



# 16<sup>th</sup> Generation Dell PowerEdge servers + VMware vSphere 8.0: Improve workload resilience & streamline management of modern apps

The combination of the latest hardware and software platforms can give DevOps and infrastructure admin teams better reliability and more provisioning options for containers—with simple and efficient management

## Introduction

Developer and operations teams that adopt DevOps philosophies and methodologies often rely on container technologies for the flexibility, efficiency, and consistency containers provide. Compared to running multiple VMs, which can be more resource-intensive, containers require less overhead and offer greater portability.<sup>1</sup> When teams need to spin up their workloads, containers are a convenient way for infrastructure admins to provide those resources and a speedy way for developers to test applications. These teams also need orchestration and management software to manage those containerized workloads. Kubernetes® is one such orchestration platform. It helps admins fully manage, automate, deploy, and scale containerized applications and workloads across an environment.<sup>2</sup>

VMware® developed vSphere® with Tanzu to integrate Kubernetes containers with vSphere and to provide self-service environments for DevOps applications. With vSphere 8.0, the latest version of the platform, VMware has further enhanced vSphere with Tanzu.

Using an environment that included the latest 16<sup>th</sup> Generation Dell™ PowerEdge™ servers, we explored three new vSphere with Tanzu features that VMware released with vSphere 8.0:

1. the ability to span clusters with Workload Availability Zones to increase resiliency and availability
2. the introduction of the ClusterClass definition in the Cluster API, which helps streamline Kubernetes cluster creation
3. the ability to create custom node images using vSphere Tanzu Kubernetes Grid (TKG) Image Builder.

We found that these features could let DevOps teams save time while also increasing availability, helping pave the way for more efficient application development processes.

## Understanding containers and their use in DevOps

Containers are packages of software that contain all the elements to run in a computing environment. Those elements include applications and services, libraries, and configuration files, among others. Containers do not typically contain operating system images but rather share elements of operating systems on the device—in our test cases, VMs running on PowerEdge servers. Using containers can benefit developers and DevOps applications because they are more lightweight and portable than single-function VMs.

Additionally, containers offer quick, reliable, and consistent deployments across different environments because they run as resource-isolated processes. They can run in a private data center, in the public cloud, or on a developer's personal laptop.

An organizational advantage of containers is that they can help teams focus on their tasks: Developers can focus on application development and code issues, and operations teams can focus on the physical and virtual infrastructure. Some organizations may choose to combine developers and operations teams in a single DevOps group that manages the entire application development lifecycle, from planning to deployment and operations. In other organizations, separate development and operational teams may choose to adopt DevOps methodologies to streamline processes. For our test purposes, we assumed that a DevOps team is responsible for provisioning and managing virtual infrastructure for developers but not necessarily using that infrastructure themselves.

### 16<sup>th</sup> Generation Dell PowerEdge servers

The latest line of servers from Dell, PowerEdge servers come in multiple form factors for workload acceleration. Also in 16<sup>th</sup> Generation Dell PowerEdge servers is the Dell PowerEdge RAID controller 12 (PERC 12 Series), which offers expanded support and capabilities compared to previous versions, including support for SAS4, SATA, and NVMe<sup>®</sup> drive types.

Dell PowerEdge servers, including those from the 16<sup>th</sup> Generation, offer resiliency in the form of Fault Resistant Memory (FRM).<sup>3</sup> FRM creates a fault-resilient area of memory that protects the hypervisor against uncorrectable memory errors, which in turn helps safeguard the system from unplanned downtime. On the virtual side, VMware vSAN™ 8.0 allows DevOps and infrastructure teams to pool storage to create smaller fault domains.

To learn more about the latest 16<sup>th</sup> Generation Dell PowerEdge servers, visit <https://www.dell.com/en-us/shop/dell-poweredge-servers/sc/servers>.

To see more of our recent work with the latest Dell PowerEdge servers, visit <https://www.principledtechnologies.com/portfolio-marketing/Dell/2023>.

