



Retail Edge: A Modernized Edge Architecture for Retail

Table of contents

Executive summary	3
Overview	3
Audience	3
Business case	4
Use cases	4
Key benefits	5
Solution components	5
Dell EMC vSAN Ready Node	6
VMware Edge Compute Stack.....	6
VMware SASE	7
Retail edge solution architecture design.....	7
Virtual infrastructure design.....	9
vCenter Server Design	9
ESXi host design	10
vSAN witness design	10
Local vSAN witness design	10
Platform scaling	11
vSphere Lifecycle Manager	12
OpenManage integration for VMware vCenter	13
Network design	13
Tanzu Kubernetes grid design.....	15
Technical specifications	16
Hardware	16
Software	17
Conclusion	17
Additional Resources/References.....	17
Author	17
Acknowledgements	17

Executive summary

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 are produced and consumed. These organizations are looking to combine data originating at the store with machine learning, analytics, and artificial intelligence 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 EMC and VMware developed a combined retail edge solution architecture, built with VMware Edge Compute Stack and Dell EMC 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. Dell EMC portfolio includes several options that are short-depth, ruggedized or tower based, delivering the same benefits as the solution architecture portrayed 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 the white paper is to define this retail edge solution architecture and help understand design principles of the combined solution from Dell EMC and VMware.

Audience

This solution architecture whitepaper is intended for architects, engineers, consultants, and IT administrators responsible for designing, implementing, and operating edge infrastructure.

Business case

As business applications and workloads evolve and expand from data center and the cloud to the locations where data are produced and consumed, new edge compute, network technologies and architectures will be 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 artificial intelligence (AI)/machine learning (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, 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 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, many times there are WAN-related concerns including bandwidth congestion, connectivity loss, or delay in transfer of these files. Retailers need a WAN solution that alleviates these concerns.

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 which also ensures application performance through reliable and secure connectivity at the edge. It should come as no surprise that organizations value efficiency, ease of use and a simplified environment. VMware and Dell Technologies are uniquely positioned to deliver a multi-cloud 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 secure remote access with assured application performance and secure the applications from various security 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 EMC vSAN Ready Nodes are pre-configured, 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 EMC also offers premiere vSAN Ready Node configurations; 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 paper we will focus on Dell EMC vSAN Ready Node XE2420 in a retail edge solution architecture.

Use cases

The retail edge solution architecture provides a mini software-defined data center (SDDC) solution for the retail industry. 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
- A 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 EMC and VMware, so your company can leverage existing VMware skillsets

Solution components

The retail edge solution architecture includes the following VMware components. The versions of VMware components are described in the [Technical specifications](#) section of this document.

Solution components
Dell EMC vSAN Ready Node
VMware Edge Compute Stack <ul style="list-style-type: none"> • VMware vCenter Server® Standard • VMware vSphere® Hypervisor (ESXi) Enterprise Plus • VMware vSAN™ Advanced • VMware Tanzu™ Standard
VMware SD-WAN

Dell EMC vSAN Ready Node

Dell EMC PowerEdge servers are designed to provide the highest performance for a diverse set of workloads. PowerEdge servers configured for vSAN (Dell EMC 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 EMC vSAN Ready Node portfolio is one of the broadest in the industry, offering 1 to 4 sockets, in 1U/2U rack, tower and modular solutions for data center, cloud and edge workloads. For more detail, download the [Dell EMC vSAN Ready Node solution brief](#).

An example of this portfolio specifically designed for edge computing is Dell EMC vSAN Ready Node XE2420. The XE2420 is a ruggedized dual-socket, 2U, short-depth, front-accessible server that is designed to support demanding edge applications such as streaming retail analytics, manufacturing logistics, 5G cell processing applications, etc. Edge and non-data center environments can place great environmental stresses on the electrical and physical operation of the servers. The XE2420 is purpose-built to operate under extended system operating temperatures from 5° C to 40° C (41° F to 104° 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 multi-cloud 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 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, informally known as TKG, is a multi-cloud 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, please refer to [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, 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 EMC vSAN Ready Node XE2420 and VMware SD-WAN to deliver a comprehensive solution that enables a reliable, secure and ensure a higher level of application performance across the data center to edge in remote, distributed locations.

