

DELL EMC READY BUNDLE FOR VIRTUALIZATION — WITH VMWARE VSAN INFRASTRUCTURE

Design Guide

APRIL 2017

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 © 2017 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 4/2017

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

EXECUTIVE SUMAMRY	4
INTRODUCTION	5
AUDIENCE AND SCOPE	5
SUPPORTED CONFIGURATION	5
Compute server.....	5
Management server	6
DESIGN PRINCIPLES	7
DESIGN OVERVIEW	8
COMPUTE DESIGN	8
NETWORK DESIGN	9
Dell Networking S4048-ON	9
Network connectivity of Dell EMC PowerEdge FX2 Servers	9
Network configuration of Dell PowerEdge rack servers.....	10
Network configuration for LAN traffic with VMware vSphere Distributed Switch (vDS)	10
STORAGE DESIGN	11
Boot device	11
Storage controller	11
Disk Groups.....	11
PowerEdge FC630 with PowerEdge FD332	12
PowerEdge FC430 with PowerEdge FD332	13
PowerEdge rack servers	13
Fault tolerance.....	13
Supported Disk Configurations.....	14
MANAGEMENT DESIGN	15
Management Infrastructure	15
Management Components	15
VMWARE VREALIZE AUTOMATION DESIGN	15
vRealize Automation components	16
VMware vRealize Automation deployment model sizing	16
SCALING THE READY BUNDLE	18
REFERENCES	19

Executive Sumamry

This Design Guide outlines the architecture of the Dell EMC Ready Bundle for Virtualization with VMware vSAN storage. A high-level overview of the steps required to deploy the components is also provided in this document.

The Dell EMC Ready Bundle for Virtualization is the most flexible Ready Bundle in the industry. This hyper-converged solution can be built to include Dell PowerEdge rack servers and PowerEdge FX2 converged platforms. Virtualization and storage are powered by VMware vSphere with vSAN to form a robust hyper-converged platform for enterprise workloads.

Introduction

The Dell EMC Ready Bundle for Virtualization with VMware is a family of converged and hyper-converged systems that combine servers, storage, networking, and infrastructure management into an integrated and optimized system that provides a platform for general-purpose virtualized workloads. The components undergo testing and validation to ensure full compatibility with VMware vSphere. The Ready Bundle uses infrastructure best practices, and converged server architecture. This document provides an overview of the design and implementation of the Dell EMC Ready Bundle for Virtualization with VMware vSAN storage.

Audience and scope

This document is intended for stakeholders in IT infrastructure and service delivery who have purchased or are considering purchase of Dell EMC Ready Bundle for Virtualization with VMware vSAN storage. This document provides information and diagrams to familiarize you with the components that make up the bundle.

Supported Configuration

The following table provides an overview of the supported components:

Table 1. Supported Configuration

Components	Details
Server platforms	Choice of: <ul style="list-style-type: none">• Dell EMC PowerEdge R630,• Dell EMC PowerEdge R730,• Dell EMC PowerEdge R730xd,• Dell EMC PowerEdge FX2 with Dell EMC PowerEdge FD332 and Dell EMC PowerEdge FC630 or Dell EMC PowerEdge FC430
LAN connectivity	(2) Dell Networking S4048-ON
Out-of-band (OOB) connectivity	(1) Dell Networking S3048-ON
Storage	vSAN
Management server platform	(3) Dell EMC PowerEdge R630
Management software components	<ul style="list-style-type: none">• VMware vCenter Server Appliance• VMware vRealize Automation• VMware vRealize Business• VMware vRealize Log Insight• VMware vRealize Operations Manager

Compute server

Ready Bundle for Virtualization offers Customers a choice of rack servers or modular enclosures for their compute and vSAN based software-defined storage infrastructure. The following table details the supported compute servers:

Table 2. Compute server configurations

Platform model	PowerEdge R630, PowerEdge R730, and PowerEdge R730xd	PowerEdge FX2 with PowerEdge FC630	PowerEdge FX2 with PowerEdge FC430
Processor	(2) Intel® Xeon® E5 v4 Broadwell processors	(2) Intel® Xeon® E5 v4 Broadwell processors	(2) Intel® Xeon® E5 v4 Broadwell processors
Memory	2400 MT/s RDIMMs	2400 MT/s RDIMMs	2400 MT/s RDIMMs
Network adapter	Intel® X710 quad port 10Gb network adapter or QLogic® 57800 10 Gb quad port network adapter	Intel® X710 quad port 10Gb network adapter or QLogic® 57800 10 Gb quad port network adapter	QLogic® 57800 10 Gb dual port network adapter
Storage controller	Dell PowerEdge HBA330	Dell PERC FD33xD RAID controller	Dell PERC FD33xD RAID controller
Supported disk configuration	Hybrid and All-Flash disk configurations	All-Flash disk configurations	All-Flash disk configurations
Boot device	SATADOM	Dell PERC H330 with local drives	On blade SATA controller with local drives
OOB management	iDRAC8 Enterprise	iDRAC8 Enterprise	iDRAC8 Enterprise
Hypervisor	VMware ESXi 6.5	VMware ESXi 6.5	VMware ESXi 6.5
FX2 chassis configuration	Not applicable	CMC and 2x Dell Networking FN IOAs	CMC and 2x Dell Networking FN IOAs

Management server

The management cluster consist of three PowerEdge R630 servers will the following configuration:

Table 3. Management server configuration

Components	Details
Number of servers	(3) Dell EMC PowerEdge R630
Processor	(2) Intel® Xeon® E5 v4 Broadwell processors
Memory	2400MT/s RDIMMs
Network adapter	Intel® X710 quad port 10Gb network adapter or QLogic® 57800 10 Gb quad port network adapter
Storage controller	PowerEdge HBA330
Disk configuration	Hybrid
Boot device	SATADOM
OOB management	iDRAC8 Enterprise

Design Principles

The following principles are central to the design and architecture of the Dell EMC Ready Bundle for Virtualization. The Ready Bundle for Virtualization is built and validated using these design principles.

- No single point-of-failure: Redundancy is incorporated in the critical aspects¹ of the solution, including server high availability features, networking, and storage.
- Integrated Management: Provide integrated management using vRealize Operations Manager with associated plugins.
- Designed for Hyper Converged Infrastructure (HCI): Hyper converged nodes are designed to meet unique requirements of HCI such as boot devices, cache and capacity drives, storage controller and fault domains.
- Hardware configuration for virtualization: The system is designed for general use case virtualization. Each server is equipped with appropriate processor, memory, host bus, and 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.
- vRealize Suite Enabled: The system supports VMware vRealize Suite for managing heterogeneous environments and hybrid cloud.
- Flexible configurations: Ready Bundle for Virtualization can be configured to suit most customer needs for a virtualized infrastructure. The solution supports options, such as configuring server model, server processors, server memory, type of storage (Fibre Channel, iSCSI and SDS), and network switches based on customer needs.

