Virtual Desktop Infrastructure with VMware Horizon on Dell Three-Tier Infrastructure with SmartFabric Services

Design Guide

Abstract

This design guide describes the architecture and design of the Dell Validated Design for Virtual Desktop Infrastructure (VDI) with VMware Horizon brokering software, based on a 3-tier architecture using Dell infrastructure, including Dell PowerEdge servers, Dell storage, and Dell PowerSwitch networking with SmartFabric Services.

Dell Technologies Solutions

Dell
Validated Design

Part Number: H19329
October 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

Topics:

Solution introduction

Overview

Dell Validated Designs for Virtual Desktop Infrastructure (VDI) on 3-tier architectures using Dell PowerEdge servers, Dell storage, and Dell PowerSwitch networking provide a flexible and customized approach to simplify and extend your VMware environment. These solutions are ideal for VDI applications because they combine abstract architectural compute, storage, networking, and virtualization layers and allow for integrated management of the components using OpenManage Systems Management Solutions and SmartFabric Services.

Installing VMware Horizon 8 with its VDI components on 3-tier architectures based on Dell Technologies products and solutions enables organizations to quickly deliver Microsoft Windows virtual desktops or server-based hosted shared sessions on a wide variety of endpoint devices.

Document purpose

This document introduces the architecture, components, design options, best practices, and configuration details for successful VDI deployments for a 3-tier architecture composed of PowerEdge servers, PowerStore storage, PowerSwitch networking, and SmartFabric Services with VMware Horizon 8.

Audience

This document is intended for decision makers, managers, architects, developers, and technical administrators of IT environments who want an in-depth understanding of the value of the Dell Validated Designs for VDI that deliver Microsoft Windows virtual desktops using VMware Horizon 8 VDI components on Dell PowerEdge servers, Dell storage, and Dell PowerSwitch networking.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Solutions team by email.

Authors: Dell Validated Designs for VDI team

What's new?

This solution brings the following new features:

- 4-socket PowerEdge server validation
- SmartFabric Services integration

Design guide introduction

Dell Technologies offers comprehensive, flexible, and efficient VDI solutions that are designed and optimized for your organization's needs. These VDI solutions are easy to plan, deploy, and run. To start or extend your data center, you can start with just one component from each of the server, storage, and networking tier and scale up or out from there.

PowerEdge servers are designed for modern IT environments of any size, anywhere. The PowerEdge portfolio can help you adapt to change quickly, incorporating the latest technology without disruption and empowering your organization from home offices to multiple data centers and clouds to distant edge locations.

Dell storage products are integrated into a mature monitoring and reporting suite that is designed to minimize management effort. With robust reporting and notification services, storage management can be as minimal as needed. The integrated dashboards also enable at-a-glance insight into the health and performance of the platform.

Dell networking products are designed for building high-capacity network fabrics that are easy to deploy, are cost-effective, and provide a clear path to a software-defined data center. The Dell SmartFabric OS10 Enterprise Edition is a network operating system that leverages multilayered disaggregation of the network functionality to support multiple architectures and environments. This allows for freedom and flexibility in selecting networking, monitoring, management, and orchestration applications to meet customers' needs. Together with SmartFabric Services, customers can quickly and easily deploy and automate data center networking fabrics using leaf-spine fabrics.

Dell Validated Designs for VDI offer several key benefits:

- **Predictable costs, performance, and scalability to support a growing workforce**—The cost of deploying applications and desktops can be reduced when deploying from a data center.
- **Rapid deployments**—Dell Validated Designs for VDI offer rapid automated deployment at the infrastructure layer through Dell Technologies Services.
- **Rapid scaling to serve enterprises of any size**—Quickly scale with workload-specific configurations per platform. The solution is scalable up to 96 nodes per cluster, supporting thousands of virtual machines (VMs). The VDI environment can continue scaling with pod architecture and by combining multiple clusters. VxRail Manager makes it easy to insert or remove nodes from your cluster to meet your business needs.
- **Granular infrastructure scaling**—Each layer of the design can be scaled independently using a scale-up or scale-out architecture to allow for a flexibility in initial design and deployment and the accommodation of future growth and changing performance needs.
- **Dell Technologies support**—Dell Validated Designs for VDI are engineered and tested specifically for VDI and related applications. Deployment services can provide installation of this turnkey VDI appliance to ensure a rapid deployment with linear and predictable scalability. Dell Technologies offers single-company-support models for VDI solutions.
- **User experience (UX)**—NVIDIA virtual GPU (vGPU) technology brings the power of NVIDIA GPUs to virtual desktops, apps, and workstations, accelerating graphics and compute to enable virtualized workspaces to perform just like physical PCs. Remote workers in a variety of locations can have a consistently excellent user experience.

2

Solution Architecture

Topics:

Architecture overview

Overview

This section provides an architecture overview and guidance on managing and scaling a VMware Horizon environment on Dell PowerEdge servers, storage, and PowerSwitch networking.

Solution architecture

The following figure depicts an example of a 3-tier architecture. It shows the validated solution, including the compute and graphics, management, storage, and networking layers. This architecture aligns with the VMware Horizon pod and block design. A pod is divided into multiple blocks and each block consists of one or more vSphere clusters and a vCenter Server.