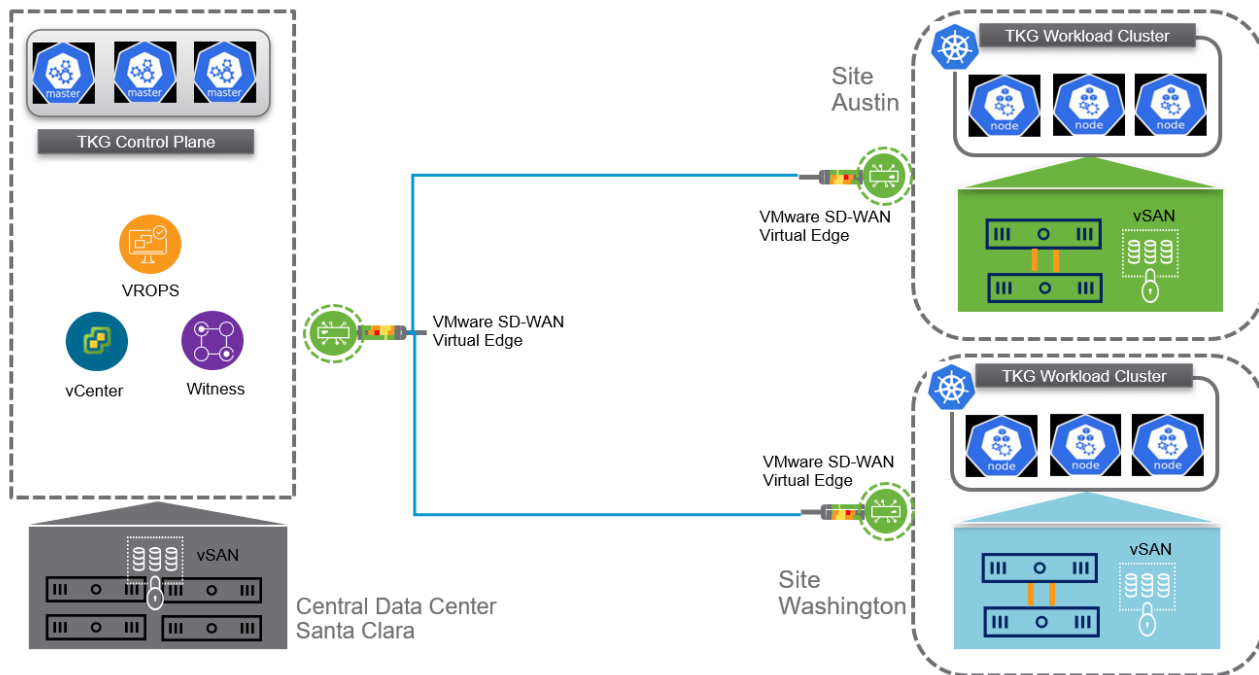


Figure 2: Retail edge solution architecture

The retail edge is a 2-node cluster with two directly connected the XE2420 nodes, and a shared VMware vSAN witness host along with a VMware SD-WAN Virtual Edge. The 2-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 XE2420s 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, to the data center or to the cloud. VMware SD-WAN provides enhancements to an edge deployment that is 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 XE2420 infrastructure at the edge with VMware SD-WAN as network connectivity.