## A quick look at vSphere with Tanzu

VMware created the Tanzu Kubernetes Grid (TKG) framework to simplify the creation of a Kubernetes environment. VMware includes it as a central component in many of their Tanzu offerings.<sup>4</sup> vSphere with Tanzu takes TKG one step further by integrating the Tanzu Kubernetes environments with vSphere clusters. The addition of vSphere Namespaces (collections of CPU, memory, and storage quotas for running TKG clusters) allows vSphere admins to allocate resources to each Namespace to help balance developer requests with resource realities.<sup>5</sup>

vSphere admins can use vSphere with Tanzu to create environments in which developers can quickly and easily request and access the containers they need to develop and maintain workloads. At the same time, vSphere with Tanzu allows admins to monitor and manage resource consumption using familiar VMware tools. (To learn more about how Dell and VMware solutions can enable a self-service DevOps environment, read our report “Give DevOps teams self-service resource pools within your private infrastructure with Dell Technologies APEX cloud and storage solutions” at <https://facts.pt/MO2uvKh.>)

## New features with VMware vSphere 8.0

As we previously noted, VMware vSphere 8.0 brings new and improved features to vSphere with Tanzu. These include the ability to deploy vSphere with Tanzu Supervisor in a zonal configuration, the inclusion of ClusterClass specification in the TKG API, and tools that let administrators customize OS images for TKG cluster worker nodes. In this section, we describe these features and how they function; later, we cover how we tested them and what our results might mean to your organization.

### Deploying vSphere with Tanzu Supervisor in a zonal configuration provides high availability

In previous versions of vSphere, administrators could deploy vSphere with Tanzu Supervisor to only a single vSphere host cluster. vSphere 8.0 introduces the ability to deploy the Supervisor in a configuration spanning three Workload Availability Zones, which map directly to three respective clusters of physical hosts. This zonal deployment of the Supervisor provides high availability at the vSphere cluster layer and keeps Tanzu Kubernetes workloads running even if one cluster or availability zone is temporarily unavailable. Figure 1 shows how we set up our environment to test the functionality of Workload Availability Zones.

