Store IT Modernization with VMware

White Paper

H19145

Abstract

This document describes how retailers can use a Dell/VMware Edge validated solution to modernize their store IT infrastructure and increase digital maturity by virtualizing their siloed application stacks and orchestrating the deployment and management of applications at scale from a single-pane-of-glass, control panel approach.

Dell Solutions

April 2022

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Introduction

Overview

Edge computing in the retail industry plays a key role in transforming our world. Today, organizations in the retail industry want to deliver new services and digital experiences for their employees and customers close to the locations where data is produced and consumed. These organizations want to combine data originating at the store with machine learning (ML), analytics, and artificial intelligence (AI) to help improve operations, become more agile, and capture next-generation business opportunities. To ensure success, organizations require retail edge infrastructure solutions that enable fast deployment of new IT services, simplified IT management, and security of edge infrastructure at scale.

Customers need a retail edge infrastructure solution that is simple, flexible, and easy to deploy. Efficient and resilient network connectivity is critical to ensure faster deployment of new IT services across data centers, branch offices, and store locations.

To meet this need, Dell and VMware developed a combined retail edge solution architecture, which is built with VMware Edge Compute Stack and Dell hyperconverged platforms that span from the fully curated experience of Dell EMC VxRail to the customer-managed VMware vSAN Ready Node built on Dell PowerEdge Servers. The Dell portfolio includes several options that are short-depth, ruggedized, or tower-based, delivering the same benefits as the solution architecture in this paper. These solutions deliver a cost-efficient, simple, and flexible architecture that ensures faster deployment of new IT services for the customers in the retail industry.

The objective of this white paper is to define this retail edge solution architecture and help readers understand the design principles of the combined solution from Dell and VMware and the various outcomes enabled by this solution, with an emphasis on the retail industry.

Audience

This solution architecture white paper is intended for architects, engineers, consultants, and IT administrators responsible for designing, implementing, and operating edge infrastructure, as well as Dev Ops teams in charge of managing lifecycle application development and delivery.

Business case

As business applications and workloads evolve and expand from the data center and the cloud to the locations where data is produced and consumed, new edge compute, network technologies, and architectures are needed to help companies deliver edge infrastructure everywhere. Customers in the retail industry are looking to combine data originating at the store with data analytics and AI/ML technologies to help improve operations and deliver new services and digital experiences. To achieve this goal, retail customers need to modernize their stores with infrastructure designed for the edge. They need faster compute to enable data-driven personalization experiences and store security, and low-latency connections for faster data analytics.

The edge is a distributed digital infrastructure for running workloads across multiple locations that are placed close to endpoints producing and consuming data. Edge computing helps businesses process, collect, and transfer data at the edge in near-real time with ultra-low latency. As a result, not only must these distributed platforms be correctly implemented, managed, and maintained, but they must also be secured to protect the edge infrastructure, customer data integrity, and business operations. As an example, when massive files need to be quickly transferred to distributed end users, there are often WAN-related concerns including bandwidth congestion, connectivity loss, or delay in the transfer of these files. Retailers need a WAN solution that alleviates these concerns.

The edge landscape at most retailer stores today is complex and fragmented with siloed physical full-stack architectures for every app, no single source of truth for data, and a lack of interoperability and efficiency. This results in increased IT infrastructure management and labor costs, as well as increased time to value for new experiences and services to be launched at scale across multiple retail store locations. As retailers add applications and enable new capabilities in this existing framework, the problems compound infinitely. This is a critical impediment to retailer success, and solving this is fundamental to technology maturity and digital transformation.

Retailers want to move to an edge architecture and ecosystem that is simple, highly available, convenient, scalable, secure, and seamless across multiple clouds and edge devices. They want the ability to consolidate all apps and workloads on a single appliance that effectively optimizes compute, storage, networking, and security. Such a solution should offer capabilities to monitor and manage various elements of the edge ecosystem at scale from a single pane of glass. Retailers are obviously sensitive to costs, so edge solutions must be designed to meet the cost and performance envelopes specific to various retailers.

As customers look at modernizing their edge and remote office/branch office (ROBO) infrastructure, they would prefer to deploy a pre-integrated and fully validated solution that is easy to manage and scale and that also ensures application performance through reliable and secure connectivity at the edge. Organizations value efficiency, ease of use, and a simplified environment. VMware and Dell Technologies are uniquely positioned to deliver a multicloud edge platform. Together, we are jointly developing edge solution architectures today to help customers to:

- Build, run, and manage edge-native applications at the near and far edge.
- Leverage their infrastructure and operations across multiple clouds.
- Enable remote access with assured application performance and secure the applications from various threats in a distributed environment.

We developed a combined retail edge solution architecture including a Dell PowerEdge-based hyperconverged platform, Dell EMC vSAN Ready Node built on the industry-leading PowerEdge servers, VMware vSphere, VMware vSAN, VMware Tanzu, and VMware SD-WAN. This solution delivers a cost-efficient, simple, flexible, and easy-to-use architecture that ensures deployment of new IT services quickly, efficiently, and securely.