¹ Out of band management is not considered critical to user workload and does not have redundancy

Design Overview

This section provides an overview of the architecture including computer, storage, network, and management architecture. The following figure provides a high-level overview of the architecture, including compute servers (showing flexible compute nodes), management servers, software-defined storage, LAN switches, and out-of-band switches:

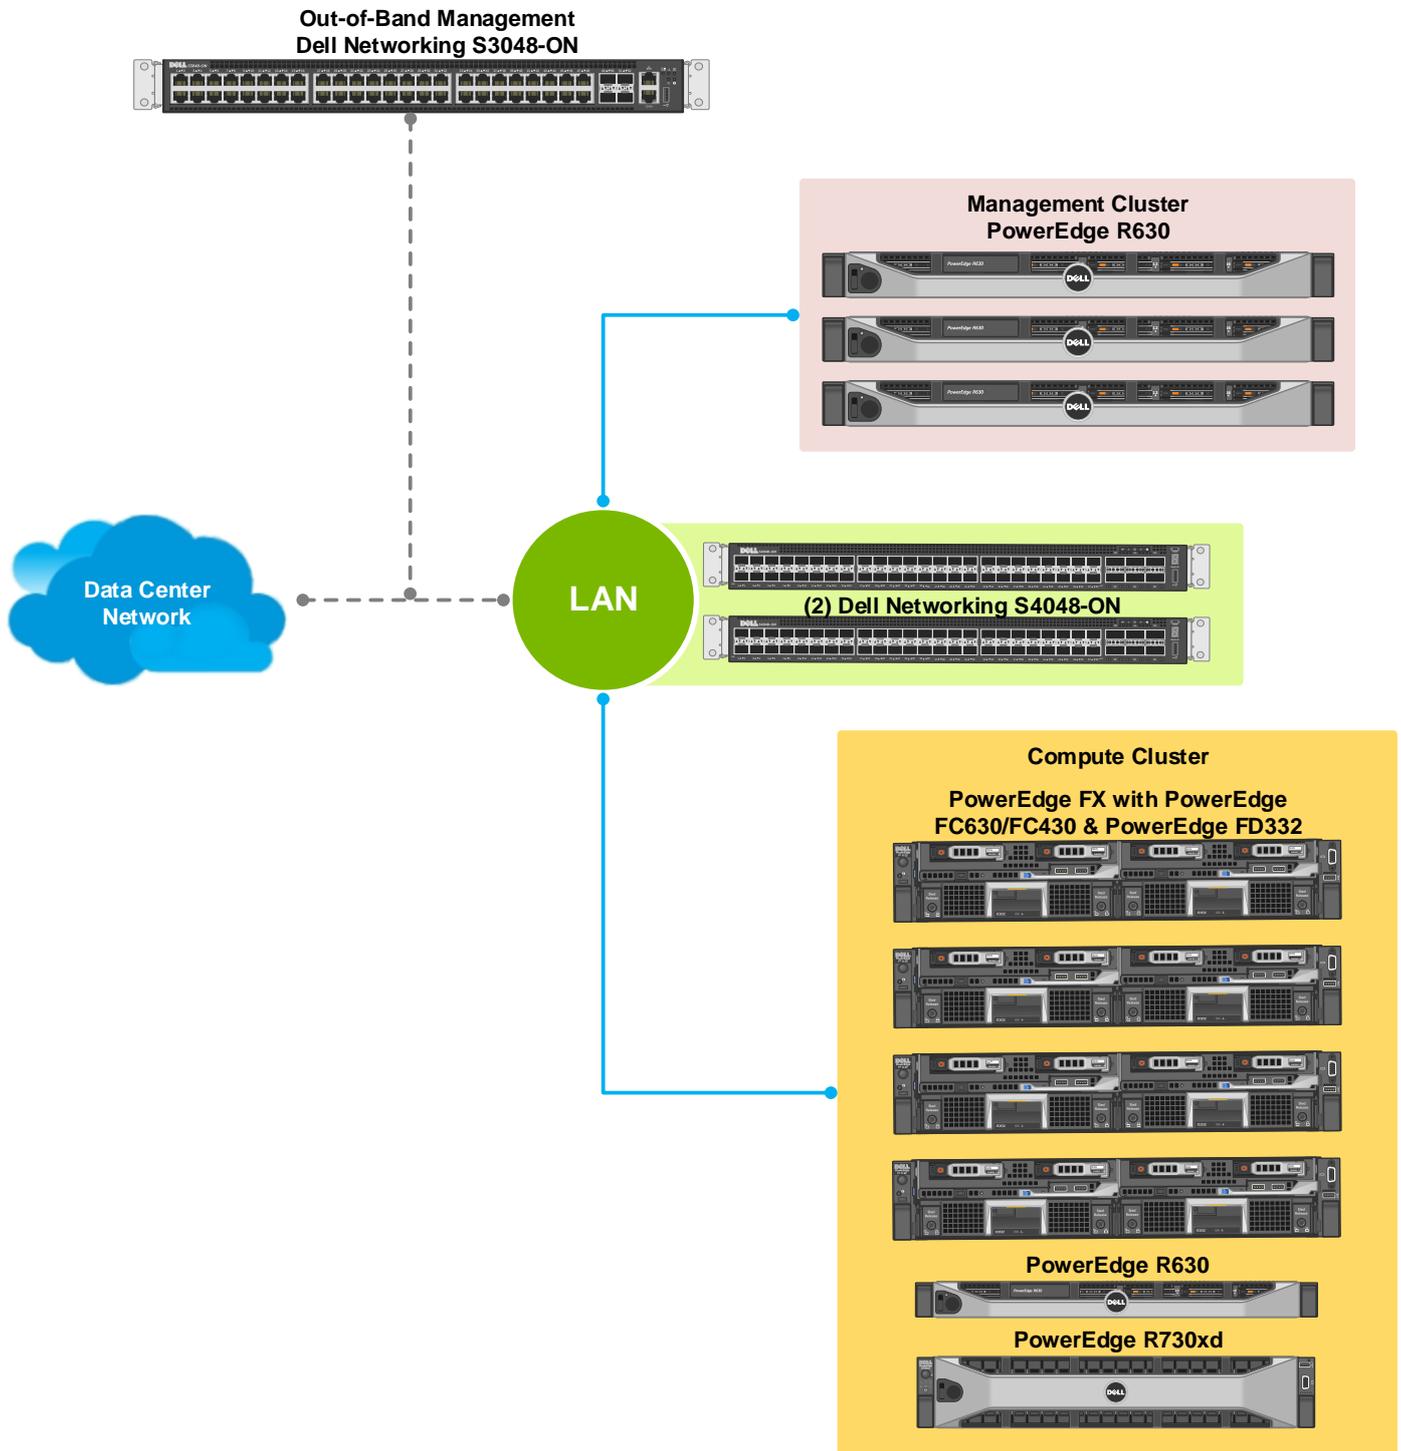


Figure 1. Ready Bundle for Virtualization with VMware vSAN – Design Overview

Compute Design

The latest Intel® Xeon v4 Broadwell generation processors power the Ready Bundle for Virtualization with VMware. With up to 22 cores per CPU and clock speeds of up to 3.5GHz in the Dell EMC PowerEdge rack servers you can reduce CPU socket-based licensing costs and achieve greater VM density. The Dell EMC Converged FX2 platform provides a powerful balance of resource density and power consumption in a dense form factor. These server nodes are configured with processors that have a lower thermal design power (TDP) value resulting in lower cooling and electrical costs.

PowerEdge rack 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, 2400 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 Dell fault-resilient mode are available.

The Intel® X710 and QLogic® network adapters are supported to provide 10 Gb network connectivity to the top of rack switches. PowerEdge servers support various BIOS configuration profiles that control the processor, memory, and other configuration options. Dell recommends the default profile of performance for the Dell EMC Ready Bundle for Virtualization with VMware.