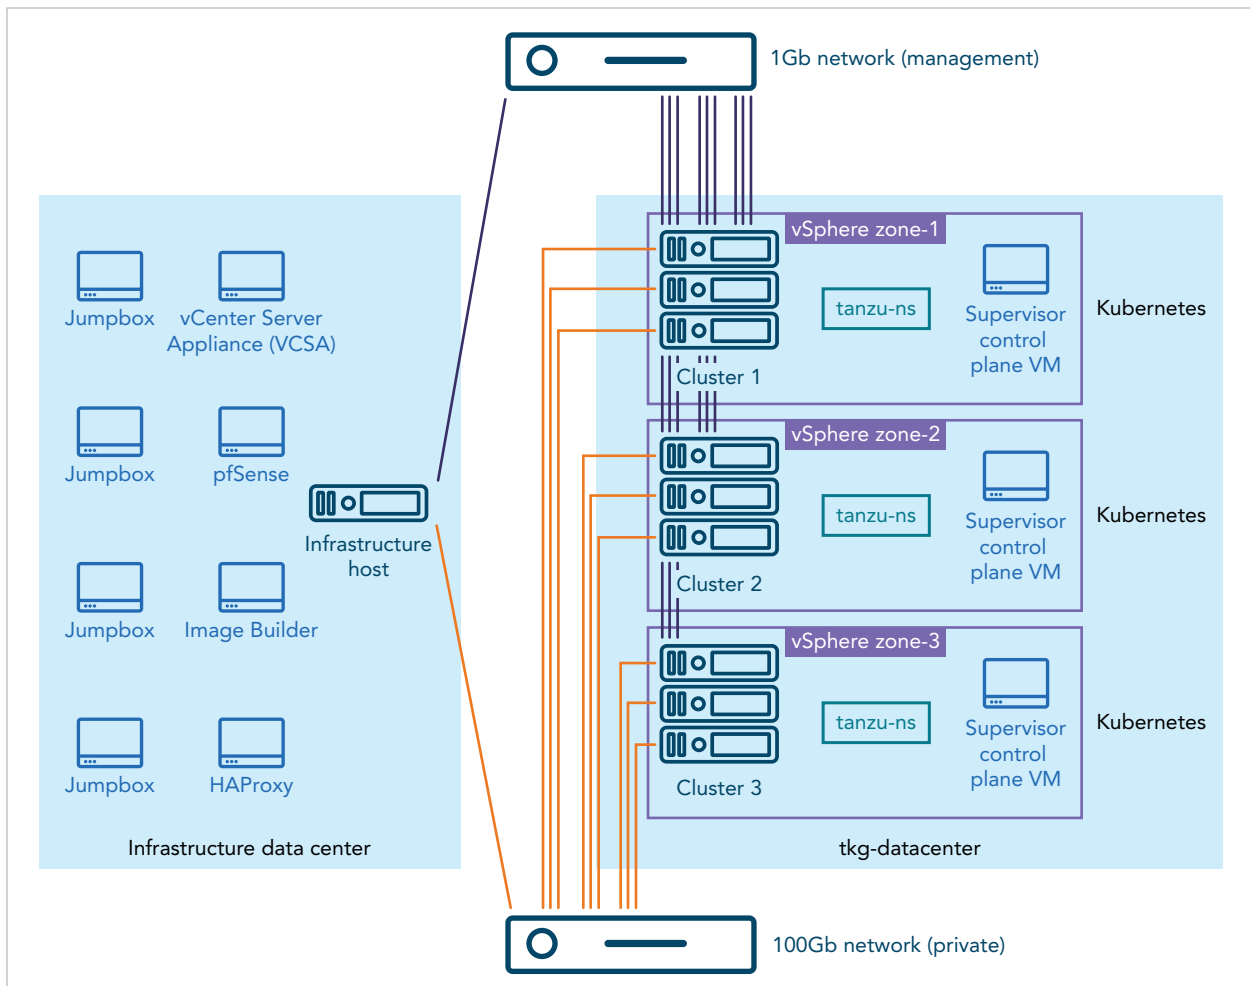


Figure 1: The zonal setup of our testbed. Source: Principled Technologies.

The vSphere with Tanzu Supervisor is a vSphere-managed Kubernetes control plane that operates on the VMware ESXi™ hypervisor layer and functionally replaces a standalone TKG management cluster. vSphere administrators can configure and deploy the Supervisor via the Workload Management tool in VMware vCenter Server®. Once an administrator has deployed the Supervisor cluster and associated it with a vSphere Namespace, the cluster acts as a control and management plane for the Namespace and all the TKG workload clusters it contains.

The Namespace provides role-based access control (RBAC) to the virtual infrastructure and defines the hardware resources—CPU, memory, and storage—available to the Supervisor for provisioning workload clusters. It also allows infrastructure administrators to specify the available VM classes, storage policies, and Tanzu Kubernetes Releases (TKRs) available to users with access to the Namespace. Users can then provision their workload clusters as necessary to host Kubernetes applications and services in a way that the hypervisor and vSphere administrators can see.

Based on our testing, having the Supervisor span three Workload Availability Zones provided increased resiliency to our Tanzu Kubernetes workload clusters. We tested by deploying a webserver application to a workload cluster spanning all three availability zones. Even after we powered off one of the physical

hosts that comprised the three Workload Availability Zones physically, the application remained functional and accessible. After we powered on the hosts, the Supervisor repaired itself and brought all missing control plane and worker nodes back online automatically without administrator intervention.

## Defining clusters with the ClusterClass specification in the TKG API is easier

Another significant improvement to vSphere with Tanzu in vSphere 8.0 is the introduction of the ClusterClass resource to the TKG API. vSphere with Tanzu in vSphere 8.0 uses TKG 2, which implements the ClusterClass abstraction from the Kubernetes Cluster API and provides a modular, cloud-agnostic format for configuring and provisioning clusters and cluster lifecycle management.

Previous versions of vSphere with Tanzu used the TKG 1.x API. This API used a VMware proprietary cluster definition called tanzukubernetescluster (TKC) that was specific to vSphere. With the introduction of ClusterClass to the v1beta1 version of the open-source Kubernetes Cluster API and its implementation in the TKG 2.x API, cluster provisioning and management in vSphere with Tanzu are now more directly aligned with the broader multi-cloud Kubernetes ecosystem and provide a more portable and modular way to define clusters regardless of the specific cloud environment. The ClusterClass specification also allows users to define which Workload Availability Zones any specific workload cluster will span and is less complicated than the previous TKC cluster specification.