Dell vSAN Ready Nodes are preconfigured, tested, and certified to run VMware vSAN. Each Ready Node includes just the right amount of CPU, memory, network I/O controllers, HDDs, and SSDs. Dell also offers premiere vSAN Ready Node configurations, in which each model boasts an Identity Module that self-identifies the server as a vSAN Ready Node upon boot-up to streamline deployment, updates, and more. In this white paper, we will focus on Dell vSAN Ready Node XR11 in a retail edge solution architecture.

Use cases

Retailers can use the purpose-built edge store IT modernization solution, combining Dell hardware and VMware Edge Compute Stack software, to efficiently virtualize various existing applications and workloads. This provides retailers with capabilities to manage adding, deleting, or updating services and customer experiences quickly, at scale, from one centralized command center and administrator within the solution. This solution enables retailers to streamline IT infrastructure and processes, simplify distributed edge/core/cloud ecosystems, and successfully support agile execution of digital transformation initiatives.

Some of the use cases are:

- Point of sale
- Inventory management
- Theft detection and prevention
- Real-time data analytics
- Video analytics and computer vision AI
- Self-service and mobile payment
- Cashier-less shopping

Figure 1. Use cases for the retail edge solution architecture

Key benefits

- **Resilient operations** using pre-integrated and validated solution design for retail applications and workloads.
- **Faster response time** deploying standardized infrastructure that supports both VMs and containers.
- **Secure and reliable** SD-WAN connection to remote sites, assuring application performance through optimized connectivity and PCI compliance.
- **Cost effective** solution architecture that is comprehensive, extensible, and scalable.
- **Purpose-built** Dell EMC vSAN Ready Node built on PowerEdge servers. Reimagined for edge computing with powerful compute and support for multiple accelerators, this compact server has the agility needed for retail edge applications and workloads.
- **Integrated lifecycle management** from Dell and VMware, so your company can leverage existing VMware skillsets.

2

Solution Components

The retail edge solution architecture includes the following VMware components. The versions of VMware components are described in Technical Specifications.

Solution components:

- Dell EMC vSAN Ready Node
- VMware Edge Compute Stack
  - vCenter Server Standard
  - vSphere Hypervisor (ESXi) Enterprise Plus
  - vSAN Advanced
  - Tanzu Standard
- VMware SD-WAN

Dell vSAN Ready Node

Dell PowerEdge servers are designed to provide the highest performance for a diverse set of workloads. PowerEdge servers configured for vSAN (Dell vSAN Ready Nodes) are jointly engineering-validated and certified, which reduces deployment risks, improves storage efficiency, and quickly and easily scales as needed. The Dell vSAN Ready Node portfolio is one of the broadest in the industry, offering one to four sockets, in 1U/2U rack, tower, and modular solutions for data center, cloud, and edge workloads. For more detail, see the Solution Brief: Build your hybrid cloud with a truly flexible foundation.

An example of this portfolio that is specifically designed for edge computing is the Dell vSAN Ready Node XR11.

The PowerEdge XR11 is a ruggedized, dual-socket, 1U server that is designed to support demanding edge applications such as streaming retail analytics, manufacturing logistics, 5G cell processing applications, and so on. Edge and non-data center environments can place great environmental stresses on the electrical and physical operation of the servers. The XR11 is purpose-built to operate under extended system operating temperatures from –5° C to 55° C (23° F to 131° F). There is an optional filter built into the bezel that is designed to remove particulate matter. This build is best suited for:

- Demanding applications at the edge
- Retail applications and analytics
- Manufacturing and logistics applications

VMware Edge Compute Stack

VMware Edge Compute Stack provides a consistent multicloud edge platform that enables customers to build, run, manage, connect, and protect their industry-specific edge-native applications at the near and far edge, while also leveraging consistent infrastructure and consistent operations across their data centers and cloud. VMware Edge Compute Stack is a purpose-built integrated stack for small scale VM and container compute, with edge HCI and SD-WAN that enables organizations to deploy edge-native applications at the far edge. It comes in three editions: Standard, Advanced, and Enterprise. The solution leverages vSphere, vSAN, and Tanzu Standard (Tanzu Kubernetes Grid and Tanzu Mission Control), and optionally VMware SD-WAN to deliver a platform on which customers can develop and innovate for edge-native applications.

VMware Tanzu Kubernetes Grid (TKG) is a multicloud Kubernetes footprint that you can run both on-premises in vSphere and at the edge. TKG provides a consistent, upstream-compatible implementation of Kubernetes that is tested, signed, and supported by VMware. TKG provisions Kubernetes clusters in a declarative manner that is familiar to Kubernetes operators and developers, and manages the lifecycle of the Kubernetes clusters. In addition, TKG includes signed and supported versions of open-source applications to provide the registry, networking, monitoring, authentication, ingress control, and logging services that a production Kubernetes environment requires. For TKG details, see VMware Tanzu Kubernetes Grid Documentation.

VMware SASE

VMware SASE (Secure Access Service Edge) is a cloud-first, cloud-native platform that combines industry-leading SD-WAN capabilities with cloud-delivered security, including cloud web security, zero trust network access, and firewalling. It provides branch, home, and remote users secure, optimized, and reliable access to modern applications deployed in public/private clouds, SaaS, or in the future, the edge.

