       |
|                      | OS Volume                          | (2) 16 GB microSD in Internal SD Module |
|                      | Network                            | Mellanox CX4 LX Dual Port NDC           |
|                      | HBA                                | QLogic 2692 Dual Port FC PCIe Card      |
|                      | OS                                 | VMware ESXi 6.5 U1                      |
| Management Hosts     | (2) Dell EMC PowerEdge R640 Server |   |
|                      | Processor                          | Intel Xeon SP-6152 Family               |

| Component          | Details  |   |
|--------------------|--|---|
|                    | Memory   | 512 GB: (16) 32 GB DDR4 RDIMMs          |
|                    | Storage Adapter  | PERC H730p                              |
|                    | Local Storage  | (2) 120 GB RI SSD                       |
|                    | OS Volume  | (2) 16 GB microSD in Internal SD Module |
|                    | Network  | Mellanox CX4 LX Dual Port NDC           |
|                    | HBA  | QLogic 2692 Dual Port FC PCIe Card      |
|                    | OS   | VMware ESXi 6.5 U1                      |
| Storage Arrays     | Dell EMC Unity 450F All-Flash<br>(25) 3.84TB RI SSD drives |   |
| Networking         | (2) Dell EMC Networking S5048-ON Ethernet Switches (LAN)   |   |
|                    | (2) Connectrix DS6505B Fibre Channel Switches (SAN)        |   |
|                    | Dell EMC Networking S3048-ON Ethernet Switches (OOB)       |   |
| Licenses           | (9) iDRAC9 Enterprise                                      |   |
|                    | (16) vSphere Enterprise Plus 1CPU, 3Years                  |   |
|                    | VMware vCenter Standard License, 3Year                     |   |
|                    | (9) OpenManage Integration for VMware vCenter, 3Year       |   |
| Cable Requirements | (20) 25 GbE SFP28 SR DAC Cables                            |   |
|                    | (22) LC-LC FC Cables                                       |   |
|                    | (22) Dell Networking Brocade Transceiver 16Gb SWL SFP      |   |
|                    | (2) 100 GbE QSFP28 Cables                                  |   |
|                    | (14) Cat5e Ethernet Network Patch Cables                   |   |



Table 10 Enterprise Large Component Specification


| Component            | Details  |   |
|----------------------|--|---|
| Virtualization Hosts | (13) Dell EMC PowerEdge R640 Server                        |   |
|                      | Processor  | (2) Intel Xeon SP-6152 Family           |
|                      | Memory   | 512 GB: (16) 32 GB DDR4 RDIMMs          |
|                      | Storage Adapter  | PERC H730p                              |
|                      | Local Storage  | (2) 120 GB RI SSD                       |
|                      | OS Volume  | (2) 16 GB microSD in Internal SD Module |
|                      | Network  | Mellanox CX4 LX Dual Port NDC           |
|                      | HBA  | QLogic 2692 Dual Port FC PCIe Card      |
|                      | OS   | VMware ESXi 6.5 U1                      |
| Management Hosts     | (2) Dell EMC PowerEdge R640 Server                         |   |
|                      | Processor  | (2) Intel Xeon SP-6152 Family           |
|                      | Memory   | 512 GB: (16) 32 GB DDR4 RDIMMs          |
|                      | Storage Adapter  | PERC H730p                              |
|                      | Local Storage  | (2) 120 GB RI SSD                       |
|                      | OS Volume  | (2) 16 GB microSD in Internal SD Module |
|                      | Network  | Mellanox CX4 LX Dual Port PCIe NIC      |
|                      | HBA  | QLogic 2692 Dual Port FC PCIe Card      |
|                      | OS   | VMware ESXi 6 U1                        |
| Storage Arrays       | Dell EMC Unity 550F All-Flash<br>(25) 7.68TB RI SSD drives |   |