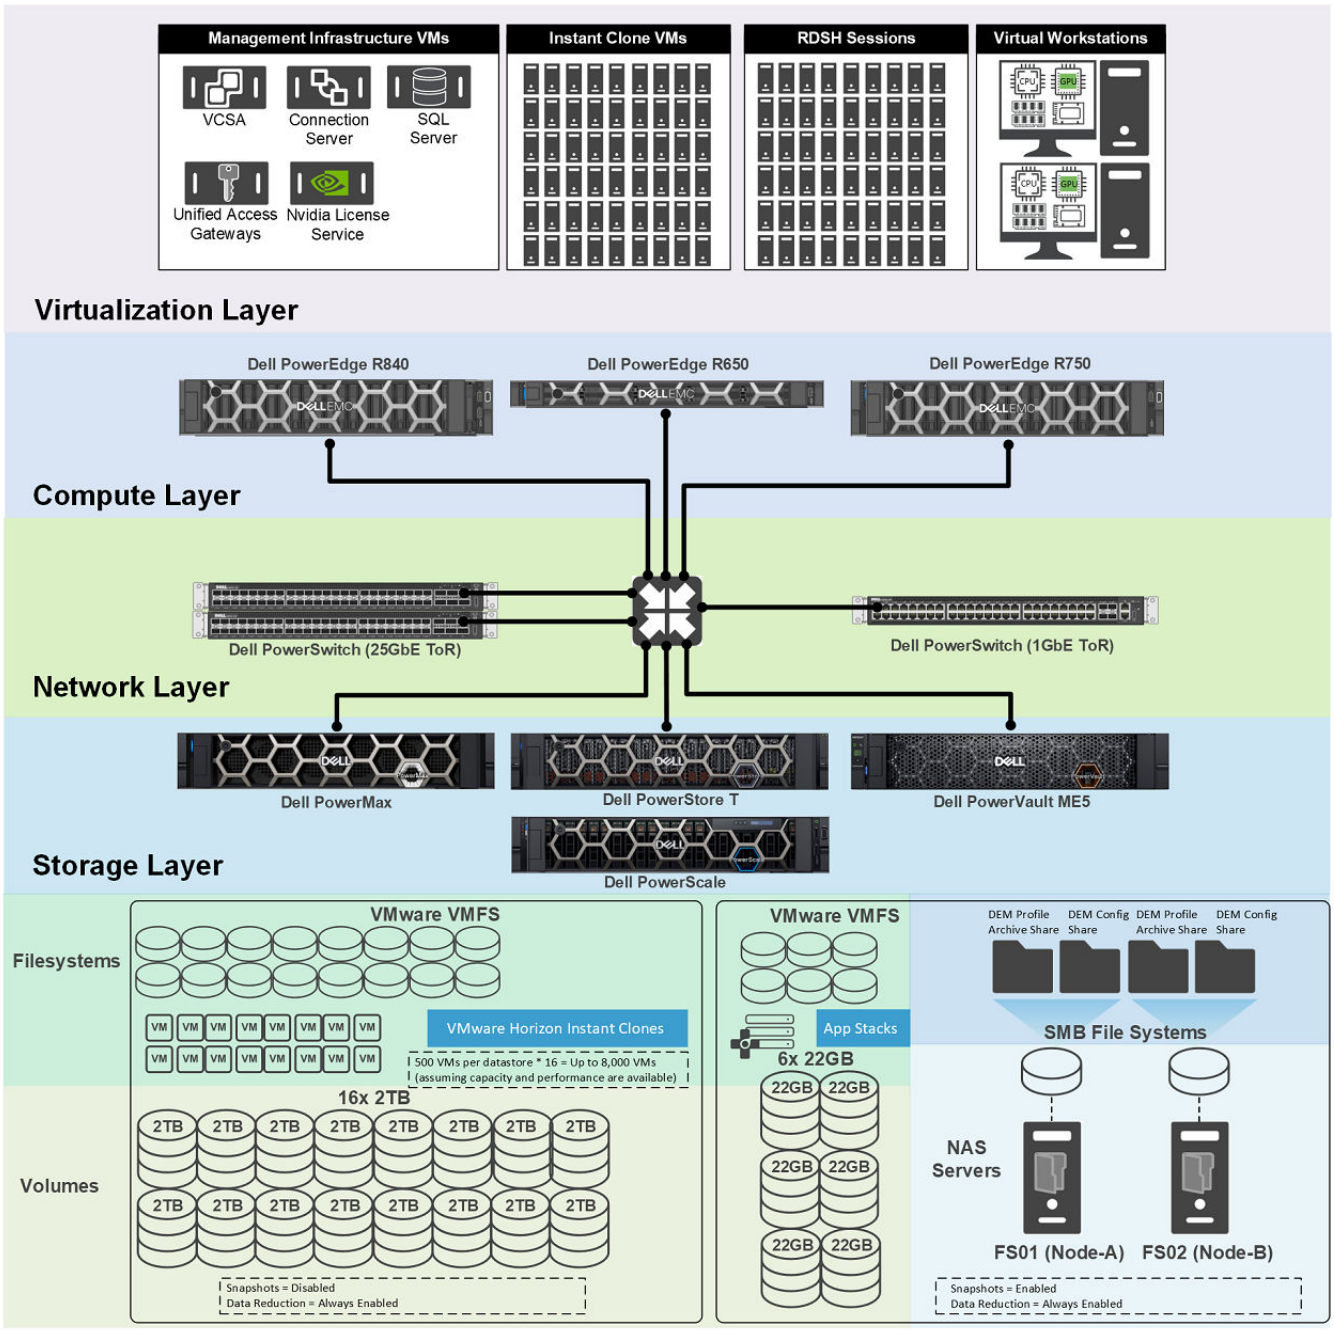

Figure 1. Solution architecture

Physical architecture

Overview

Dell Technologies offers a variety of products and solutions that can be combined into the optimal configuration to support the VDI workload that is required by your business needs. There are various PowerEdge servers, storage, and networking configurations available. This section describes the physical architecture of the PowerEdge R650, R750, and R840 servers, PowerMax, PowerStore, PowerVault, and PowerScale storage, and PowerSwitch networking products which are recommended for VDI deployments and are leveraged for "Management-Optimized," "Density Optimized" and "Virtual Workstation" configurations.

Servers

PowerEdge servers provide the highest performance for a diverse set of VDI workloads from the edge to the cloud to the core. As applications and workplaces become more complex, it is more important than ever for businesses to have end-to-end solutions that work together seamlessly. PowerEdge servers and VMware provide industry-leading solutions for a modern data center, including servers, virtualization, storage, networking, and cloud.

PowerEdge servers deliver lower total cost of ownership (TCO), scalable architectures, intelligent automation and management, and multilayer security when paired with the industry leader in enterprise virtualization software. PowerEdge and VMware help modernize, automate, and protect IT infrastructure.

Dell PowerEdge R650

The dual-socket 1U Dell PowerEdge R650 is the ideal rack server to address performance, high scalability, and density. For those who need to increase efficiency and accelerate operations, this rack server is ideal for data center standardization on a wide range of workloads. This is the recommended server model for deploying VDI management resources on but can also be configured similarly to the density optimized configuration, but with less disk capacity and GPU configurations.

Figure 2. Dell PowerEdge R650

Dell PowerEdge R750

The dual-socket 2U Dell PowerEdge R750 is a full-featured enterprise server that delivers outstanding performance for the most demanding workloads. With powerful 3rd Generation Intel Xeon Scalable processors, this server provides high-density virtualization with VDI and allows for maximum flexibility. This server provides the highest processor speeds and graphics capabilities with up to two NVIDIA A16 or A40 GPUs per node.