## Administrators can customize OS images for TKG cluster worker nodes by using new tools

vSphere 8.0 also provides tools for administrators to build their own OS images for TKG cluster worker nodes. vSphere with Tanzu includes a VMware-curated set of worker node images for specific combinations of Kubernetes releases and Linux OSes, but now administrators can create custom OS images and make them available to DevOps users via a local content library for use in provisioning workload clusters. This allows administrators to create custom OS images outside of the options distributed with TKG, including the ability to create worker nodes based on Red Hat® Enterprise Linux® 8 that TKG supports but does not currently include in the TKG distribution set of default images. To suit the needs of their DevOps users, administrators can also use this feature to preinstall binaries and packages to worker node images that the OS does not include by default.

## What can the combination of 16<sup>th</sup> Generation Dell PowerEdge servers and VMware vSphere 8.0 deliver for containerized environments?

A multi-node Dell PowerEdge server cluster running VMware vSphere 8.0 could improve workload resiliency and streamline management and security of modern apps. The following sections delve into our testing and explore how the new vSphere 8.0 features could benefit developers and infrastructure teams.

## About VMware vSphere 8.0

vSphere is an enterprise compute virtualization program that aims to bring “the benefits of cloud to on-premises workloads” by combining “industry-leading cloud infrastructure technology with data processing unit (DPU)- and GPU-based acceleration to boost workload performance.”<sup>6</sup>

This latest version introduces the vSphere Distributed Services Engine, which enables organizations to distribute infrastructure services across compute resources available to the VMware ESXi host and offload networking functions to the DPU.<sup>7</sup>

In addition to the improvements for vSphere for Tanzu and the introduction of vSphere Distributed Services Engine, the latest version of the vSphere platform offers new features for vSphere Lifecycle Manager, artificial intelligence and machine learning hardware and software, and other facets of data center operations.

To learn more about vSphere 8.0, visit <https://core.vmware.com/vmware-vsphere-8>.

## How we tested

### About the Dell PowerEdge R6625 servers we used

The PowerEdge R6625 rack server is a high-density, two-socket 1U server. According to Dell, the company designed the server for “data speed” and “HPC workloads or running multiple VDI instances,”<sup>8</sup> and to serve “as the backbone of your data center.”<sup>9</sup> Two AMD EPYC™ 4<sup>th</sup> Generation 9004 Series processors power these servers, delivering up to 128 cores per processor. For storage, customers can get up to 153.6 TB in the front bays from 10 2.5-inch SAS, SATA, or NVMe HDD or SSD drives, and up to 30.72 TB in the rear bays from two 2.5-inch SAS or SATA HDD or SSD drives. For memory, servers have 24 DDR5 DIMM slots to support up to 6 TB of RDIMM at speeds up to 4800 MT/s. Customers can also choose from many power and cooling options. The server and processor offer the following security features:<sup>10</sup>

- AMD Secure Encrypted Virtualization (SEV)
- AMD Secure Memory Encryption (SME)
- Cryptographically signed firmware
- Data at Rest Encryption (SEDs with local or external key management)
- Secure Boot
- Secured Component Verification (hardware integrity check)
- Secure Erase
- Silicon Root of Trust
- System Lockdown (requires iDRAC9 Enterprise or Datacenter)
- TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China Nation

Table 1 shows the technical specifications of the six 16<sup>th</sup> Generation PowerEdge R6625 servers we tested. Three of the servers had one configuration, and the other three had a different one. We clustered the similarly configured servers together.

Table 1: Specifications of the 16<sup>th</sup> Generation Dell PowerEdge R6625 servers we tested. Source: Principled Technologies.