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.

Dell Networking S4048-ON

The network architecture employs Virtual Link Trunking (VLT) connection between the two Top-of-Rack (ToR) switches. In a non-VLT environment, redundancy requires idle 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 Networking S4048-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

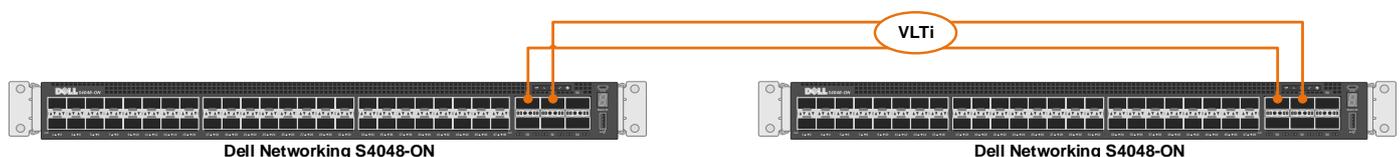


Figure 2. Dell Networking S4048-ON Virtual Link Trunk Interconnect (VLTi) Configuration

The Dell Networking S4048-ON switches each provide six 40 GbE uplink ports. The Virtual Link Trunk Interconnect (VLTi) configuration in this architecture uses two 40 GbE ports from each Top-of-Rack (ToR) switch to provide an 80 Gb data path between the switches. The remaining four 40Gb 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.

Network connectivity of Dell EMC PowerEdge FX2 Servers

This section describes the network connectivity when using PowerEdge FX2 modular infrastructure blade servers. The compute cluster Dell EMC PowerEdge FC630 and the Dell EMC PowerEdge FC430 server nodes connect to the Dell Networking S4048-ON switches through the Dell EMC PowerEdge FN410S modules shown in the following PowerEdge FX2 blade chassis figure:

- Connectivity between the PowerEdge FX2 server nodes and the Dell EMC PowerEdge FN410S: The internal architecture of Dell EMC PowerEdge FX2 chassis provides connectivity between an Intel® or QLogic® 10 GbE Network Daughter Card (NDC) in each PowerEdge FX2 server node and the internal ports of the Dell EMC PowerEdge FN410S module. The Dell EMC PowerEdge FN410S has eight 10 GbE internal ports. Each PowerEdge FX2 server can be configured with two Dell EMC PowerEdge FC630 and two Dell EMC PowerEdge FD332 storage blocks or four Dell EMC PowerEdge FC430 and two Dell EMC PowerEdge FD332 storage blocks. Dell EMC PowerEdge FC630 server nodes are configured with quad port 10 Gb network adapters that connect to the internal ports in each Dell EMC PowerEdge FN410S. The PowerEdge FC430 server nodes are configured with dual port 10 Gb network adapters that connect to the internal ports in each PowerEdge FN410S.

- Connectivity between the PowerEdge FN410S and S4048-ON switches: Two Dell EMC PowerEdge FN410S I/O Aggregators (IOA) in the architecture provides the Top-of-Rack (ToR) connectivity for the PowerEdge FX2 server nodes. Each IOA provides four external ports. Ports 9 and 10 from each FN IOA form a port-channel which in turn connects to the Top-of-Rack (ToR) switches.

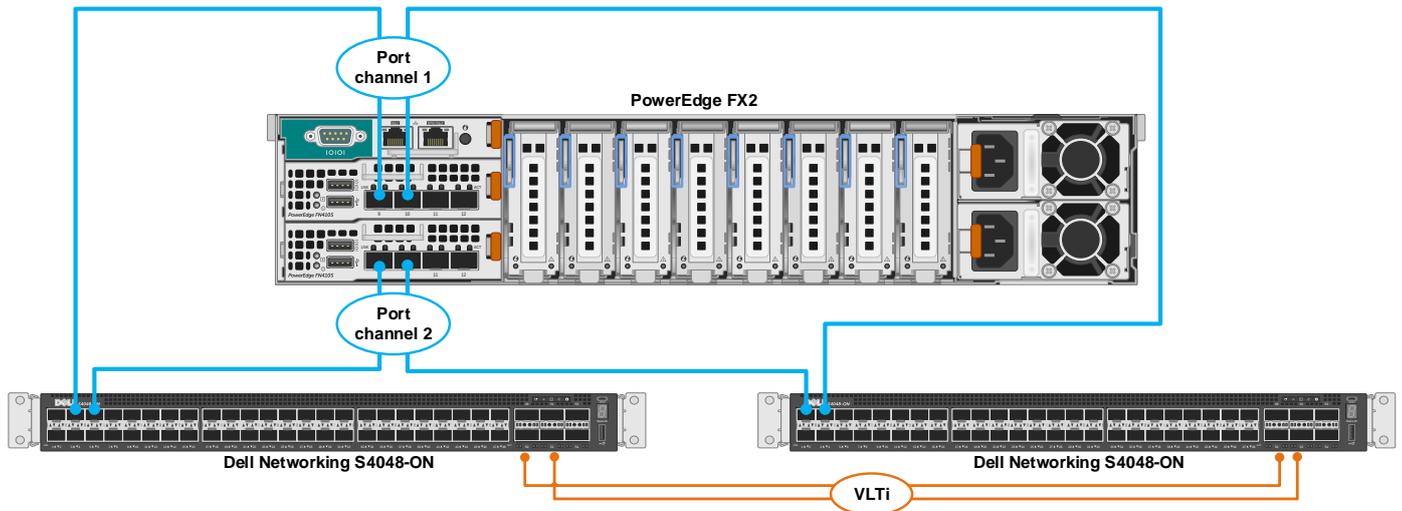


Figure 3. PowerEdge FX2 connectivity

Network configuration of Dell PowerEdge rack servers

The compute cluster can either be Dell EMC PowerEdge FX2 servers or Dell EMC PowerEdge rack servers. This section describes the network connectivity if rack servers are used for compute servers, as well as the management servers. The following image is an example of the connectivity between the compute and management PowerEdge rack servers and Dell Networking S4048-ON switches. The compute and management rack servers have two 10 GbE connections to each of the S4048-ON switches through one quad port 10 GbE Network Daughter Card (NDC).