VMware SD-WAN

VMware SD-WAN, part of VMware SASE, optimizes the network to assure application performance with security despite network issues, be it packet loss, bandwidth congestion, latency, or jitter. VMware SD-WAN can deliver this kind of network optimization with one of its differentiating features, Dynamic Multipath Optimization (DMPO). By combining application recognition, traffic prioritization, and shaping with the ability to measure network path performance, the solution steers traffic on a packet-by-packet basis to achieve the highest quality of experience for end users.

VMware SD-WAN simplifies operations through a cloud-hosted management platform that centralizes network and security policy creation, distribution, and control. The solution visualizes application delivery performance while employing AIOps to suggest remediation actions for LAN, WAN, or data center devices to streamline troubleshooting.

Retail edge solution architecture design

The retail edge solution architecture includes Dell vSAN Ready Node XR11 and VMware SD-WAN to deliver a comprehensive solution that is reliable, secure, and ensures a higher level of application performance across the data center to the edge in remote, distributed locations.

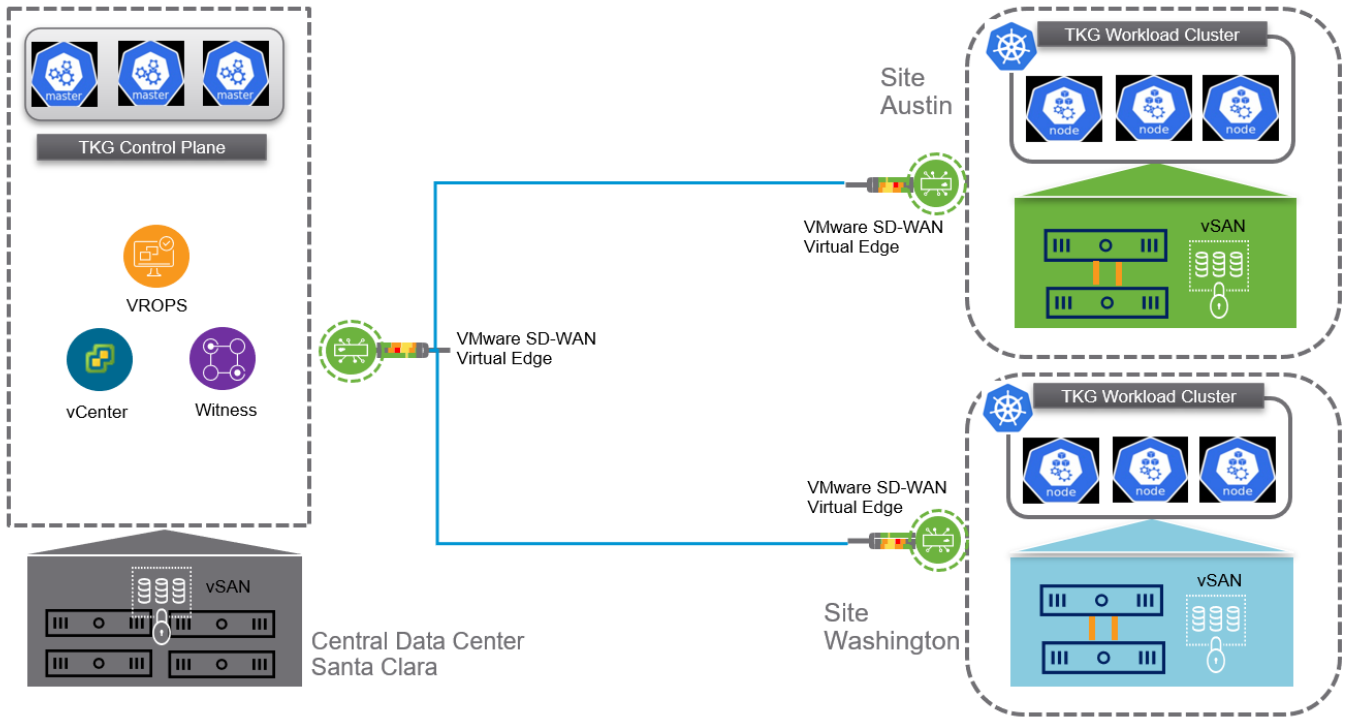

Figure 2. Retail edge solution architecture

The retail edge is a two-node cluster with two directly-connected XR11 nodes, a shared VMware vSAN witness host, and a VMware SD-WAN Virtual Edge. The two-node cluster is deployed at each edge location and managed by VMware vCenter Server from a central data center.

This retail edge solution architecture provides hardware in the form of two physical XR11s to host VMs or containers and has enough resources to place additional virtual appliances such as VMware SD-WAN Edge for a secure tunnel to your edge, the data center, or the cloud. VMware SD-WAN provides enhancements to an edge deployment that are complementary to the native networking provided by service providers. VMware SD-WAN delivers increased network agility and cost reduction by using a software-defined approach to abstracted network hardware and by transporting characteristics from the applications that are using the data center and edge network.

The following figure shows the three-site deployment architecture options of two-node XR11 infrastructure at the edge with VMware SD-WAN as network connectivity.