| Component          | Details  |
|--------------------|--|
| Networking         | (2) Dell EMC Networking S5048-ON Ethernet Switches (LAN) |
|                    | (2) Connectrix DS6510B Fibre Channel Switches (SAN)      |
|                    | Dell EMC Networking S3048-ON Ethernet Switches (OOB)     |
| Licenses           | (15) iDRAC9 Enterprise                                   |
|                    | (28) vSphere Enterprise Plus 1CPU, 3Years                |
|                    | VMware vCenter Standard License, 3Year                   |
|                    | (15) OpenManage Integration for VMware vCenter, 3Year    |
| Cable Requirements | (32) 25 GbE SFP28 SR DAC Cables                          |
|                    | (34) LC-LC FC Cables                                     |
|                    | (34) Dell Networking Brocade Transceiver 16Gb SWL SFP    |
|                    | (2) 100 GbE QSFP28 Cables                                |
|                    | (20) Cat5e Ethernet Network Patch Cables                 |

## 5 Scaling Dell EMC Ready Stack


The solution can be scaled by adding multiple compute nodes (pods) in the customer data center. The Dell EMC Networking Z9100 switch can be used to create a simple yet scalable network. The Z9100 switches serve as the spine switches in the leaf-spine architecture. The Z9100 is a multiline rate switch supporting 10/25/40/50/100 Gb Ethernet connectivity and can aggregate multiple racks with little or no oversubscription.

When connecting multiple racks, using the 40/100 GbE uplinks from the rack, you can build a large fabric that supports multi-terabit clusters. The density of the Z9100 allows flattening the network tiers and creating an equal-cost fabric from any point to any other point in the network.



**Figure 15. Multiple Compute Pods Scaled Out Using Leaf Spine Architecture**

For large domain layer-2 requirements the Extended Virtual Link Trunking (eVLT) can be used on the Z9100, as shown in the following figure. The VLT pair formed can scale in terms of hundreds of servers inside multiple racks. Each rack has four 40/100 GbE links to the core network providing enough bandwidth for all the traffic between each rack.




**Figure 16. Multiple Compute PODs scaled out using eVLT**

## 6 Dell EMC Ready Stack Validation

This section provides a high-level summary of the Ready Stack validation process. The system validation requirements and methodology ensure resiliency throughout the design, protection from hardware failures, and assurance that the system will perform well under load.

### 6.1 Validation Strategy

The validation strategy for Dell EMC Ready Stack is driven by the following requirements:




**Figure 17. Dell EMC Ready Stack Validation Requirements**

- **Design principles:** The design principles provide a guideline for what needs to be validated. Each aspect of the design principle (such as high availability) must be validated.
- **Flexible components:** Dell EMC Ready Stack supports a variety of rack servers and storage models.
- **Interoperability:** Dell EMC Ready Stack supports various software and hardware components. The interoperability between these components must be validated.
- **System Deployment:** Deploying Dell EMC Ready Stack from bare metal to a virtual infrastructure ready to provision customer workloads must be validated.

## 6.2 Validation Methodology Overview

This section provides the various aspects of the validation methodology used by Dell EMC Engineering for validation of the Dell EMC Ready Stack.



**Figure 18. Dell EMC Ready Stack Validation Methodology**

### 6.2.1 Leveraging existing certification programs

The first step in validation is to select only the components that have been certified by VMware IOVP. This ensures the most trouble-free interoperability while minimizing any potential issues. The I/O Vendor Partner package (IOVP) offers members a comprehensive set of tools and resources needed to develop, certify and release I/O device drivers in the following categories: Network (NIC) device drivers, Storage (HBA) device drivers and RDMA (Remote Direct Memory Access) device drivers. Partners initiating device driver development use the NDDK (Native Device Driver Kit) to develop native I/O drivers, which is the preferred driver model in terms of performance and support. All the Dell EMC Ready Stack components that provide support for VMware products are compliant with the VMware IOVP certification program.

### 6.2.2 Versions Matrix

The second step is to determine the correct versions of all supported components. Versions matrix includes BIOS and firmware versions of all validated hardware components, hypervisor version, plugin versions and versions of all software components including drivers. This matrix is built by consulting each individual product support and interoperability matrix.