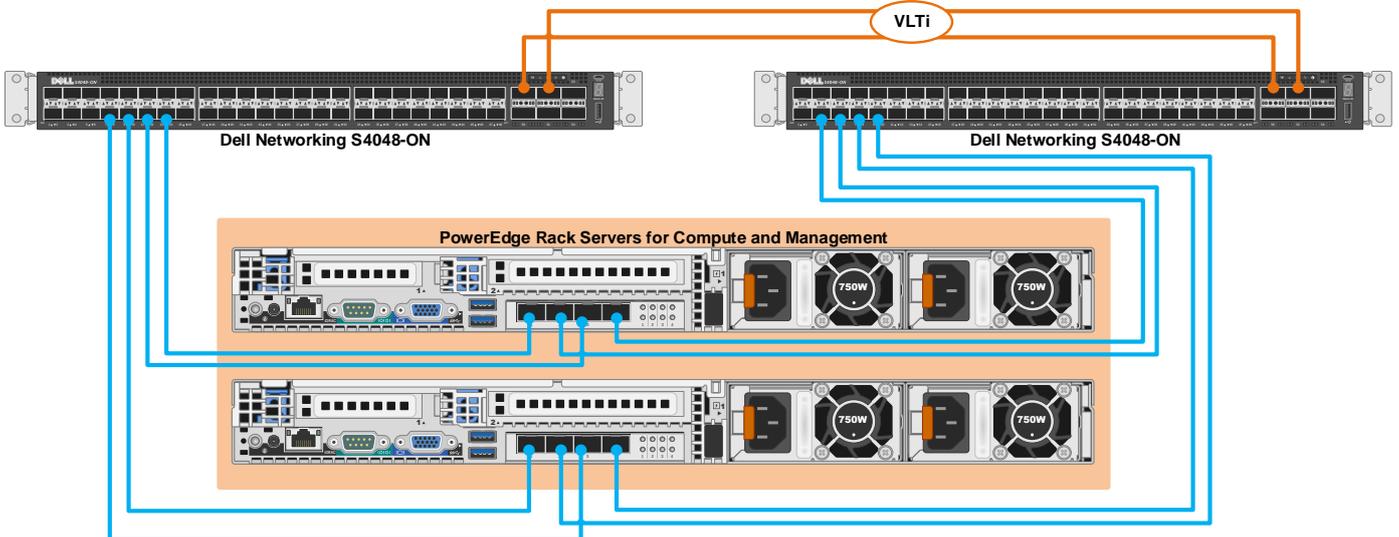


Figure 4. PowerEdge rack server connectivity

Network configuration for LAN traffic with VMware vSphere Distributed Switch (vDS)

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 figures show the virtual distributed switch configuration for a Dell EMC PowerEdge FC430 (with dual 10Gb Ethernet ports) and Dell EMC PowerEdge FC630/rack servers (with quad 10Gb Ethernet ports).

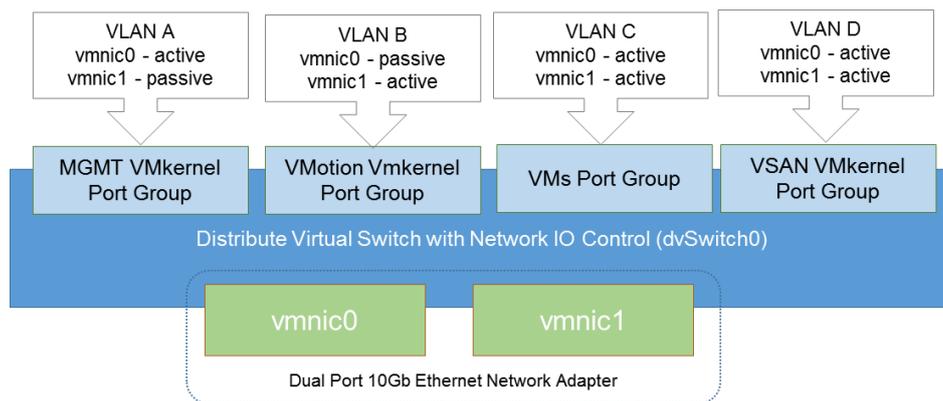


Figure 5. Distributed Virtual Switch for dual port configuration

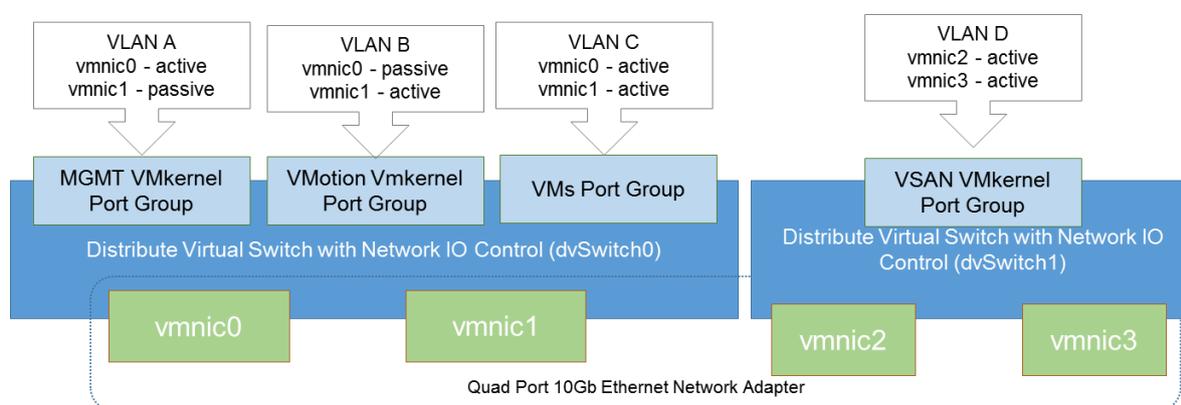


Figure 6. Distributed Virtual Switch for quad port configuration

Storage Design

Designing a hyper-converged infrastructure requires careful selection of cache/capacity drives, storage controller, and boot device to ensure optimal performance. The three management servers form a vSAN cluster while the compute nodes are part of a separate vSAN cluster. The following concepts and recommendations form a basis for component selection for both compute and management servers:

Boot device

VMware vSAN introduced two primary requirements for the boot device. First, the boot device cannot be connected to the same RAID controller or Host Bus Adapter (HBA) as the drives that vSAN consumes. Second, vSAN generates trace logs that need to be stored on a persistent storage device. IDSDM and USB boot options lack the endurance to be used as persistent storage, thus these do not support storing trace logs and necessitate the use of an NFS datastore or other persistent storage. For this reason Dell recommends SATADOM or local drives on a separate controller for the ESXi OS installation.

Storage controller

VMware vSAN supports RAID and HBA storage controllers. A RAID controller allows the drives to be configured as individual RAID 0 volumes or it can also operate in HBA mode. Hardware-based RAID controllers are not required for data protection. VMware vSAN provides object level redundancy to protect against drive failures. Dell recommends using a HBA to reduce complexity and cost.

NOTE: Local drives used for the ESXi installation or for creating a VMFS datastore not powered by vSAN should be connected to a separate storage controller.

Disk Groups

VMware vSAN is structured around the concept of logical pools of drives organized into functional management units called disk groups (DGs). These DGs must always follow certain design guidelines.

- Each DG always has a single cache drive, which is always a solid state drive (SSD).
- Each DG always contains one (1) to seven (7) capacity drives, which can either be all magnetic disks (hybrid mode) or all SSDs (all flash mode).
- The minimum number of DGs per host is one (1), and the maximum DGs per host is five (5).
- The cache drive serves as a read and write buffer for hybrid configurations, but only as a write buffer for all flash implementations.
- The cache drive size should meet or exceed 10% of the useable capacity in the DG.
- The performance and endurance requirements of the cache drive are greater than those of the capacity drives.
- VMware recommends that the configuration of each node should closely match that of the other nodes in the vSAN cluster. This is referred to as a “balanced configuration”.

The broad range of disk group configurations selected for Dell EMC Ready Bundle for Virtualization with VMware vSAN have been designed with variables such as cost, performance, availability, and capacity in mind. This enables the offering to cover a diverse spectrum of use cases.

Multiple disk groups per node generally offer improved performance and fault tolerance compared to single DG designs thanks to the additional high performance cache drive targets and multiple DG fault domains. It is important to note that because the performance and endurance requirements are higher for cache drives, these devices are typically more expensive. This can add up quickly when spread across many nodes in a clustered solution such as this, thus creating the need for a range of configurations for different budgets and use cases.