<b>Cluster 1</b>	
Hardware	3x Dell PowerEdge R6625
Processor	2x AMD EPYC 9554 64-core processor, 128 cores, 3.1 GHz
Disks (vSAN)	4x 3.2TB Dell Ent NVMe v2 AGN MU U.2
Disks (OS)	2x 600GB Seagate ST600MM0069
Total memory in system (GB)	128
Operating system name and version/build number	VMware ESXi, 8.0.1, 21495797
<b>Cluster 2</b>	
Hardware	3x Dell PowerEdge R6625
Processor	2x AMD EPYC 9554 64-core processor, 128 cores, 3.1 GHz
Disks (vSAN)	4x 6.4TB Dell Ent NVMe v2 AGN MU U.2
Disks (OS)	2x servers: 2x 600GB Seagate ST600MM0069 1x server: 2x 600GB Toshiba AL15SEB060NY
Total memory in system (GB)	128
Operating system name and version/build number	VMware ESXi, 8.0.1, 21495797

In addition to the 16<sup>th</sup> Generation Dell PowerEdge R6625 servers, we included previous-generation PowerEdge R7525 servers in our testbed (see Table 2). Our use of the older servers in testing reflects actual use cases where organizations might include them for less-critical applications or as failover options.

Table 2: Specifications of the Dell PowerEdge R7525 servers we tested. Source: Principled Technologies.

<b>Cluster 3</b>	
Hardware	3x Dell PowerEdge R7525
Processor	2x AMD EPYC 7513 32-Core Processor, 64 cores, 2.6 GHz
Disks (vSAN)	4x 3.2TB Dell Ent NVMe v2 AGN MU U.2
Disks (OS)	2x servers: 2x 240GB Micron MTFD DAV240TDU 1x server: 2x 240GB Intel SSD SCKKB240G8R
Total memory in system (GB)	512
Operating system name and version/build number	VMware ESXi, 8.0.1, 21495797

For more information on the servers we used in testing, see the science behind this report.

To learn more about the Dell PowerEdge R6625, visit [https://www.dell.com/en-us/shop/cty/pdp/spd/poweredge-r6625/pe\\_r6625\\_tm\\_vi\\_vp\\_sb](https://www.dell.com/en-us/shop/cty/pdp/spd/poweredge-r6625/pe_r6625_tm_vi_vp_sb).

## Setting up the VMware vCenter environment

We started with six Dell PowerEdge R6625 and three Dell PowerEdge R7525 servers and installed the VMware vSphere 8.0.1 release of the Dell-customized VMware ESXi image on each server. To support our environment, we deployed the latest version of VMware vCenter on a separate server. We then created three separate vSAN-enabled vSphere clusters, as we defined in Table 1, and deployed vSphere with Tanzu with a zonal Supervisor cluster spanning the three vSphere clusters, each set as a vSphere Zone. Figure 2 shows the details of our testbed.