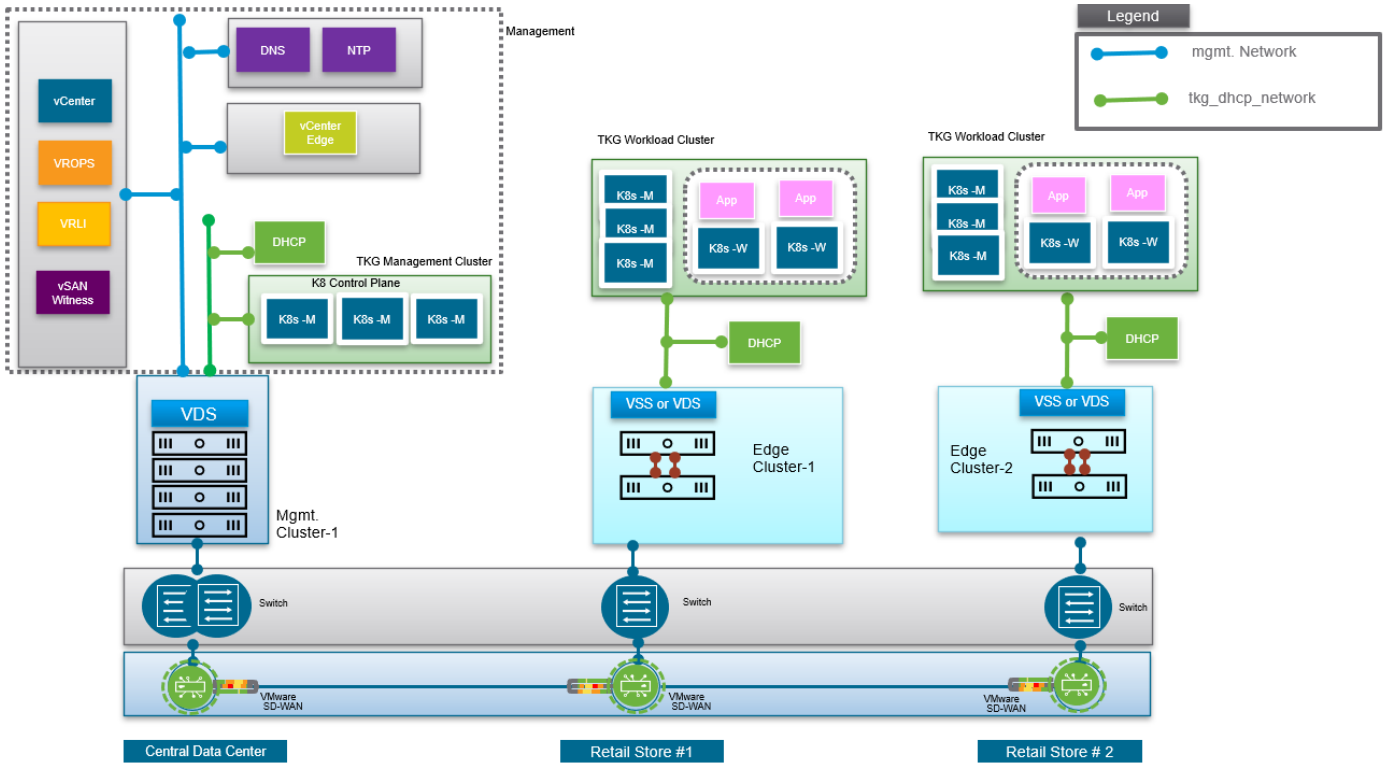

Figure 3. Architecture from the central data center to the edge

- In this deployment, the central data center hosts all management components, including DNS, VMware vCenter Server, vSAN witness appliance, VMware vRealize Operations Manager, and vRealize Log Insight.
- A two-node retail edge cluster is connected back-to-back with 10 GbE for vSAN and vMotion traffic and hosted at each retail store.
- The central data center and edge clusters are connected through a VMware SD-WAN L3-routed network.
- The vSAN witness appliance is hosted on the management cluster at the central data center and witness traffic is connected with a L3-routed network over the VMware SD-WAN solution.
- Kubernetes control plane nodes are hosted on the management cluster at a central data center, and worker nodes are hosted on a two-node retail edge cluster for retail applications.
- Kubernetes clusters use a dedicated local DHCP server.
- A dedicated edge vCenter is hosted on the management cluster at the central data center and manages all the two-node retail edge clusters for high availability and reliability. It also provides better management and separation of vCenter roles between management and workloads.
- A dedicated management vCenter is hosted on the management cluster at the central data center for managing all management resources, clusters, and management VMs.

NOTE: vRealize Operations and vRealize Log Insight are optional components.

NOTE: If you would like to add more nodes, add them in groups of two in a cluster. Once you reach three nodes, you can remove the vSAN witness as a third node.

Virtual infrastructure design

vCenter server design

The vCenter Server design includes the design for all vCenter Server instances. For the retail edge, determine the number of instances, their sizes, networking configuration, vSphere cluster layout, redundancy, and security configuration.

Although vCenter Server is deployed at a central data center and manages all the retail edge clusters, it is critical to ensure that your vCenter is designed appropriately before you onboard retail edge clusters and your applications.