Multiple disk groups can also enable additional capacity, but only if an abundance of drive bays are available. If drive bays are limited, multiple DGs will actually reduce capacity because two or more drive slots are being dedicated to cache. Consider a server configuration like the AF-6 FC430 in Table 4 below, which has eight (8) drive bays available to vSAN. The maximum capacity configuration would be one (1) cache drive plus seven (7) capacity drives. To implement dual disk groups would involve two (cache) drives and only leave six (6) drive slots remaining for capacity.

All of these factors are considered in the design of the disk group configurations for Dell EMC Ready Bundle for Virtualization with VMware vSAN. Entry level configurations include a single disk group with 6Gbps SATA SSDs, the priority being cost containment. High end designs offer dual DGs with 12Gbps SAS SSDs, the focus being performance and availability. Across the board various capacities are offered.

PowerEdge FC630 with PowerEdge FD332

Dell EMC PowerEdge FD332 provides the storage block for Dell EMC PowerEdge FC630. Each hyper-converged compute node consists of a PowerEdge FC630 with an associated PowerEdge FD332 as shown in the figure. PowerEdge FD332 is configured with Dell PERC FD33xD (Dual PERC) and up to 16 2.5" disk drives, which serve as the VMware vSAN controller and drives respectively. Dell PowerEdge RAID Controller H330 on board the FC630 blade is used as the boot controller with two 10k SAS drives in RAID-1 are used as the boot device.

NOTE: The Dell EMC PowerEdge FC630 with the Dell EMC PowerEdge FD332 support only all flash configuration.

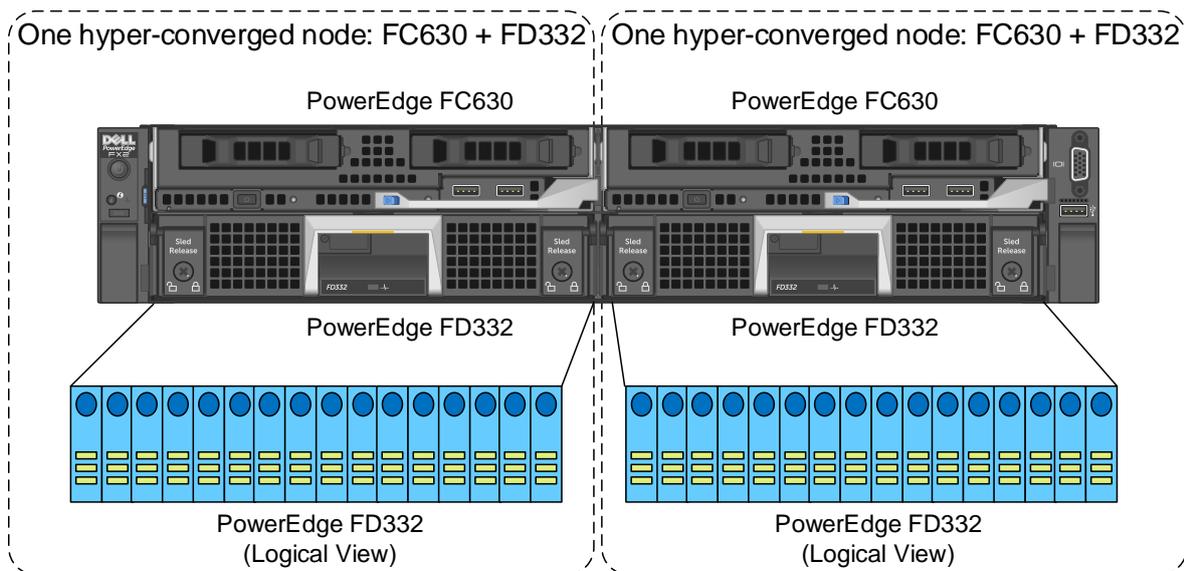


Figure 7. PowerEdge FC630 with PowerEdge FD332 – Disk Layout

PowerEdge FC430 with PowerEdge FD332

The Dell EMC PowerEdge FD332 provides the storage block for Dell EMC PowerEdge FC430. Each hyper-converged compute node consists of a PowerEdge FC430 associated with eight drives in a PowerEdge FD332 as shown in the figure. PowerEdge FD332 is configured with Dell PERC FD33xD (Dual PERC) and up to 16 2.5" disk drives, which serves as the VMware vSAN controller and drives respectively. On board SATA controller is used as the boot controller with one local 1.8" SSD drive used as the boot device.

NOTE: The Dell EMC PowerEdge FC430 with the Dell EMC PowerEdge FD332 support only all flash configuration.

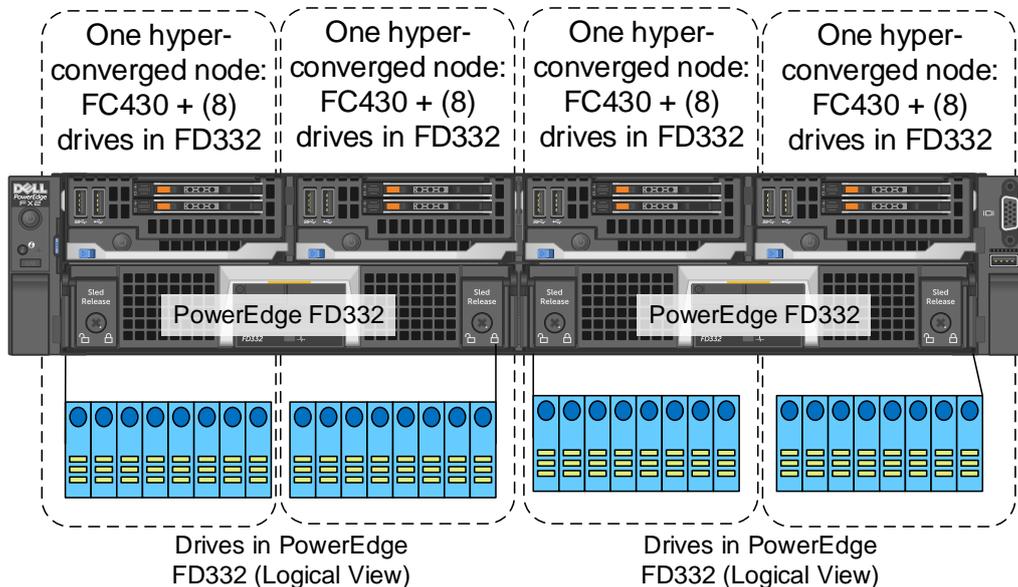


Figure 8. PowerEdge FC430 with PowerEdge FD332 – Disk Layout

PowerEdge rack servers

The Dell EMC PowerEdge R630, PowerEdge R730, and PowerEdge R730xd can be selected as a compute server. The PowerEdge R630 is used as a management server. SATADOM is a flash device that uses SATA interface and is the boot device for the Dell PowerEdge rack servers. SATADOM is cost efficient, reliable, and supports ESXi installation and VMware vSAN traces with no need to redirect vSAN traces.