Table 11 below represents the firmware and software versions that were validated as part of Dell EMC Ready Stack.

Table 11 Dell EMC Ready Stack Solution Validated Hardware and Software

| Layer    | Device   | Version(s)                                |
|----------|--|---|
| Server   | PowerEdge R640/R740/740xd                          | BIOS 1.3.7<br>iDRAC 3.15.17.15            |
|          | Mellanox CX4 LX Dual-port rNDC                     | Firmware 14.20.18.20, mlx5-core 4.16.10.3 |
|          | Qlogic 2692 Dual-port Fibre HBA                    | Firmware 14.02.13, qlnativefc 2.1.57.0-1  |
| Network  | Dell EMC Networking S3048 OS 9                     | FTOS 9.13.0.1P1                           |
|          | Dell EMC Networking S5048 OS 9                     | FTOS 9.12.1.0                             |
| Storage  | Dell EMC Unity x50F                                | 4.2.2                                     |
|          | Connectrix DS6500 Fabric OS                        | 8.1.2a                                    |
| Software | VMware vSphere ESXi                                | 6.5.0 U1, Build 7388607                   |
|          | VMware vCenter Server Appliance                    | 6.5.0 U1, Build 7312210                   |
|          | Dell EMC Virtual Storage Integrator                | 7.3.2                                     |
|          | Dell EMC OpenManage Integration for VMware vCenter | 4.1                                       |
|          | Dell EMC Avamar Virtual Edition                    | 7.5.1                                     |
|          | Dell EMC Data Domain Virtual Edition               | 6.1.1.5                                   |

### 6.2.3 Deployment Validation

The next step in validation is to deploy the system from bare metal. Best practices of each individual components, as well as the Dell EMC Ready Stack's design guidelines are followed in deploying the system. The Dell EMC Ready Stack Deployment Guide provides instructions on how the system should be deployed by Dell services at a customer site.

### 6.2.4 Simulate Workload

Once Dell EMC Ready Stack is deployed, virtual machines are installed to simulate a workload. These virtual machines are used to validate the systems. The objective is to generate resource consuming

workloads that will aid in verifying the soundness of the deployment. Deploying one virtual machine per physical core in the compute cluster is a good starting point for the workload simulation. The following tools are used within each virtual machine to simulate workloads:

- **IOmeter:** IOmeter is an I/O subsystem measurement and characterization tool for single and clustered systems. It is easily configured to replicate the behavior of many popular applications. For workload simulation purposes, IOmeter is used in situations where VMs need to be scaled up in a cluster without stopping the workload to see what the cluster can handle.
- **TestLimit:** TestLimit is a console based program which can allocate specified amounts of system resources such as RAM, GDI handles, large pages and more to stress test Windows based computers and applications. For instance, TestLimit can be used to allocate all but 1GB of RAM, to find out how applications work in low memory situations. TestLimit can also be used to simulate a memory leak, allocating a certain number of megabytes every few seconds.
- **Network file copy:** To simulate network file copies, a script is used to generate a large number of files that vary in size from 100KB to 80MB at a network location. The files are then copied to the VMs that are running workloads to stress them further as well as simulate real world environments better.

## 6.2.5 Design Validation

Once we have the workload running, we validate the system by testing each design aspect. Here is a list of some of the test cases that are used to validate the system. While this is not a comprehensive list, it provides a general idea of what is being validated.

### 6.2.5.1 Host Layer

- **Validation of VMware High Availability (host failover):** The objective is to ensure VMware HA is configured correctly and VMs can recover upon host failures. The host failure can be simulated by shutting down the host using the iDRAC. VMware High Availability functionality is verified by ensuring that the virtual machines on the failed host restart on other hosts in the cluster.
- **Validation of VMware Distributed Resource Scheduler:** The objective is to ensure that the virtual machines are evenly distributed across the hosts in the cluster and each host's resource consumption (CPU and memory) is relatively even across the cluster. A new host is added to the cluster, and it is verified that the DRS migrates some of the VMs to the newly added host for an even resource consumption among the hosts in the cluster.