A vCenter Server deployment can consist of two or more vCenter Server instances according to the scale, number of VMs, and the continuity requirements for your environment. You must protect the vCenter Server system, as it is the central point of management and monitoring. You can protect vCenter Server according to the maximum downtime tolerated. Use the following methods to protect the vCenter Server instances:

- Automated protection using vSphere HA
- Automated protection using vCenter Server HA

Recommended vCenter server design

The following bullets detail the vCenter server design recommendations by describing each design decision, as well as the justification:

- Deploy two vCenter Server systems.
  - Isolates vCenter Server failures to management or compute workloads.
- One vCenter Server supports the management workloads.
  - Isolates vCenter Server operations between management and compute workloads.
- Another vCenter Server supports the compute workloads at the edge.
  - Supports a scalable vSphere cluster design where you might reuse the management components as more compute workload domains are added.
  - Simplifies capacity planning for compute workloads because you do not consider management workloads for the Compute vCenter Server.
  - Improves the ability to upgrade the vSphere environment and related components by the separation of maintenance windows.
  - Supports separation of roles and responsibilities to ensure that only administrators with proper authorization can attend to the management workloads.
  - Facilitates quicker troubleshooting and problem resolution.
- Protect all vCenter Servers by using vSphere HA.
  - Supports the availability objectives for vCenter Server without the required manual intervention during a failure event.

ESXi host design

Ensure that the physical specifications of the ESXi hosts allow for successful deployment and operation of the retail edge design.

Recommended ESXi host design

The following bullets detail the ESXi host design recommendations by describing each design decision, as well as the justification:

- Ensure that all ESXi hosts with XR11 have a uniform configuration across the retail edge clusters.
  - Provides ease of management and maintenance across the cell sites.
- Set up each ESXi host with a minimum of four physical NICs.
  - Ensures full redundancy for the required two physical NICs for management.
  - The other two physical NICs will be dedicated for two-node vSAN and vMotion.
- Set up each ESXi host in the cluster with ESXI boot drive and vSAN for workloads.
  - vSAN is the primary storage solution for retail edge. Consider the disk size based on retail workloads.

- Set up each ESXi host in the retail edge location with a recommended minimum of 192 GB RAM.
  - Provides a good starting point for most workloads.
  - Allows for ESXi and other management overhead.

NOTE: See Hardware for more details.

vSAN witness design

VMware vSAN two-node architecture is a perfect solution for organizations that have many small branch offices or retail sites. It is also beneficial for small businesses and startups who want to avoid the significant up-front costs associated with storage hardware. The shared witness host appliance reduces the physical resources that are needed at the central data center, resulting in a greater level of savings for **a** large number of two-node retail store deployments.

vSAN data traffic requires a low-latency, high-bandwidth link. Witness traffic can use a high-latency, low-bandwidth, and routable link. To separate data traffic from witness traffic, configure a dedicated VMkernel network adapter for vSAN witness traffic, as shown in the Network design section of this document.

NOTE: Minimum network connectivity requirements for the witness virtual appliance are:

- 1.5 Mbps bandwidth connectivity
- 500 milliseconds latency RTT

Local vSAN witness design

An alternate vSAN witness design allows for the witness VM to be deployed locally at the same location as the vSAN cluster. This design uses a Dell VEP appliance for SD-WAN and witness VMs.

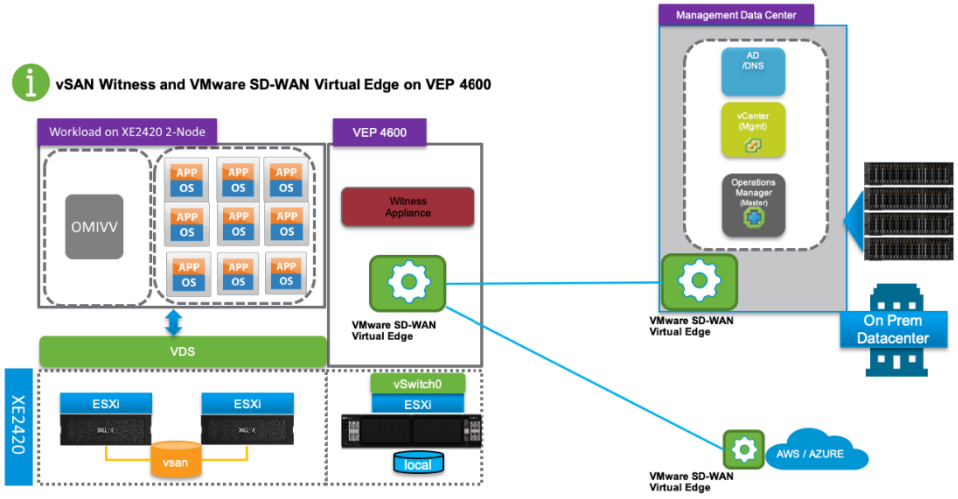

Figure 4. Local vSAN witness design

Platform scaling

Retail stores at edge locations scale depending on the resource and availability requirements of applications.

One single shared witness can be shared across a maximum of 64 two-node clusters, supporting up to 64,000 components, and requires at least six CPUs and 32 GB memory allocation for the witness.

When designing the retail edge solution architecture for scale, review the VMware maximum supported configurations for your scaling requirement.

Some of the supported configurations of VMware maximums are:

Table 1. VMware configuration maximums

Type	Maximums
Hosts per vCenter	2,500
VDS per vCenter	128
Hosts per VDS	2,000
1 x vSAN witness (shared)	64 x two-node clusters

The following is a high-level example of scaling requirements based on VMware maximums.

NOTE: Consult with your Dell account representative before you consider the scale and sizing of your environment.