The Dell PowerEdge HBA330 is used as the storage controller for vSAN drives. The PowerEdge HBA330 is a 12 Gbps SAS HBA card that integrates the latest enhancements in PCI Express® (PCIe) and SAS technology. Eight lanes of PCI Express 3.0 provide fast signaling for vSAN technology. PowerEdge HBA330 does not support hardware RAID, which is not required for vSAN, thereby reducing complexity.

Fault tolerance

The number of Failures to Tolerate (FTT) is a parameter that must be configured in VMware vSAN, indicating the number of fault domains (defined as a single node by default) that can fail without loss of data. It is recommended to configure FTT to a value of 1 for general purpose. FTT can be configured higher for specific virtual machine and application needs. FTT=0 is not recommended, as it provides no redundancy.

Supported Disk Configurations

The following table lists all the supported disk configurations:

Table 4. Supported Disk Configurations

Type	Server	Cache Drive	Capacity Drive	Disk Groups
HY-2	PowerEdge R630	(1) 400 GB SAS WI HUSMM	(4) 1 TB 7.2K NL-SAS SAS 10K RPM	1
HY-4	PowerEdge R630	(1) 400 GB SAS WI HUSMM	(7) 1 TB 7.2K NL-SAS SAS 10K RPM	1
HY-6	PowerEdge R730	(1) 800 GB SAS WI HUSMM	(7) 1.2 TB 10K SAS 10K RPM	1
HY-6	PowerEdge R730xd	(2) 400 GB SAS WI HUSMM	(12) 1 TB 7.2K NL-SAS 10K RPM	2
HY-8	PowerEdge R730	(2) 400 GB SAS WI HUSMM	(14) 1.2 TB SAS 10K RPM	2
HY-8	PowerEdge R730xd	(3) 400 GB SAS WI HUSMM	(18) 1.2 TB SAS 10K RPM	3
AF-4	PowerEdge R630	(1) 400 GB SAS WI HUSMM	(4) 1.92 TB SAS RI PM1633a	1
AF-4	PowerEdge R730	(1) 800 GB SAS WI HUSMM	(7) 1.92 TB SAS RI PM1633a	1
AF-6	PowerEdge R630	(2) 800 GB SAS WI HUSMM	(8) 1.92 TB SAS RI PM1633a	2
AF-6	PowerEdge R730xd	(2) 800 GB SAS WI HUSMM	(8) 1.92 TB SAS RI PM1633a	2
AF-8	PowerEdge R730	(2) 800 GB SAS WI HUSMM	(12) 1.92 TB SAS RI PM1633a	2
AF-8	PowerEdge R730xd	(2) 800 GB SAS WI HUSMM	(12) 1.92 TB SAS RI PM1633a	2
AF-6	PowerEdge FC430	(1) 800 GB SAS WI HUSMM	(6) 1.92 TB SAS RI PM1633a	1
AF-6	PowerEdge FC630	(1) 800 GB SAS WI HUSMM	(6) 1.92 TB SAS RI PM1633a	1
AF-8	PowerEdge FC630	(2) 800 GB SAS WI HUSMM	(12) 1.92 TB SAS RI PM1633a	2

Management Design

Management Infrastructure

The management infrastructure consists of three R630 servers that form a VSAN cluster. Management components are virtualized to provide high availability. The bundle further protects these components using the dedicated management cluster. Redundant 10 Gb Ethernet uplinks to the network infrastructure combined with vSphere High Availability ensure that management components stay online. A Dell Networking S3048 switch is used for OOB connectivity. iDRAC ports in each management and compute cluster connect to this switch.

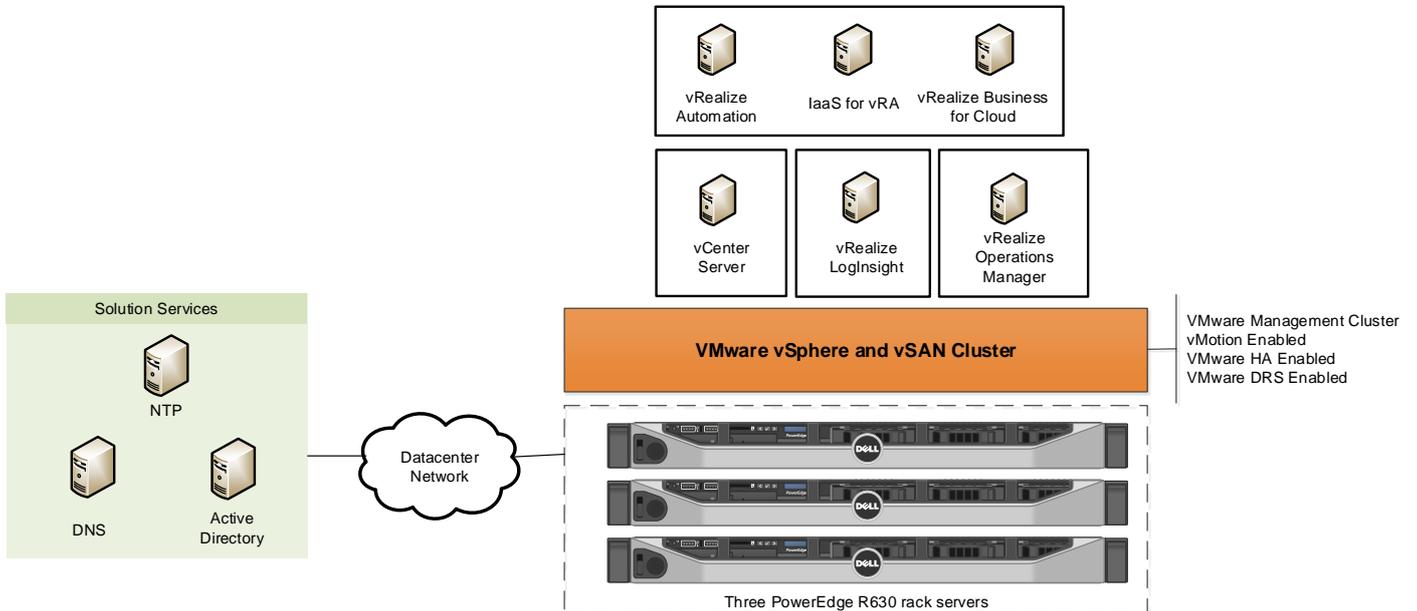


Figure 9. Management Design

Management Components

The following are management components:

- VMware vCenter Server Appliance
- VMware vRealize Automation
- VMware vRealize Business
- VMware vRealize Log Insight
- VMware vRealize Operations Manager

The management software components run on virtual machines which reside in your management cluster. The following table lists the management components in the bundle and the VM sizing of those components:

Table 5. Management components sizing

Component	VMs	CPU cores	RAM (GB)	OS (GB)	Data (GB)	NIC
VMware vCenter Server	1	4	12	12	256	1
vRealize Operations Manager	1	4	16	20	254	1
vRealize Automation (vRA)	1	4	18	50	90	1
vRealize Business	1	4	8	50	0	1
vRealize Log Insight	1	4	8	20	511	1
vRealize IaaS (for vRA)	1	4	6	80	0	1

VMware vRealize Automation Design

The VMware vRealize Automation architecture can be deployed using several models to provide flexibility in resource consumption, availability, and fabric management. The small, medium, and large deployment models all scale up as needed.