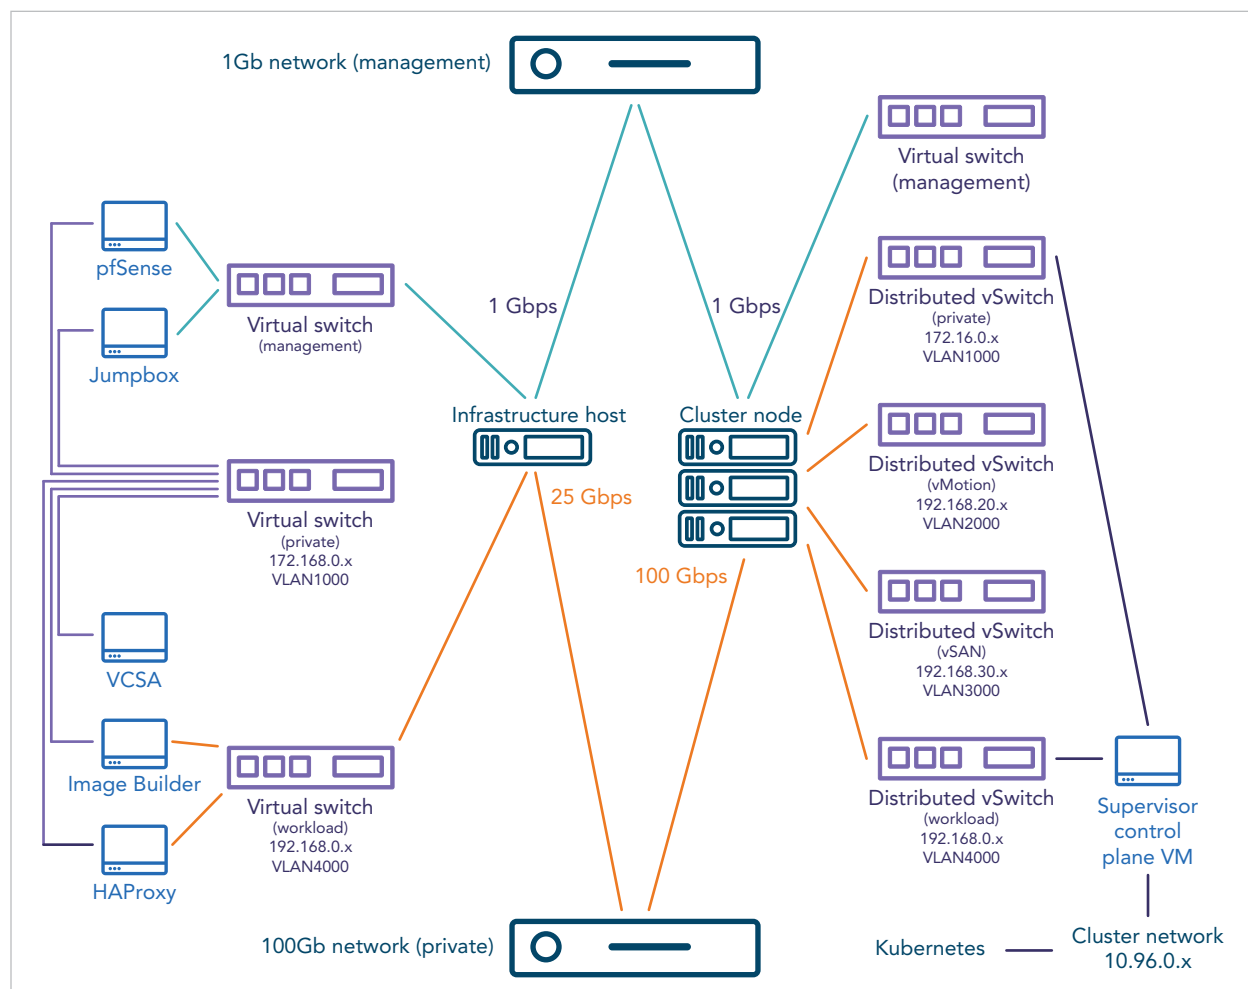


Figure 2: Our testbed in detail. Source: Principled Technologies.



## Testing Workload Availability Zones

To demonstrate the utility of a zonal Supervisor deployment spanning three Workload Availability Zones, we deployed a simple “hello world” NGINX web server application to a workload cluster that spanned all three zones. Our deployment defined three replicas of the application, which the Supervisor distributed across the three zones automatically. Then, to simulate a system failure—such as might occur with a power outage—we powered off the physical hosts of one vSphere cluster mapped as an availability zone.

Even after this loss of power on the cluster, the workload continued to function. The missing Supervisor and worker nodes and their services automatically came back online after we powered on the hosts and they reconnected to the vCenter Server.

For clusters hosting any kind of key workload, whether user-facing or internal, availability is critical. If a cluster goes down completely—whether due to natural disaster, human error, or malicious intent—the resulting downtime can have severe consequences, from hurting productivity to losing revenue. As an example for our use case, downtime for containerized microservices that are part of an ecommerce app could lead to miscalculations in inventory or the inability to process an order, potentially causing customer dissatisfaction and lost revenue. Workload Availability Zones offer a resiliency that can ease the minds of both development and infrastructure teams by ensuring that containerized applications and services have the uptime to meet usage demands.

## Testing ClusterClass

We used the v1beta1 Cluster API implementation of ClusterClass to deploy workload clusters using TKG 2.0 on a multi-zone vSphere with Tanzu Supervisor. The Cluster API, implemented in the vSphere with Tanzu environment as part of TKG, provides provisioning and operating of Kubernetes clusters across different types of infrastructure, including vSphere, Amazon Web Services, and Microsoft Azure, among others. We used YAML configuration files in our test environment to define and test two sample ClusterClass cluster definitions representing different options available to users when provisioning a workload cluster deployment in vSphere with Tanzu: One spanning all three Workload Availability Zones, and one that uses the Antrea Container Network Interface (CNI) rather than the default Calico CNI.