Figure 3. Dell PowerEdge R750

Dell PowerEdge R840

The four-socket 2U Dell PowerEdge R840 delivers consistent, high performance results for data-intensive applications and business critical workloads including VDI. With powerful 2nd Generation Intel® Xeon® Scalable processors and up to 112 cores, the R840 can handle even the most compute-intensive VDI workloads to drive your business forward faster. Create an optimal configuration of NVMe, SSD, HDD, and GPU resources to address your most demanding workloads — all in a 2U chassis.

Figure 4. Dell PowerEdge R840

VDI-optimized configurations

We have designated common configurations as Management-Optimized, Density-Optimized, Virtual Workstation, and Compute-Dense. These designations are outlined in the following tables and are referenced throughout the document.

Table 1. Common configurations

Configuration	Platform	CPU	RAM	GPU (optional)	Description
Management-Optimized	PowerEdge R650	1 x Intel Xeon Gold 6330 (28-core @ 2.0 GHz)	256 GB (8 x 32 GB @ 2,933 MHz)	None	Offers density and value to provide a dedicated environment to deploy virtualized management infrastructure.
Density-Optimized	PowerEdge R750	2 x Intel Xeon Gold 6348 (28-core @ 2.6 GHz)	1024 GB (16 x 64 GB @ 3,200 MHz)	Up to 2 x full length, dual width (FLWD) Up to 6 x half length, single width (HLSW)	Offers an abundance of high-performance features and tiered capacity that maximizes user density.
Virtual Workstation	PowerEdge R750	2 x Intel Xeon Gold 6354 (18-core @ 3.0 GHz)	512 GB (16 x 32 GB @ 3,200 MHz)	Up to 2 x FLDW Up to 6 x HLSW	Offers even higher performance at the tradeoff of user density. Typically for ISV or high-end graphics workstations.
4-socket Density-Optimized	PowerEdge R840	4 x Intel Xeon Gold 6248 (20-core @ 2.5 GHz)	1536 GB (24 x 64 GB @ 2,933 MT/s)	None	Offers support for a wide variety of datacenter workloads including VDI, Data Analytics, Databases, Data Warehousing, Business Intelligence, and CRM.

Tiered memory (PMem) configurations

Starting with VMware vSphere 7.0 update 2, Intel Optane Persistent Memory (PMem) can be deployed with 2nd and 3rd Generation Intel Xeon processors to enable memory tiering. When deployed, a subset of traditional DRAM can be combined with PMEM and configured to run in "Memory Mode." Tiered memory configurations have been shown to provide nearly equivalent performance for many VDI workloads to a traditional DRAM configuration while potentially reducing overall cost.

For additional information, see vSphere/vSAN Support for Intel's Optane Persistent Memory (PMem) (67645) and Memory Tiering: A New Approach to Solving Modern Data Challenges. The following table shows tiered memory (PMem) configurations:

Table 2. Tiered memory (PMem) configurations

Configuration	Platform	CPU	RAM	GPU (optional)	Description
Management-Optimized	PowerEdge R650	1 x Intel Xeon Gold 6330 (28-core @ 2.0 GHz)	1024 GB DRAM - (16 x 16 GB @ 2,933 MT/s) PMem 200 – (8 x 128 GB @ 2,933 MT/s)	None	Offers density and value to provide a dedicated environment to deploy virtualized management infrastructure.
Density-Optimized	PowerEdge R750	2 x Intel Xeon Gold 6348 (28-core @ 2.6 GHz)	1024 GB DRAM – (16 x 16 GB @ 3,200 MT/s) PMem 200 – (8 x 128 GB @ 3,200 MT/s)	Up to 2 x full length, dual width (FLWD)[a] Up to 4 x half length, single width (HLSW)	Offers an abundance of high-performance features and tiered capacity that maximizes user density.
Virtual Workstation	PowerEdge R750	2 x Intel Xeon Gold 6354 (18-core @ 3.0 GHz)	1024 GB DRAM - (16 x 16 GB @ 3,200 MT/s)	Up to 2 x FLDW Up to 4 x HLSW	Offers even higher performance at the tradeoff of user density. Typically for ISV or high-end graphics workstations.

Table 2. Tiered memory (PMem) configurations (continued)

Configuration	Platform	CPU	RAM	GPU (optional)	Description
			PMem 200 – (8 x 128 GB @ 3,200 MT/s)		
4-socket Density-Optimized	PowerEdge R840	4 x Intel Xeon Gold 6248 (20-core @ 2.5 GHz)	2048 GB DRAM - (24 x 32 GB @ 2,666 MT/s) PMem 100 - (16 x 128 GB @ 2,666 MT/s)	None	Offers support for a wide variety of datacenter workloads including VDI, Data Analytics, Databases, Data Warehousing, Business Intelligence, and CRM.

a. When using air cooling, FLDW GPUs are supported only on the PowerEdge R750 8x 2.5" NVME, 16x2.5" SAS, and 16x2.5" NVMe chassis with a maximum datacenter ambient temperature of 30°C.

NVIDIA GPUs

You can configure Dell Validated Designs for VDI with the following NVIDIA GPUs:

- **NVIDIA A40**—These GPUs provide an increase in performance and multi-workload capabilities for the data center, combining superior professional graphics with powerful compute and AI acceleration to meet today's design, creative, and scientific challenges. Driving the next generation of virtual workstations and server-based workloads, NVIDIA A40 brings features for ray-traced rendering, simulation, virtual production, and more to professionals anytime, anywhere.
- **NVIDIA A16**—These GPUs combined with NVIDIA Virtual PC (vPC) or NVIDIA RTX Virtual Workstations (vWS) software enables remote desktops and workstations with the power and performance to tackle any project from anywhere. Purpose-built for high-density, graphics-rich VDI and leveraging the NVIDIA Ampere architecture, the A16 provides double the user density compared to the previous generation, while ensuring the best possible user experience.
- **NVIDIA A2**—The NVIDIA A2 provides a low power, small footprint, high performance GPU that when combined with Dell PowerEdge R650 servers provide the perfect combination for edge deployments at scale, instantly upgrading existing entry level CPU servers to handle a variety of graphical VDI workloads. Featuring a low-profile PCIe Gen4 card and a low 40-60 watt (W) configurable thermal design power (TDP) capability, the A2 brings adaptable graphics acceleration to any compatible server at an entry-level price point.
- **NVIDIA T4**—Powered by NVIDIA Turing Tensor Cores, the NVIDIA T4 provides a low-profile, low power design at 70 W and provides for VDI graphics enablement at the edge when combined with Dell PowerEdge R650 servers. The T4 GPU is flexible enough to run Knowledge Worker VDI or professional graphics workloads and is optimized for maximum utility in enterprise data centers and the cloud.