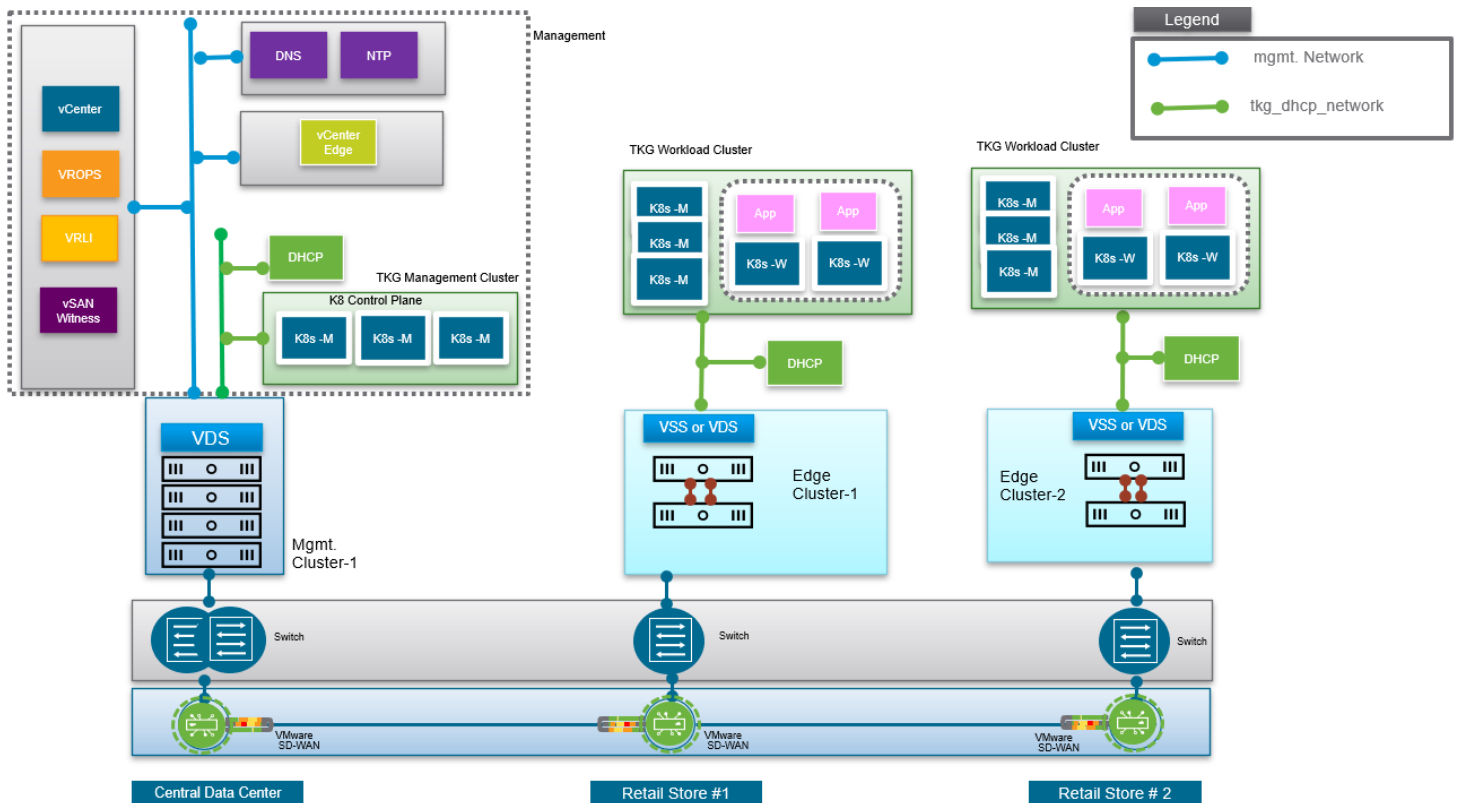


Figure 3: Architecture from the central data center to the edge

- In this deployment, the central data center is hosting all management components including DNS, VMware vCenter Server™, vSAN witness appliance, VMware vRealize Operations Manager™ and vRealize Log Insight.
- 2-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 via VMware SD-WAN L3 routed network.
- vSAN witness appliance is hosted on the management cluster at central data center and witness traffic is connected with an L3 routed network over VMware SD-WAN.
- Kubernetes control plane nodes are hosted on the management cluster at central data center and worker nodes are hosted on 2 node retail edge cluster for retail applications.
- Dedicated local DHCP server for Kubernetes clusters.