Table 2. Scaling requirements

Number of sites	Requirements
1,000 retail edge sites	1 x Management vCenter
	1 x Edge vCenter
	16 x vSAN witness appliance **NOTE:** Each witness appliance can support 64 two-node clusters.
2,000 retail edge sites	1 x Management vCenter
	1 x Edge vCenter
	32 x vSAN witness appliance

vSphere Lifecycle Manager

vSphere Lifecycle Manager (vLCM) manages software and firmware lifecycles of the ESXi hosts in a cluster with a single image. vSphere Lifecycle Manager images are a new functionality that provides a simplified and unified workflow for patching and upgrading ESXi hosts. You can also use vSphere Lifecycle Manager images for bootstrapping purposes and firmware updates.

An image defines the exact software stack to run on all ESXi hosts in a cluster. When you set up an image, you select an ESXi version and a vendor add-on from the vSphere Lifecycle Manager depot. If no ESXi base images and vendor add-ons are available in the vSphere Lifecycle Manager depot, you must populate the depot with software updates by synchronizing the depot or by manually uploading updates to the depot.

vSphere Lifecycle Manager remediates hosts that are part of a vSAN cluster sequentially. By design, only one host from a vSAN cluster can be in maintenance mode at any time.

For edge locations, instead of accessing the vSphere Lifecycle Manager depot in vCenter Server, clusters in retail edge deployments can download data from a depot that is local for them. You can configure vSphere Lifecycle Manager to use local depots for any cluster that uses images. A retail edge cluster has limited or no access to the Internet or limited connectivity to vCenter Server. As a result, clusters in retail edge deployments might have limited access to the vSphere Lifecycle Manager depot during the compliance check, remediation pre-check, and remediation operations. With vSphere Lifecycle Manager images, you can use a local depot for Retail Edge clusters and configure vSphere Lifecycle Manager to use the local depot during the compliance check, remediation pre-check, and the remediation tasks. The local depot overrides the vSphere Lifecycle Manager depot. Using local depots with ROBO clusters saves time and network bandwidth.

For each cluster that you manage with a single image, you can add and use multiple local depots instead of the default vSphere Lifecycle Manager depot. You can also delete the depot overrides that you configure. If depot overrides are not active for a cluster, the cluster uses the general vSphere Lifecycle Manager depot in vCenter Server.

Prerequisites:

- Set up an online depot to which the cluster can connect.
- Export an offline bundle with components from a vSphere Lifecycle Manager image and import the offline bundle to the target local depot.
- Verify that you have the proper privileges. See the *vSphere Lifecycle Manager Privileges for Using Images* section in the vSphere Managing Host and Cluster Lifecycle documentation.

OpenManage integration for VMware vCenter

Dell EMC OpenManage gives customers system management capabilities and tools to assist in full lifecycle management and inventory control of a fleet of PowerEdge servers. OpenManage provides capabilities in monitoring, catalogs and change management, APIs, and ecosystem integrations and connections. The foundation of the OpenManage industry-leading portfolio is the integrated Dell Remote Access Controller (iDRAC).

OpenManage Integration with VMware vCenter (OMIVV) is the Dell EMC PowerEdge server administration tool that manages directly within the VMware vCenter environment. OMIVV provides:

- Monitoring and alerts views in one place, both physical and virtual
- Firmware updates managed from within vCenter
- Expedited server deployment

vLCM and OMIVV together simplify firmware and software management and accelerate hypervisor and firmware updates from hours to minutes for customers.

For more details, see Dell EMC OpenManage Ecosystem Portfolio: Integrations, Connections and RESTful APIs and the vSphere 7.0 Compatibility Study from Principled Technologies.

Network design

The following figure shows the network design of a central data center and edge site with a two-node XR11, vSAN witness, and VMware SD-WAN.