NOTE: All Ampere ("A") GPUs are supported with vGPU on vSphere 7.0 update 2 and newer.

The following table summarizes some of the relevant specifications of the previously mentioned GPUs:

Table 3. GPU specifications

GPU	Architecture	Maximum number of vGPU users per card	Thermal design	Dimensions	Memory	VDI use cases
NVIDIA A40	Ampere	32	300W	Full height, full length (FHFL) dual slot	48 GB GDDR6	For mid-range to high-end 3D design and creative workflows.
NVIDIA A16	Ampere	64	250W	FHFL dual slot	4 x 16 GB GDDR6	Office productivity applications, graphics-rich virtual desktops, and low to mid-range creative workflows.
NVIDIA A2	Ampere	16	40-60W	1-slot Low-Profile PCIe	16 GB GDDR6	Office productivity applications, graphics-rich virtual desktops, and low

Table 3. GPU specifications (continued)

GPU	Architecture	Maximum number of vGPU users per card	Thermal design	Dimensions	Memory	VDI use cases
						to mid-range creative workflows.
NVIDIA T4	Turing	16	70W	1-slot Low-Profile PCIe	16 GB GDDR6	Office productivity applications, graphics-rich virtual desktops, and low to mid-range creative workflows. Use in space-constrained environments such as 1U servers or edge deployments or where versions of vSphere prior to 7.0 update 2 are required.

Client components

Users can access the virtual desktops through various client components. The following table lists the client components that Dell Technologies recommends:

Table 4. Recommended client components

Component	Description	Recommended use	More information
Latitude laptops and 2-in-1 PCs	• Biggest screens in a smaller footprint with a wide array of ports to connect peripherals and enjoy speakerphone experience • More responsive apps with Dell Optimizer and intelligent audio for a better conference experience • Better connectivity including 4G LTE, Wi-Fi 6, and eSIM • 5G design on the Latitude 9510 • Smart antenna design on select products for better connections	• Mobility and space-saving devices • Allows users to be productive and stay connected with versatile, space-saving mobile solutions • Offers a modern portfolio that is built to prioritize customer experience and keep employees productive wherever they work with a selection of laptops, 2-in-1s, and ecosystem products	www.delltechnologies.com/Latitude
OptiPlex business desktops and All-in-Ones	• Intel 9th Generation core processors, providing 2 x system responsiveness with Intel Optane Memory • Flexible expansion options, including rich CPU, SSD, and PCIe NVMe • Many innovative form factors with versatile mounting options, including the industry's only zero-footprint modular desktop hidden in plain sight, and space-saving AIOs • Rich interaction with display technology, including 4k UHD AiO and matching multi-monitor support	• The ultimate modular solution • Ideal for desk-centric and remote workers in fixed environments who require varying degrees of performance and expandability	www.delltechnologies.com/OptiPlex

Table 4. Recommended client components (continued)

Component	Description	Recommended use	More information
Precision workstations	• The most complete workstation portfolio with towers, racks, and mobile form factors • Powerful workstations for the most demanding applications, scalable storage, and RAID options • Smallest, most intelligent, and highest-performing mobile workstation portfolio • Rack workstations delivering shared or dedicated resources • Ensures peace of mind with ISV-certified, reliable performance	• High-end graphics and extreme performance • Precision workstations designed to run processor- and graphic-intensive applications and activities with mission-critical reliability such as analytics, simulations, and modeling	www.delltechnologies.com/Precision

Storage

Dell PowerMax

PowerMax is the industry's most secure mission-critical storage platform and the latest release of software (PowerMaxOS 10) delivers over 200 new features with industry-leading storage automation and advanced cyber resiliency to help companies accelerate their journey through digital transformation. It is the next-generation multi-node scale-out architecture with end-to-end NVMe and 100 Gb per second Infiniband (Dynamic Fabric technology). PowerMax 2500 and PowerMax 8500 with end-to-end NVMe design increases performance, scalability, and flexibility for customer innovation, providing IT organizations with reliable mission-critical storage that remains continuously modern over time. PowerMax delivers up to 2x faster performance, up to 14x higher storage density, and includes 4:1 data reduction guaranteed for open systems configurations.

Figure 5. Dell PowerMax

Dell PowerStore

To enable IT organizations to address their data needs and improve agility to enable digital transformation, Dell Technologies introduces PowerStore, a modern storage appliance that is designed for the data era. PowerStore provides customers with data-centric, intelligent, and adaptable infrastructure that supports both traditional and modern workloads. PowerStore is designed to support any workload by delivering unified storage (physical or virtual, file-based, or container-based) in a performance-optimized appliance that can scale up and out when demands increase.

Figure 6. Dell PowerStore

Dell PowerVault

Optimized for SAN, PowerVault ME5 storage delivers the performance, capacity and operational simplicity that price-sensitive, small to medium-sized businesses demand. PowerVault ME5 supports many types of workloads, including VDI. PowerVault ME5 brings radical simplicity and exceptional affordability to the entry market. It is simple to install and configure, comes with all-inclusive software, and includes an intuitive web-based HTML5 GUI for "anywhere" management. ME5 Series storage flexibly

configures as all-flash or hybrid and you can mix HDD and SSD drives in the configuration. It also includes built-in compatibility with VMware, simplifying physical and virtual resource management.

Figure 7. Dell PowerVault

Network storage

In recent years, the increased growth in the amount of data that is stored in file shares and user home directories across IT environments has resulted in an increased focus on better management of this unstructured data. As a result, many organizations are deploying dedicated file workload solutions with capabilities such as cloud file tiering and single file system namespaces across their IT infrastructure, including for file workloads in a VDI environment.

Dell Technologies provides several solutions for different types of file workloads that you can leverage for user profile management and user data.

Dell PowerStore T

Dell PowerStore T storage is simple, unified storage that enables flexible growth with intelligent scale-up and scale-out capabilities and public cloud integration.