- A dedicated edge vCenter is hosted on the management cluster at the central data center and manages all the 2-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.

Notes:

1. vRealize Operations and vRealize Log Insight are optional components.
2. If you would like to add more nodes, then you have to add them in groups of two in a cluster. Once you reach three nodes, then you may remove the vSAN witness as third node.

Virtual infrastructure design

vCenter Server Design

The vCenter Server design includes the design for all the 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 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 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	
Design decision	Design justification
<ul style="list-style-type: none"> • Deploy two vCenter Server systems • One vCenter Server supports the management workloads. • Another vCenter Server supports the compute workloads at the edge 	<ul style="list-style-type: none"> • Isolates vCenter Server failures to management or compute workloads. • Isolates vCenter Server operations between management and compute workloads. • 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	
Design decision	Design justification
Ensure that all ESXi hosts with XE2420 have a uniform configuration across the retail edge clusters.	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 2 physical NICs for management Other two physical NICs will be dedicated for 2-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. Please 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	<ul style="list-style-type: none"> • A good starting point for most workloads • Allows for ESXi and other management overhead

Note: Please refer to the [Technical specifications](#) section of this document for the hardware details.

vSAN witness design

VMware vSAN 2-node architecture is a perfect solution for organizations that have many small branch offices or retail sites. It is also very 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 amount of physical resources needed at the central data center, resulting in a greater level of savings for a large number of 2-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, you can configure a dedicated VMkernel network adapter for vSAN witness traffic as shown in the network design section in 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.