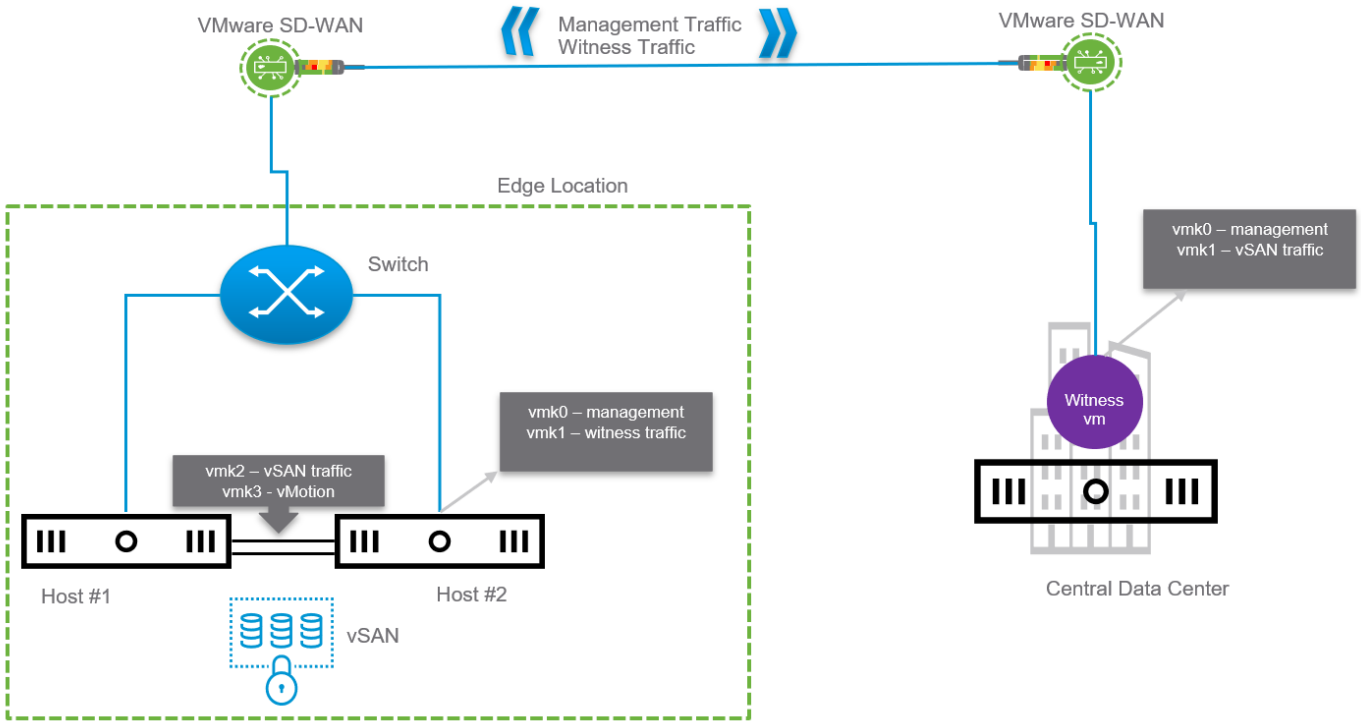

Figure 5. Retail edge network connectivity diagram

The configuration in the figure is as follows:

Table 3. Retail edge cluster: Host #1 and Host #2

VMkernel Network	Traffic Type
vmk0	Tagged for management traffic
vmk1	Tagged for witness traffic[a]
vmk2	Tagged for vSAN traffic

Table 3. Retail edge cluster: Host #1 and Host #2 (continued)

VMkernel Network	Traffic Type
vmk3	Tagged for vMotion traffic

a. This must be done using esxcli vsan network ip add -i vmk1 -T=witness on each host.

Table 4. vSAN witness appliance

VMkernel Network	Traffic Type
vmk0	Tagged for management traffic
vmk1	Tagged for vSAN traffic

The following figure shows how each physical NIC is mapped to each virtual switch in the ESXi host and each VMkernel network.

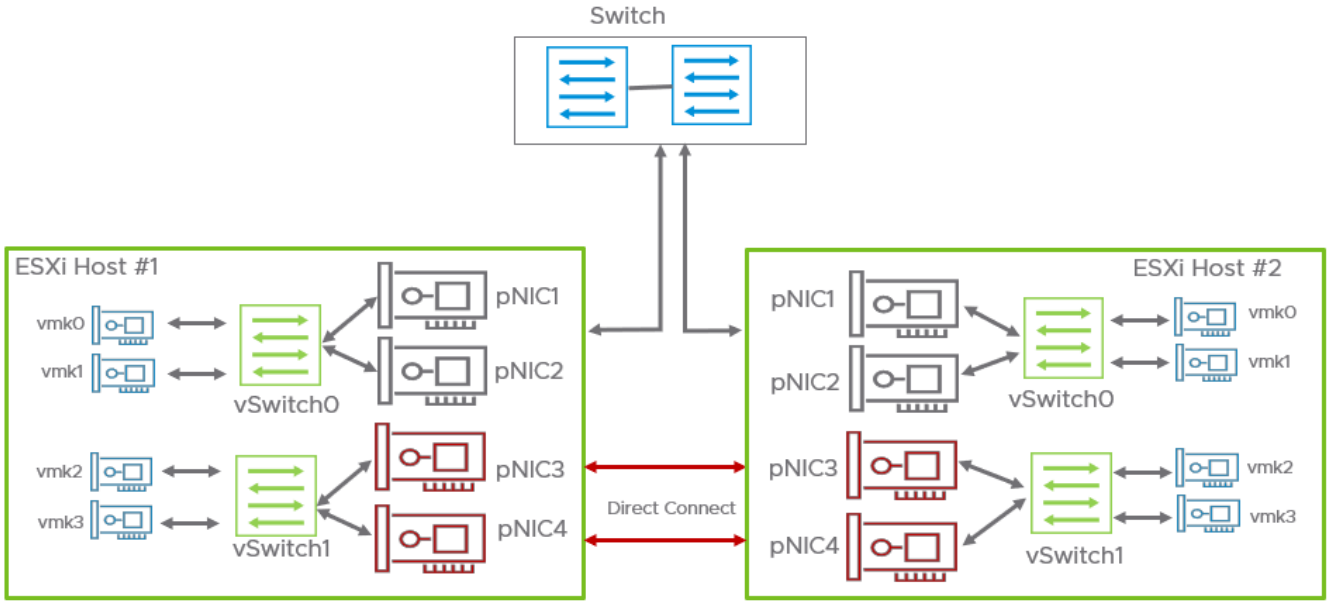

Figure 6. Physical NIC and virtual switch mapping

- vSwitch0 is mapped to pNIC1 and pNIC2, and vmk0 and vmk1 is configured with vSwitch0.
- vSwitch1 is mapped to pNIC3 and pNIC4, and vmk2 and vmk3 is configured with vSwitch1.
- pNIC1 and pNIC2 is connected to upstream management switch.
- ○ Witness
  pNIC3 and pNIC4 is directly connected to each host in the cluster.