Dell PowerStore T is ideal for general-purpose NAS or SAN mixed workload consolidation, smaller file workloads (including small to midsized VDI environments), and transactional databases.

Dell Technologies recommends that you deploy a separate PowerStore T storage system with a vSphere HA cluster or block when you are deploying Dell PowerStore T in a VDI environment. Each PowerStore T system can scale up to four appliances per cluster. This structure provides the greatest scalability, resiliency, and flexibility when deploying and maintaining file services for the overall user pod. As unstructured data storage needs grow over time, the capacity of each PowerStore T storage system can be scaled up or out independently with minimal user impact. You have the choice to deploy alternative architectures to the one suggested here, but you should carefully consider the tradeoffs.

For guidance about selecting an appropriate PowerStore T storage solution for your file workload requirements, see the Dell PowerStore website.

Dell PowerScale

Dell PowerScale storage is a scale-out NAS solution for any file workload. The PowerScale system is ideal for a wide range of file workloads (including large-scale enterprise VDI environments requiring a single file system namespace), high-performance computing (HPC), archiving, and infrastructure consolidation.

Dell Technologies recommends that you deploy a separate PowerScale system with a vSphere HA cluster or block when you are deploying a PowerScale storage system in a VDI environment. This structure provides the greatest scalability, resiliency, and flexibility for deploying and maintaining file services for the overall user pod. As unstructured data-storage needs grow over time, you can scale up the capacity of each PowerScale storage system independently with minimal user impact. In addition to scaling up each PowerScale chassis, you can also scale out a PowerScale system by using the Dell OneFS operating system. Multiple PowerScale systems can provide a single volume and namespace that all user pods in a data center can access.

For guidance about selecting an appropriate PowerScale storage solution for your file workload requirements, see the Dell PowerScale website.

Figure 8. Dell PowerScale

Networking

Dell Validated Designs for VDI enable flexibility in networking selections. VDI validations have been successfully performed with the following hardware, although several other choices are available:

- **Dell Networking S5248F-ON (25 GbE ToR switch)**—This switch provides optimum flexibility and cost-effectiveness for demanding compute and storage traffic environments. This ToR switch features 48 x 25 GbE SFP28 ports, 4 x 100 GbE QSFP28 ports, and 2 x 100 GbE QFSP28-DD ports. The S5248F-ON switch supports ONIE for zero-touch installation of network operating systems.

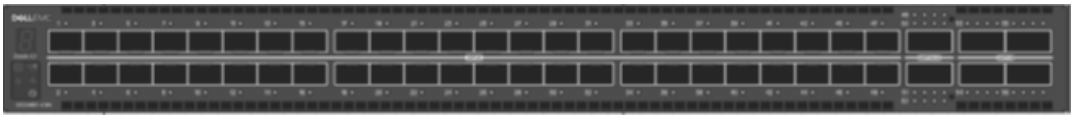

Figure 9. Dell Networking S5248F-ON switch

- **Dell Networking S4048-ON (10 GbE ToR switch)**—This switch optimizes your network for virtualization with a high-density, ultra-low-latency ToR switch that features 48 x 10 GbE SFP+ and 6 x 40 GbE ports (or 72 x 10 GbE ports in breakout mode) and up to 720 Gbps performance. The S4048-ON switch supports ONIE for zero-touch installation of alternate network operating systems.

Figure 10. Dell Networking S4048-ON switch

For more information about these switches, see Dell PowerSwitch S Series 10GbE Switches and Dell PowerSwitch S Series 25/40/50/100 GbE Switches.

Topology

This architecture is designed for true linear scaling and uses a leaf-spine network architecture. This consists of two network tiers: an L2 leaf and an L3 spine. Both tiers are based on 100 GbE and non-blocking switches. This architecture maintains consistent performance without any throughput reduction. The following figure shows this architecture:

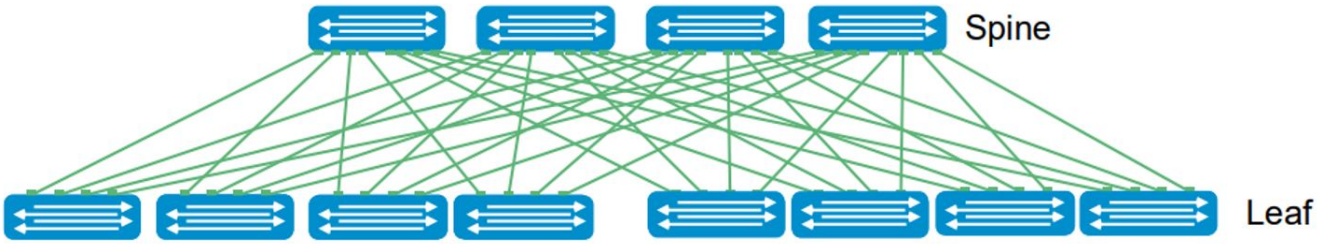

Figure 11. Leaf-spine architecture

SmartFabric Services

SmartFabric Services (SFS) is included with Dell SmartFabric OS10. With SFS, customers can quickly and easily deploy and automate data center networking fabrics. SFS automatically builds an L3 leaf-spine fabric, which enables faster time to production for private cloud environments while being fully interoperable with the existing data center infrastructure.

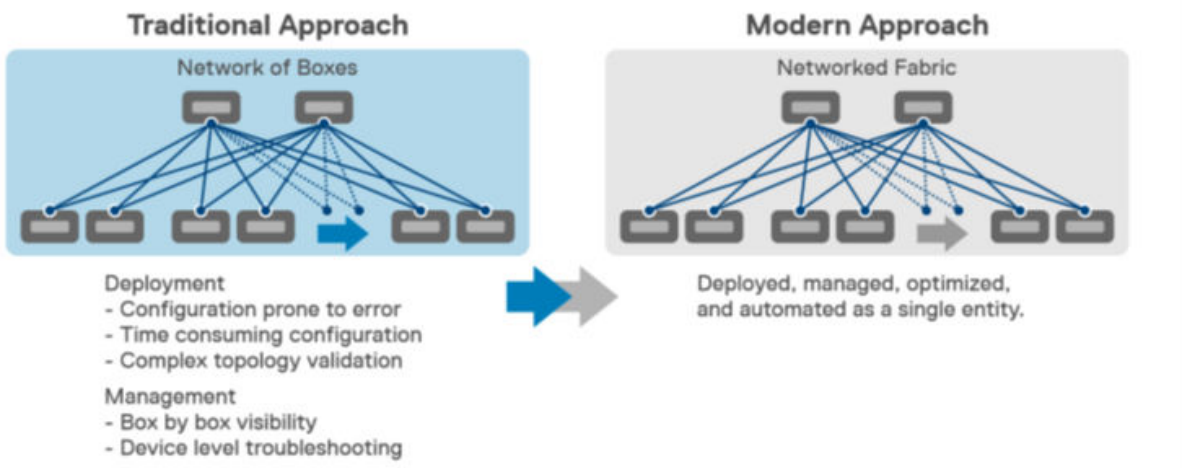

Figure 12. SmartFabric Services (SFS)

For additional configuration, deployment, and other guides for data center networking solutions, see the Data Center Networking Solution Info Hub.