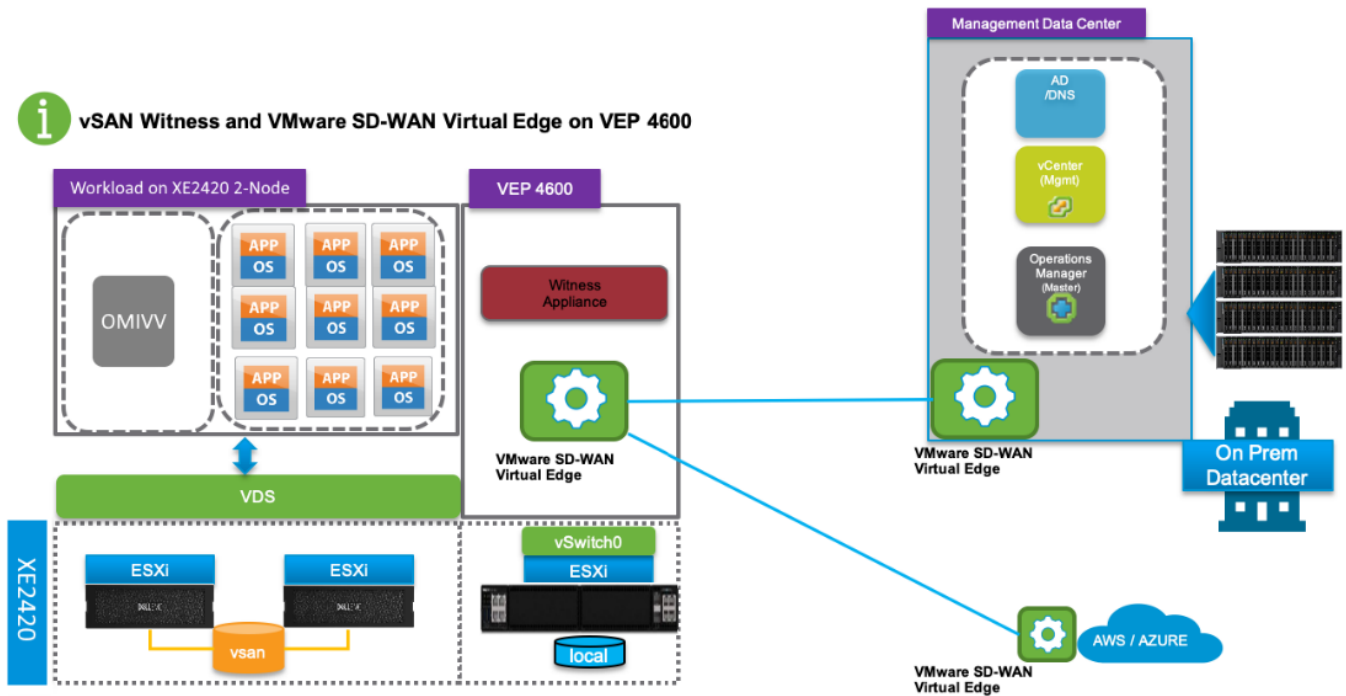


Figure 4: Local vSAN witness design

For a more detailed overview of this design, please refer to [PowerEdge XE2420 2-Node vSAN Edge-Ready HCI](#) white paper.

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 2-node clusters, supporting up to 64,000 components, and requires at least 6 CPUs and 32 GB memory allocation for the witness.

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

Some of the supported configurations of VMware maximums are:

Supported configurations of VMware maximums	
Type	Maximum
Hosts per vCenter	2500
VDS per vCenter	128
Hosts per VDS	2000
1 x vSAN witness (shared)	64 x 2-node clusters

Below is a high-level example of scaling requirements based on VMware maximums. Please discuss with your VMware account representative before you consider the scale and sizing of your environment.

High-level example of scaling requirements based on VMware maximums	
Number of sites	You need
1,000 retail edge sites	1 x Management vCenter 1 x Edge vCenter 16 x vSAN witness appliance ** Each witness appliance can support 64 2-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 upgrade of 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 uploading updates to the depot manually.

vSphere Lifecycle Manager remediates hosts that are part of a vSAN cluster sequentially. The reason is that by design only one host from a vSAN cluster can be in a 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 vSphere Lifecycle Manager Privileges for Using Images.

For more details, please refer to 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 OpenManage industry leading portfolio is integrated Dell EMC Remote Access Controller (iDRAC).

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

- Physical and virtual, monitoring and alerts views in one place
- 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, please visit [OpenManage Integrations and Connections](#) and refer to [Principled Technologies Study, August 2020](#).

Network design

The following figure shows the network design of a central data center and edge site with 2-node XE2420, vSAN witness and VMware SD-WAN.