NOTE: You have a choice to configure VMware Distributed Switch for this, however the VDS can only support 128 sites per vCenter.

Determine the number of networks or VLANs that are required, depending on the type of traffic. This includes the vSphere operational traffic and traffic that supports the organization's services and applications.

For retail edge, we use a dedicated VLAN for the following:

- Management
- Uplink
- Virtual machine

For the VMware vSAN witness, an additional witness VLAN and VM port group is configured on the two-node cluster.

Tanzu Kubernetes grid design

The Tanzu Kubernetes clusters are deployed in the edge vCenter.

Retail edge application consumes resources from the edge vCenter. Resource pools provide guaranteed resource availability to workloads. Resource pools are elastic, allowing more resources to be added as its capacity grows. Each Kubernetes cluster can be mapped to a vSphere resource pool. In a retail edge design with a central data center and edge sites, a Kubernetes control plane node can be placed on a vSphere cluster at the central data center, and worker nodes can be placed on a vSphere edge cluster at the edge to support application workloads. The vSphere clusters and vSphere host can be managed by edge vCenter, which is hosted in the central data center. See Architecture from the central data center to the edge for the architecture.

The following table describes design decisions to address when implementing Tanzu Kubernetes clusters.

Table 5. Tanzu Kubernetes clusters design recommendations

Design decision	Design justification
Map the Tanzu Kubernetes clusters to the vSphere Resource Pool in the edge vCenter.	Enables resource guarantee and resource isolation.
• Create dedicated DHCP IP subnet pools for the Tanzu Kubernetes cluster management network. • Dedicate a static IP for Kubernetes endpoint API. • After the cluster is deployed, create a DHCP reservation for each node deployed in the TKG cluster so that the IP addresses do not change.	• Simplifies the IP address assignment to Kubernetes clusters. • Exclude a block of IP addresses from the DHCP pool for static IP addresses to assign to the Kube-VIP address of each cluster.
Place the Kubernetes cluster management network on a virtual network, which is routable to the management network for vSphere.	vSAN is the primary storage solution for retail edge. Consider disk size based on retail workloads.
Set up each ESXi host in the retail edge location with a recommended minimum of 192 GB RAM.	• Provides connectivity to the vSphere infrastructure. • Simplifies the network design and reduces the network complexity.

NOTE: TKG does not use Kube-VIP as a load balance for workloads in workload clusters. Kube-VIP is used solely for the API server of the cluster. For an application ingress, consider deploying Contour within each edge cluster. Contour is an open-source ingress controller, primarily backed by VMware. Using an ingress controller, you can limit the number of external network routes to each cluster, then manage traffic using the standard Kubernetes configurations within to route requests to the correct application. See Implementing Ingress Control with Contour and Contour Reference Architecture for more information.

You can also propagate the OVA images for cluster deployment, which includes deploying the Kubernetes node OS template and HA proxy template to each edge location by using Harbor replication to propagate all the TKG extensions and required container images (for example, CSI, CNI) to the edge sites. This way, clusters can be created in limited network bandwidth and latency.

Consult with your Dell Technologies representative to implement the above design principles.

4

Technical Specifications

Hardware

The following table shows the technical specifications of this solution.

Table 6. Solution specifications

Specification	Quantity
Dell vSAN Ready Node XR11	2
VMware SD-WAN	1
10 GbE network switch	1

NOTE: For this solution validation, we used the VMware virtual SD-WAN solution. However, Dell recommends the use of a dedicated VMware SD-WAN appliance for each store for network reliability and resiliency.

NOTE: You have the option to use any managed 10 GbE network switch for your environment.

Software

The following table shows the minimum software version requirements for the retail edge two-node cluster.

Table 7. Software requirements

Components	Version	Build number
ESXi	7.0U2a	17867351
vCenter	7.0U2a	17920168
VSAN Witness Appliance	7.0U2	17630552
VMware SD-WAN Virtual Edge	3.3.2	N/A
Tanzu Kubernetes Grid	1.3.1	N/A

5

Conclusion

This solution architecture describes a recommended configuration of Dell vSAN Ready Node and VMware vSAN witness architecture and deployment in a retail edge location with an option for secure and reliable connectivity to their edge network using VMware SD-WAN. This solution continues to ease deployment and management. It enables an efficient centralized management with reduced hardware and software costs, while meeting the needs of environments with limited space, limited budget, and IT personnel constraints.

A

References

vSAN 2-Node Cluster Guide

VMware Tanzu at the Edge: Solution Architecture for ROBO Topology

Solution Brief: Build your hybrid cloud with a truly flexible foundation

VMware Tanzu Kubernetes Grid Documentation

Managing Host and Cluster Lifecycle

Dell EMC OpenManage Ecosystem Portfolio: Integrations, Connections and RESTful APIs

vSphere 7.0 Compatibility Study from Principled Technologies

Implementing Ingress Control with Contour

Contour Reference Architecture