Software

This section provides a high-level overview of the components needed for creating and deploying a VDI environment. Successful deployment requires a deep understanding of the architecture when you are designing the environment.

VMware vSphere

VMware vSphere provides a flexible and secure foundation for business agility, with the following benefits for VDI applications:

- **Improved appliance management**—The vCenter Server Appliance Management Interface provides CPU and memory statistics, network and database statistics, disk space usage, and health data. This reduces reliance on a command-line interface for simple monitoring and operational tasks.
- **VMware vCenter Server native high availability**—This solution for vCenter Server Appliance consists of active, passive, and witness nodes that are cloned from the existing vCenter Server instance. You can enable, disable, or destroy the vCenter HA cluster at any time. Maintenance mode prevents planned maintenance from causing an unwanted failover. The vCenter Server database uses native PostgreSQL synchronous replication, while key data outside the database uses separate asynchronous file system replication.
- **Backup and restore**—Native backup and restore for vCenter Server Appliance enables users to back up vCenter Server and Platform Services Controller appliances directly from the vCenter Server Appliance Management Interface or API. The backup consists of a set of files that is streamed to a selected storage device using the SCP, HTTP(S), or FTP(S) protocol. This backup fully supports vCenter Server Appliance instances with both embedded and external Platform Services Controller instances.
- **VMware vSphere HA support for NVIDIA vGPU-configured VMs**—vSphere HA protects VMs with the NVIDIA vGPU shared pass-through device. In the event of a failure, vSphere HA tries to restart the VMs on another host that has an identical NVIDIA vGPU profile. If no available healthy host meets this criterion, the VM fails to power on.
- **VMware vSAN Enterprise Edition**—Includes all-flash space-efficiency features (deduplication, compression, and erasure coding), software-defined, data-at-rest encryption, and stretched clusters for cost-efficient performance and greater hardware choice.
- **VMware Log Insight**—Provides log management, actionable dashboards, and refined analytics that enable deep operational visibility and faster troubleshooting.

NOTE: vSphere Enterprise Edition (or vSphere Desktop) is required to support NVIDIA graphics cards.

VMware Horizon

The architecture described here is based on VMware Horizon 8, which provides a complete end-to-end solution that delivers Microsoft Windows virtual desktops to users on a wide variety of endpoint devices. Virtual desktops are dynamically assembled on demand, providing pristine, yet personalized, desktops each time a user logs in.

VMware Horizon 8 provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the VDI.

NOTE: For more information, see the Horizon resources page and VMware Horizon Frequently Asked Questions.

The core Horizon components include:

- **Horizon Connection Server (HCS)**—Installed on servers in the data center, HCS brokers client connections, authenticates users, entitles users by mapping them to desktops or pools, establishes secure connections from clients to desktops, supports single sign-on, and sets and applies policies.
- **Horizon Administrator**—Provides administrator functions such as deploying and managing Horizon desktops and pools, setting and controlling user authentication, and more.
- **Horizon Agent**—Provides a means of communication with Horizon clients. The agent is installed on all VMs, physical machines, and Terminal Service servers that are used as a source for Horizon desktops. On VMs, the agent communicates with the Horizon client to provide services such as USB redirection, printer support, and more.
- **Horizon Client**—Installed on endpoints, the client creates connections to Horizon desktops that can be run from tablets, Windows, Linux, or Mac PCs or laptops, thin clients, and other devices.
- **Unified Access Gateway**—Provides a way to securely deliver connections that require a higher level of security to access, such as remote internet connections.
- **Horizon Portal**—Provides access to download links for full Horizon clients. The portal enables the HTML access feature to run a Horizon desktop inside a supported browser.
- **vCenter Server**—Provides centralized management and configuration to the entire virtual desktop and host infrastructure. It facilitates configuration, provisioning, and management services.

Horizon desktop cloning

VMware Horizon 8 offers the following methods for cloning desktops:

- **Full clones**—These are typically used for testing purposes or to create management VMs. Full clones are not ideal for VDI because full copies have no connection to the original VM. You must update each VM with this approach.
- **Instant clones**—These are available with Horizon Universal Subscription, Horizon Standard Subscription, and Horizon Enterprise Edition (TERM) licenses. This technology provisions a VM immediately after a user requests one. This is a far easier approach to operating system updates and patch management because the VM is created when it is needed. You can use the combination of products such as VMware App Volumes and Dynamic Environment Manager to emulate persistence.

NOTE: Horizon Linked Clones and Composer were deprecated in Horizon 8 2006 and were removed in Horizon 8 2012.

NVIDIA vGPU

NVIDIA vGPU virtualizes GPU hardware acceleration to share GPUs between multiple virtual desktops or aggregate and assign them to a single virtual desktop, without compromising the graphics experience. NVIDIA vGPU offers the following software variants to enable graphics for different virtualization techniques:

- **NVIDIA Virtual Applications (vApps)**—Delivers graphics accelerated applications using Remote Desktop Service Host (RDSH).
- **NVIDIA Virtual PC (vPC)**—Provides full virtual desktops with up to dual 4K monitor support or single 5K monitor support.
- **NVIDIA RTX Virtual Workstation (vWS)**—Provides workstation-grade performance in a virtual environment with support for up to four Quad 4K or 5K monitors or up to two 8K monitors.

3

Validation

Topics:

Validation

Overview

Performance analysis and characterization (PAAC) testing on Dell VDI solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and end-user experience (EUE) during load-testing. This ensures the optimal combination of EUE and cost-per-user.