Note: Do not use the Minimal Deployment model described in the vRealize Automation documentation. This model is for proof of concept deployments only and cannot scale with the environment.

vRealize Automation uses DNS entries for all components. When scaling, vRealize Automation uses load balancers to distribute the workload across multiple automation appliances, web servers, and infrastructure managers. These components can be scaled as needed and distributed in the same data center or geographically dispersed. In addition to scaling for workload, vRealize Automation with load balancers is highly available. This feature ensures that users who consume vRealize Automation XaaS blueprints are not affected by an outage.

vRealize Automation components

The following are the components for vRealize Automation:

- VMware vRealize Automation Appliance - Contains core VRA services and the vRealize Orchestrator service, AD sync connectors, and an internal appliance database. Clusters A/A except DB failover is manual.
- Infrastructure web server - Runs IIS for user consumption of vRealize services and is fully active-active.
- Infrastructure Manager Service - Manages the infrastructure deployed with VRA. For small/medium deployment, this service runs on the Infrastructure Web Server and is active-passive with manual failover.
- Agents - Agents are used for integration with external systems such as Citrix, Hyper-V, and ESXi.
- Distributed Execution Manager - Performs the orchestration tasks necessary. When using multiple DEM worker servers, Distributed Execution Orchestrators are used to assign, track, and manage the tasks given to the workers.
- MSSQL - Tracks infrastructure components, supports Always ON for SQL Server 2016, otherwise a SQL failover cluster is necessary for HA.

VMware vRealize Automation deployment model sizing

VMware vRealize Automation deployment can be small, medium, or large based on the requirements. The following table provides a summary of the three deployment options:

Table 6. vRealize Automation deployment options

	Small	Medium	Large
Managed Machines	10,000	30,000	50,000
Catalog Items	500	1,000	2,500
Concurrent Provisions	10	50	100
vRA Appliances	1	2	2
Windows VMs	1	6	8
MSSQL DB	Single server	Failover cluster	Failover cluster
vRealize Business Appliance	1	1	1
Load Balancers	0	3	3

The following images illustrate a large deployment:

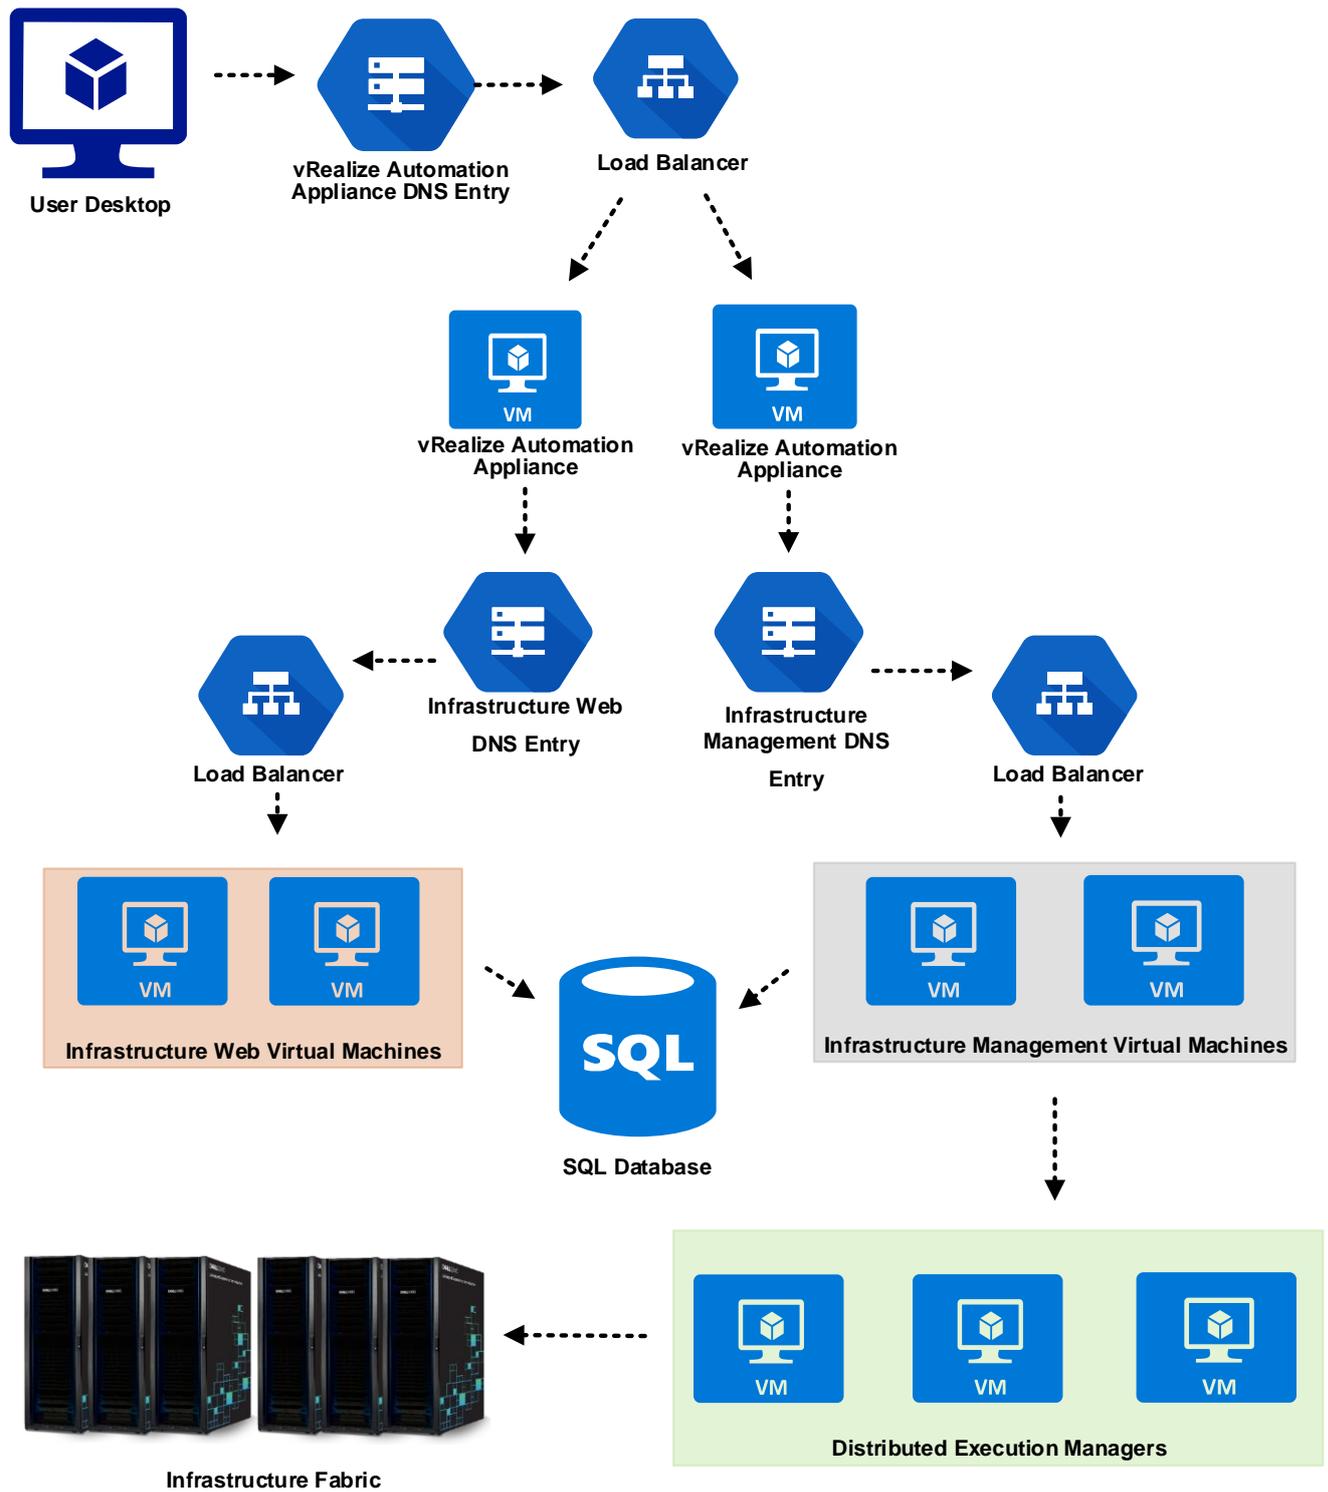


Figure 10. vRealize Automation – Large Deployment

Scaling the Ready Bundle

The solution can be scaled by adding multiple compute nodes (pods) in the customer data center. The Dell 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 Gb, 25 Gb, 40 Gb, 50 Gb, and 100 Gb Ethernet connectivity and can aggregate multiple racks with little or no oversubscription. When connecting multiple racks, using the 40 Gb Ethernet 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.

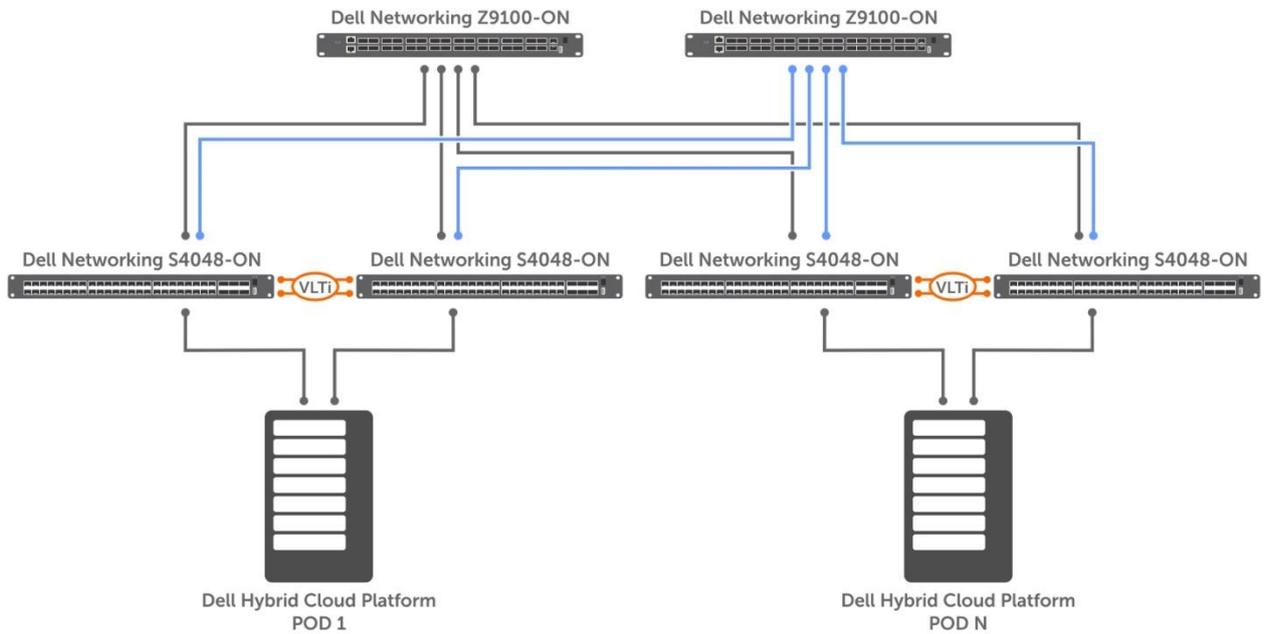


Figure 11. Multiple Compute PODs scaled out using leaf spine architecture

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 GbE links to the core network providing enough bandwidth for all the traffic between each rack.

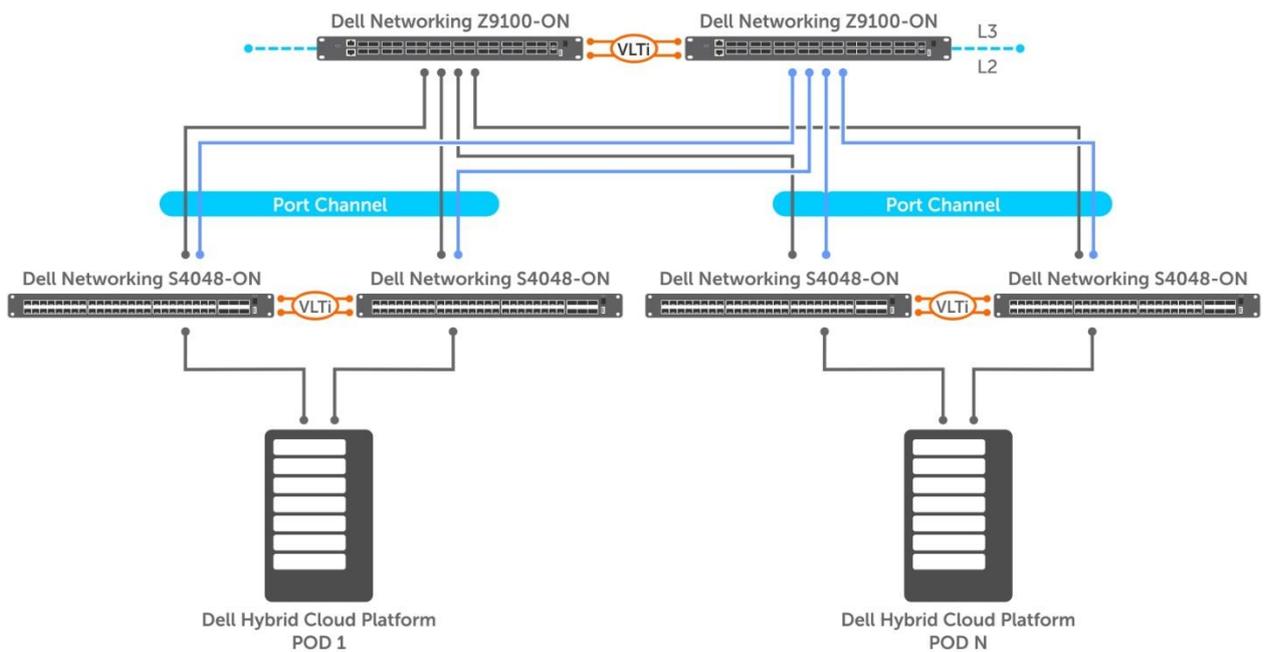


Figure 12. Multiple Compute PODs scaled out using eVLT

References

- [Administering VMware vSAN 6.5](#)
- [VMware vSAN Design and Sizing Guide](#)
- [Boot device considerations for vSAN](#)
- [vRealize Automation – Reference Architecture](#)
- [Dell VLT Reference Architecture](#)
- [Dell Configuration Guide for S4048-ON](#)