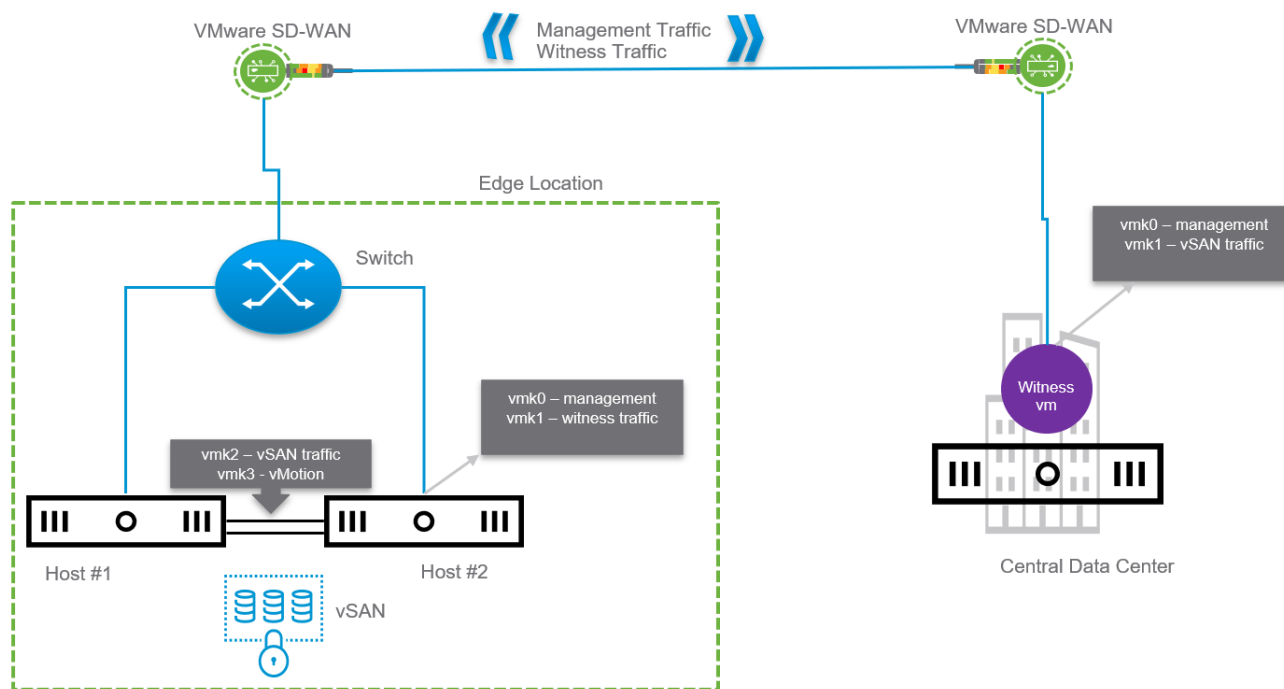


Figure 5: Retail edge network connectivity diagram

In the figure above, the configuration is as follows:

Retail edge cluster: Host #1 and Host #2	
VMkernel network	Traffic type
vmk0	Tagged for management traffic
vmk1	Tagged for witness traffic (**)
vmk2	Tagged for vSAN traffic
vmk3	Tagged for vMotion traffic

Note: (**) This must be done using `esxcli vsan network ip add -i vmk1 -T=witness` on each host.

vSAN witness appliance	
VMkernel network	Traffic type
vmk0	Tagged for management traffic
vmk1	Tagged for vSAN traffic

The figure below describes in detail how each physical NIC is mapped to each virtual switch in the ESXi host and each VMkernel network.