The implementation of the ClusterClass object in the TKG 2 API replaces the previous `tanzukubernetescluster` resource used in TKG 1.x and brings Tanzu cluster definitions into a more modular and cloud-agnostic format that aligns with user expectations and helps with portability across diverse cloud environments. Although we did not test the full extent of the feature, we found that it worked well in our testing.

## Testing node image customization on the workload cluster

Using the VMware Tanzu Image Builder, we created a custom Ubuntu 20.04 image for use with Kubernetes v.1.24.9 and vSphere with Tanzu. We added a simple utility package (curl) that the base Ubuntu installation does not include by default. To confirm that vSphere used the custom OS image on the provisioned cluster's nodes as we expected, we connected to one of the worker nodes via SSH. When we did so, we saw that the platform had installed curl—the feature had worked as intended.

This feature allows teams to build OS images to meet the specific needs of their applications or services. They can then add the custom-built OS images as a custom TKr in vSphere for Tanzu to provision workload clusters.

DevOps teams can create custom machine images for the management and workload cluster nodes to use as VM templates. Each custom machine image packages a base operating system, a Kubernetes version, and any additional customizations into an image that runs on vSphere. However, custom images must use the operating systems that vSphere for Tanzu supports; currently, it supports Ubuntu 20.04, Ubuntu 18.04, Red Hat Enterprise Linux 8, Photon OS 3, and Windows 2019.<sup>11</sup>

This kind of customization could save time for DevOps teams and help maintain Tanzu environmental consistency. It also allows users to predefine binaries included on an OS image rather than install them afterward.

## Conclusion

The combination of 16<sup>th</sup> Generation Dell PowerEdge servers and VMware vSphere 8.0 provides a resilient and efficient solution for organizations running containerized applications. Upgrading to the hardware and software combination provides helpful features that increase workload availability and streamline management. We found Workload Availability Zones, ClusterClass, and node image customization features to function well and as we expected, thus aiding both development and infrastructure teams in a DevOps context. Deploying vSphere 8.0 on Dell PowerEdge servers enables organizations to take advantage of the benefits of containerization and modernize their IT infrastructure.

1. Rosbach, Felix, "Why we use Containers and Kubernetes - an Overview," accessed July 10, 2023, <https://insights.comforte.com/why-we-use-containers-and-kubernetes-an-overview>.
2. Rosbach, Felix, "Why we use Containers and Kubernetes - an Overview."
3. Dell Technologies, "Dell Fault Resilient Memory," accessed June 28, 2023, <https://dl.dell.com/content/manual19947903-dell-fault-resilient-memory.pdf?language=en-us>.
4. VMware, "VMware Tanzu Kubernetes Grid Documentation," accessed July 7, 2023, <https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/index.html>.
5. VMware, "What is a vSphere Namespace?," accessed July 7, 2023, <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-with-tanzu-concepts-planning/GUID-638682AE-5EC9-47AA-A71C-0BECF9DC27A1.html>.
6. VMware, "VMware vSphere," accessed June 29, 2023, <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vsphere/vmw-vsphere-datasheet.pdf>.
7. VMware, "Introducing VMware vSphere Distributed Services Engine and Networking Acceleration by Using DPUs," accessed June 29, 2023, <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-installation/GUID-EC3CE886-63A9-4FF0-B79F-111BCB61038F.html>.
8. Dell Technologies, "PowerEdge R6625 Rack Server," accessed June 28, 2023, [https://www.dell.com/en-us/shop/cty/pdp/spd/poweredge-r6625/pe\\_r6625\\_tm\\_vi\\_vp\\_sb](https://www.dell.com/en-us/shop/cty/pdp/spd/poweredge-r6625/pe_r6625_tm_vi_vp_sb).
9. Dell Technologies, "PowerEdge R6625: Breakthrough performance," accessed June 28, 2023, <https://www.dell-technologies.com/asset/en-us/products/servers/technical-support/poweredge-r6625-spec-sheet.pdf>.
10. Dell Technologies, "PowerEdge R6625: Breakthrough performance."
11. VMware, "Tanzu Kubernetes Releases and Custom Node Images," accessed July 10, 2023, <https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/2/about-tkg/tkr.html>.

Read the science behind this report at <https://facts.pt/oJcF4hw> ►



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.

This project was commissioned by Dell Technologies.