Test and performance analysis methodology

Login VSI performance testing process and monitoring

We tested each user load against four runs:

- A pilot run to validate that the infrastructure was performing correctly and that valid data could be captured.
- Three subsequent runs to enable data correlation.

During testing, while the environment was under load, we logged in to a session and completed tasks that correspond to the user workload. This test is subjective, but it provides a better understanding of the EUE in the desktop sessions, particularly under high load. It also helps to ensure reliable data gathering.

To ensure that the user experience was not compromised, we monitored the following important resources:

- **Compute host servers**—Solutions based on VMware vCenter for VMware vSphere gather key data (CPU, memory, disk, and network usage) from each of the compute hosts during each test run. This data is exported to .csv files for single hosts and then consolidated to show data from all hosts. While the report does not include specific performance metrics for the management host servers, these servers are monitored during testing to ensure that they are performing at an expected level with no bottlenecks.
- **Hardware resources**—Resource overutilization can cause poor EUE. We monitored the relevant resource utilization parameters and compared them to relatively conservative thresholds. These thresholds are shown in the following table. They were selected based on industry best practices and our experience to provide an optimal trade-off between good EUE and cost-per-user while also allowing sufficient burst capacity for seasonal or intermittent spikes in demand.

Table 5. Resource utilization parameters

Parameter	Pass/fail threshold
Physical host CPU utilization	85%[a]
Physical host memory utilization	85%
Network throughput	85%
Disk latency	20 milliseconds (ms)
Login VSI failed sessions	2%

a. The Dell Validated Design for VDI team recommends that average CPU utilization not exceed 85 percent in a production environment. A 5 percent margin of error was allocated for this validation effort. Therefore, CPU utilization sometimes exceeds our recommended percentage. Because of the nature of Login VSI testing, these exceptions are reasonable for determining our sizing guidance.

Load generation

Login VSI installs a standard collection of desktop application software, including Microsoft Office and Adobe Acrobat Reader, on each VDI desktop testing instance. It then uses a configurable launcher system to connect a specified number of simulated users to available desktops within the environment. When the simulated user is connected, a login script configures the user environment and starts a defined workload. Each launcher system can launch connections to several VDI desktops (target machines). A centralized management console configures and manages the launchers and the Login VSI environment.

We used the following login and boot conditions:

- Users were logged in within a login timeframe of one hour.
- All desktops were started before users were logged in.

Login VSI workloads

The following table describes the Login VSI workloads that we tested:

Table 6. Login VSI workloads

Login VSI workload name	Workload description
Task Worker	A light workload that runs fewer applications and starts/stops them less frequently than the other workloads, resulting in lower CPU, RAM, and IO usage. The Task Worker workload uses the following applications: ● Adobe Reader ● Internet Explorer ● Microsoft Excel ● Microsoft Outlook ● Microsoft Word It also uses native Windows apps (Notepad and 7-Zip) to complete the print and zip actions used by the workload meta functions.
Knowledge Worker	Designed for virtual machines with 2 vCPUs. This workload includes the following activities: ● Microsoft Outlook—Browse messages. ● Internet Explorer—Browse websites and open a YouTube-style video (480p movie trailer) three times in every loop. ● Microsoft Word—Start one instance to measure response time and another to review and edit a document. ● Doro PDF Printer and Acrobat Reader—Print a Word document and export it to PDF. ● Microsoft Excel—Open a large, randomized sheet. ● Microsoft PowerPoint—Review and edit a presentation. ● FreeMind—Run a Java-based Mind Mapping application. ● Other—Perform various copy and zip actions.

Test configuration details

The following tables list the hardware and software components of the infrastructure that was used for the PAAC test. All host machines were updated with the latest operating system and security updates.

The following table shows the hardware components:

Table 7. Hardware components

Hardware type	Component
Compute host hardware	1 x PowerEdge R840 Servers ● 4 x Intel Xeon Gold 6248 @ 2.5 GHz, 20-core processors ● 1,536 GB Memory @ 2933 MT/s (24 x 64 GB DDR4) ● Broadcom Adv. Dual 25 Gb Ethernet 25 Gbps
Management host hardware	R730
Storage	PowerStore 1000T ● PowerStore OS 2.1.1.1 ● 6 x 1.92 TB NVMe SSD ● 25 GbE NVMe/TCP
Network	S5248-ON Switch

The following table shows the software components:

Table 8. Software components

Software type	Program/version
Display protocol	BLAST Extreme H.264 + Switch Codec
Broker	VMware Horizon 2203 (8.5)
Hypervisor	vSphere ESXi 7.0.3, 19482537
SQL	Microsoft SQL Server 2019
Desktop operating system	Microsoft Windows 10 Enterprise 64-bit (version 21H2)
Office software	Microsoft 365 Apps for Business 2205
Profile Management	Dynamic Environment Manager 2203
Management operating system	Microsoft Windows Server 2022
Login VSI version	4.1.40.1
Antivirus software	Windows Defender

User VM configurations

The following table summarizes the standard PAAC VM configurations for Login VSI currently used for the various profiles/ workloads tested:

Table 9. Login VSI configurations

Workload	VM profiles				
	vCPUs	RAM	RAM reserved	Desktop video resolution	Operating system
Knowledge Worker	2	4 GB	2 GB	1920 x 1080	Windows 10 Enterprise 64-bit

Standard VDI test results and analysis

Summary of results

The following table summarizes the host utilization metrics for the different Login VSI workloads that we tested, and the user density derived from Login VSI performance testing:

Table 10. Summary of test results

Server configuration	Workload	Operating system	User density	Average CPU	Average memory active (GB)	Average IOPS per user	Average net Mbps per user
4-socket Density-Optimized	Knowledge Worker	Windows 10, 21H2	225	84.17%	289.25	66.65	4.62

The host utilization metrics mentioned in the table are defined as follows:

- **User density**—The number of users per compute host that successfully completed the workload test within the acceptable resource limits for the host. For clusters, this number reflects the average of the density achieved for all compute hosts in the cluster.
- **Average CPU**—The average CPU usage over the steady-state phase. For clusters, this number represents the combined average CPU usage of all compute hosts. On the latest Intel processors, the ESXi host CPU metrics exceed the rated 100 percent for the host if Turbo Boost is enabled, which is the default setting. An additional 35 percent of CPU is available from the Turbo Boost feature, but this additional CPU headroom is not reflected in the VMware vSphere metrics where the performance data is gathered.