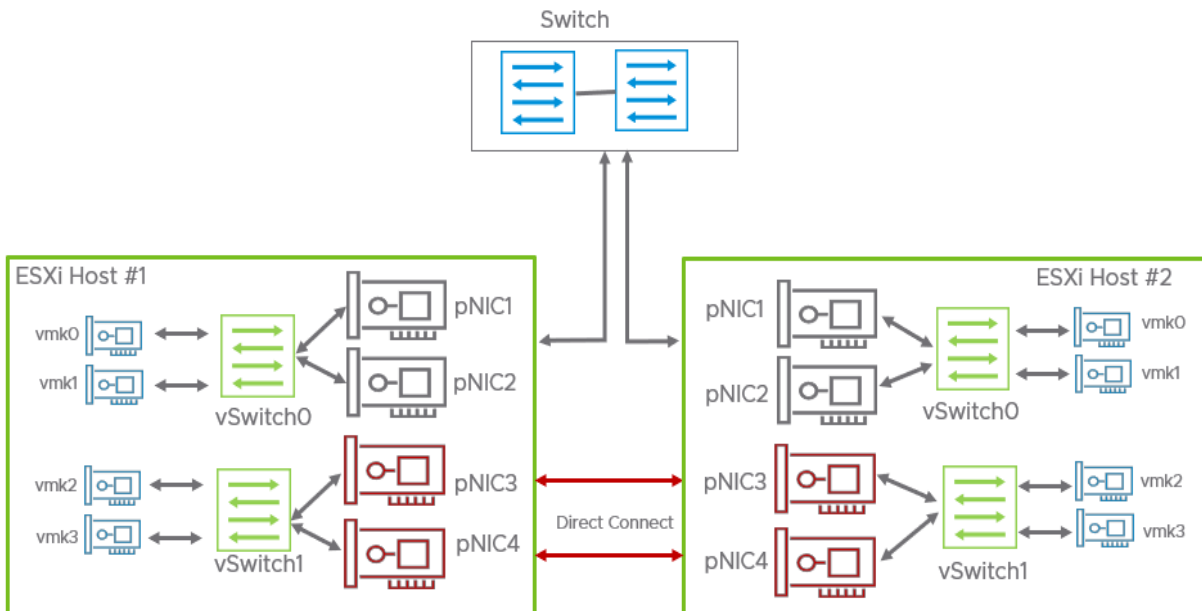


Figure:6 Physical NIC and virtual switch mapping

- 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
- 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 VDS can support only 128 sites per vCenter.

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

For retail edge we have used dedicated VLAN for the following:

- Management
- Uplink
- Virtual machine

For VMware vSAN witness, the following additional VLAN and VM port groups are configured on 2-node cluster:

- Witness

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; more resources can 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, Kubernetes control plane node can be placed on a vSphere cluster at central data center and worker nodes can be placed on a vSphere edge cluster at the edge to support application workloads. vSphere cluster and vSphere host can be managed by edge vCenter which is hosted in the central data center. Please see the figure [Architecture from the central data center to the edge](#) in this document for the architecture.

The following table describes design decisions that will need to be made when implementing Tanzu Kubernetes clusters.

Implementing Tanzu Kubernetes clusters	
Design decision	Design justification
Map the Tanzu Kubernetes clusters to the vSphere Resource Pool in the edge vCenter	Enables resource guarantee and resource isolation
<ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • 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. Please ensure size of disk is considered based on retail workloads.
Set up each ESXi host in the retail edge location with a recommended minimum of 192 GB RAM	<ul style="list-style-type: none"> • Provides connectivity to the vSphere infrastructure • Simplifies network design and reduces 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 cluster's API server. For an application ingress, you can 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 standard Kubernetes configurations within to route requests to the correct application. Please refer to [VMware Tanzu Kubernetes Grid Document](#) 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 (e.g., CSI, CNI, etc.) to the edge sites. That way, clusters can be created in limited network bandwidth and latency.

Please engage your VMware Tanzu sales solution architect to implement the above design principles.

Technical specifications

Hardware

The following table shows the technical specifications of this solution validation.

Hardware technical specifications	
Specification	Quantity
Dell EMC vSAN Ready Node XE2420 Intel(R) Xeon(R) Silver 4216 CPU @ 2.10 GHz, 16 core, 1 socket 256 GB RAM 2 x 900 GB capacity HDD 1 x 400 GB cache HDD	2
VMware SD-WAN	1
10 GbE network switch	1

Notes:

1. For this solution validation, we have used VMware virtual SD-WAN solution. However, it is recommended to use a dedicated VMware SD-WAN appliance for each store for network reliability and resiliency.
2. 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 2-node cluster.

Hardware technical specifications		
Components	Version	Build
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

Conclusion

This solution architecture describes a recommended configuration of Dell EMC 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 continues to provide consistent experience of ease of deployment and management. It enables an efficient centralized management with reduced hardware and software costs, while meeting the needs of environments with limited space, budget, and IT personnel constraints.

Additional Resources/References

- [vSAN 2-node Guide](#)
- [VMware Tanzu at the Edge Solution Architecture](#)

Author

Shree Das is a Technologist in the VMware OCTO business unit. Responsible for architecture and design of various customer reference architecture solutions. Shree's focus is on the VMware Cloud, Software-Defined Data Center, Hyperconverged Platform, Edge Computing, and other emerging technologies such as HPC and AI/ML.

Acknowledgements

The author would like to thank Muneyb Minhazuddin, Marilyn Basanta, Shefali Chinni, Jerry Haskins, Damian Karlson, Raymond Ruan, Anthony Dukes, Mike Monthei, Marco Nicosia, Kevin Willey, Christopher Laffey, Max Fishman, Linka Biaggi, William Leslie, and Inigo Olcoz for their inputs, review and feedback to make this project successful.