### 6.2.5.2 Network Layer

- **Validating Network Load Balancing and Multipathing:** Ensure that all the network paths are utilized when a load is running. VMware vCenter can be used to monitor the resource usage of both the physical network interfaces and virtual network interfaces. The following are the verification steps for various traffic types.
- **Management Network:** The ESXi management network is assigned to two *vmnic* uplinks for active/active redundancy. Management network traffic can use any use any network port (governed by NIOC). Simulating a network failure, all Management network traffic should fail to the remaining active port. Management traffic fail over can be verified using VMware vCenter.
- **vMotion Network:** The VMware vMotion network is also assigned to two *vmnic* uplinks for active/active redundancy. While VMware vMotion is in progress, determine the network port in use. Simulate a network failure while vMotion is progressing, and ensure that the traffic fails over to the remaining port and vMotion is completed without disruption.
- **Virtual Machine Network:** Ensure that all available network paths are used if there are multiple VMs in a server, with each generating sufficient network traffic. Simulate the failure of an active port and ensure that the VM traffic is uninterrupted.
- **TOR Switch:** Simulate the failure of the S5048 TOR switch by rebooting the switch. Ensure that all network traffic fails over to the other TOR switch. The network traffic includes management,

vMotion, virtual machine and vSAN (if applicable). Ensure that the traffic fails back when the switch is online.

### 6.2.5.3 Storage Layer

Fibre Channel based systems are configured with two fault domains. The following steps will ensure that load balancing and high availability is configured correctly between the two fault domains.

- **Fibre Channel Network:** The Fibre Channel based system is configured with dual port Fibre Channel HBAs (*vmhba*). Ensure that multipathing is configured correctly by verifying that both interfaces are utilized in vCenter. Simulate the failure of one storage port and ensure that the traffic fails over.
- **FC SAN Switch:** Simulate the failure of the Connectrix DS6500 switch by rebooting the switch. Ensure that all the storage traffic fails over to the other switch. Ensure that the traffic fails back when the switch is online.
- **FC Storage Controller:** Simulate the failure of a storage controller by removing one of the Unity x50F storage controllers. Ensure that all storage traffic fails over to the other controller. Ensure that the traffic fails back when the controller is back online.

### 6.2.5.4 Stress Testing

The Dell EMC Ready Stack solution infrastructure is subjected to stress testing to ensure that the system does not fail under heavy workload. This includes significantly stressing the CPU and storage resources. The objectives are to identify any bottlenecks and to make sure that the system does not falter under heavy workload. For a typical stress test, The tools mentioned in section 6.2.5.2 are run for several days. Host and cluster performance and logs are collected and reviewed to ensure the system is running as designed.

## 7 Solution Summary

In summary, Dell EMC Ready Stack provides best-in-class converged infrastructure performance in a flexible, scalable, and affordable solution. This Design Guide has provided both high-level and in-depth solution design information, for a variety of stakeholders, to guide you through the design philosophy and architectural decisions made for the Dell EMC Ready Stack.

For detailed deployment implementation please check the Dell EMC Ready Stack Deployment Guide.

## 8 References

- [Dell EMC PowerEdge R640 Installation and Service Manual](#)
- [Integrated Dell Remote Access Controller 9](#)
- [iDRAC9 Systems Management – Wiki](#)
- [Dell EMC Unity Family Installation Guide](#)
- [Dell EMC Unity: Best Practices Guide](#)
- [Dell EMC Unity Unisphere CLI User Guide](#)
- [VMware vSphere Documentation](#)
- [Dell Configuration Guide for the S3048–ON System](#)
- [Dell EMC Networking OS Configuration Guide for the S5048F–ON System](#)
- [Brocade Fabric OS Web Tools Administrator's Guide](#)
- [VSI for VMware vSphere Web Client Product Guide](#)
- [OpenManage Integration for VMware vCenter Web Client Installation Guide](#)
- [Avamar Virtual Edition for VMware System Installation Guide](#)
- [Data Domain Virtual Edition Installation and Administration Guide](#)