- **Average active memory**—For ESXi hosts, the amount of memory that is actively used, as estimated by the VMkernel based on recently touched memory pages. For clusters, this is the average amount of physical guest memory that is actively used across all compute hosts over the steady state period.
- **Average IOPS per user**—IOPS calculated from the average cluster disk IOPS over the steady state period divided by the number of users.
- **Average network usage per user**—Average network usage on all hosts calculated over the steady state period divided by the number of users.

Login VSI Knowledge Worker Workload

We performed this test with the Login VSI Knowledge Worker workload on a single PowerEdge R840 (see Table 1). We created the desktop VMs using VMware Horizon instant clone technology and used the VMware Horizon Blast Extreme display protocol. We populated the compute host with 225 desktop VMs.

CPU usage

The following graphs show the CPU utilization of the compute host during testing. CPU usage with all VMs powered on was approximately 5.5 percent before the test started. The CPU usage steadily increased during the login phase, as shown in the following figure.

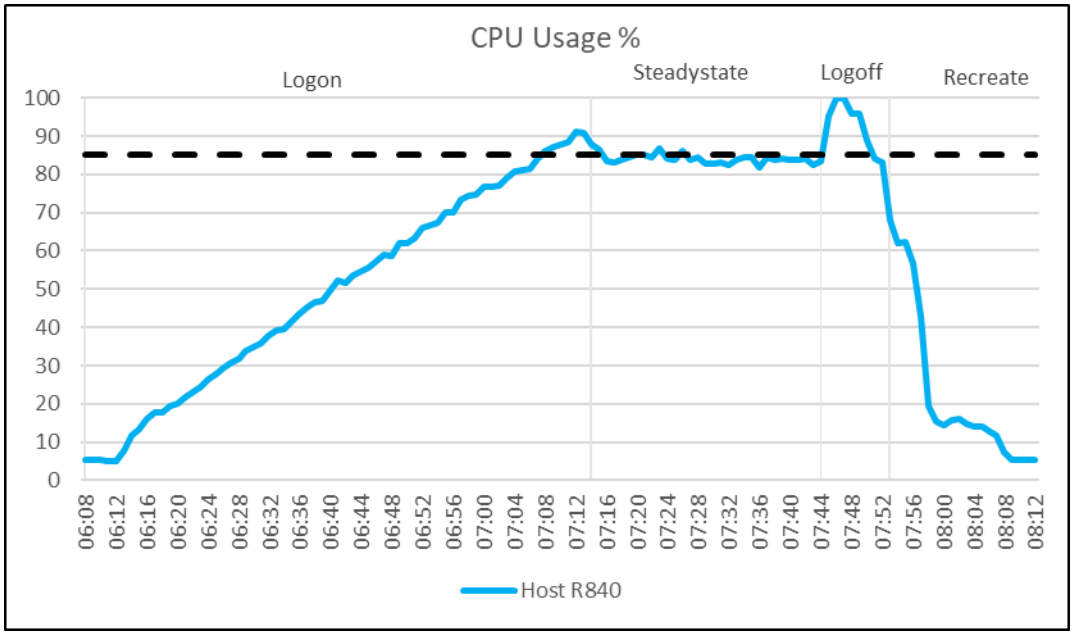

Figure 13. CPU usage

During the steady-state phase, an average CPU utilization of 84.17 percent was recorded. This value is close to the pass/fail threshold that we set for average CPU utilization (see Table 4). To maintain good EUE, do not exceed this threshold. You can load more user sessions while exceeding this threshold for CPU, but you might experience a degradation in user experience.

As shown in the following figure, the CPU readiness was well below the 5 percent threshold that we set.

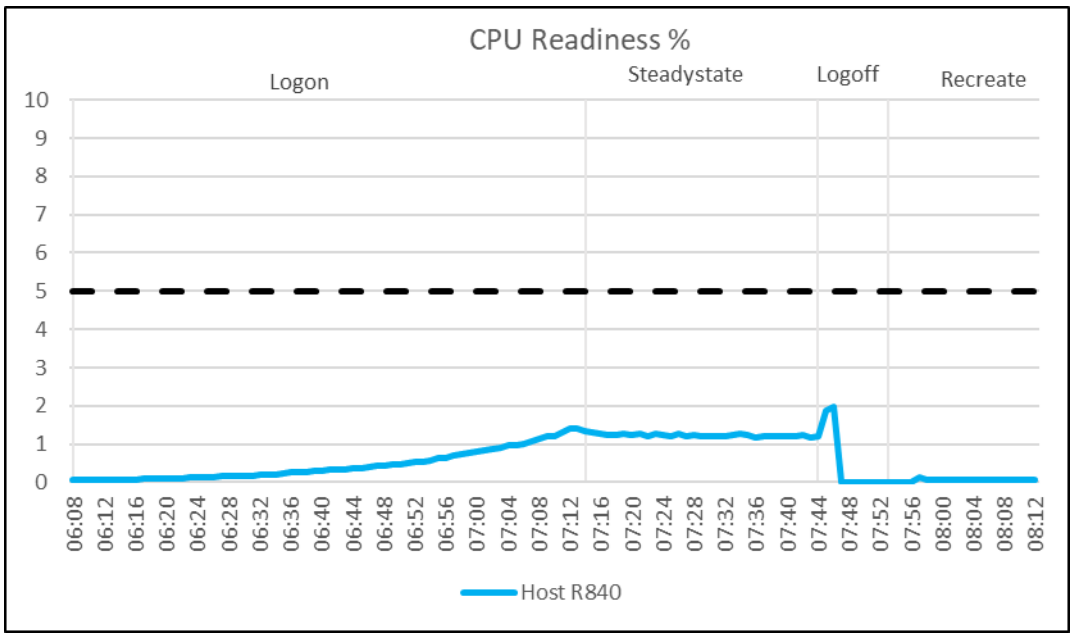

Figure 14. CPU readiness

Memory

We observed no memory constraints during the testing on the compute host. Out of 1,536 GB of available memory per node, the compute host reached a maximum consumed memory of 917.83 GB and a steady-state average of 908 GB. Active memory usage reached a maximum active memory of 311.16 GB and recorded a steady-state average memory of 367 GB. There was no memory ballooning or swapping on the hosts. The following figures show consumed and active memory.

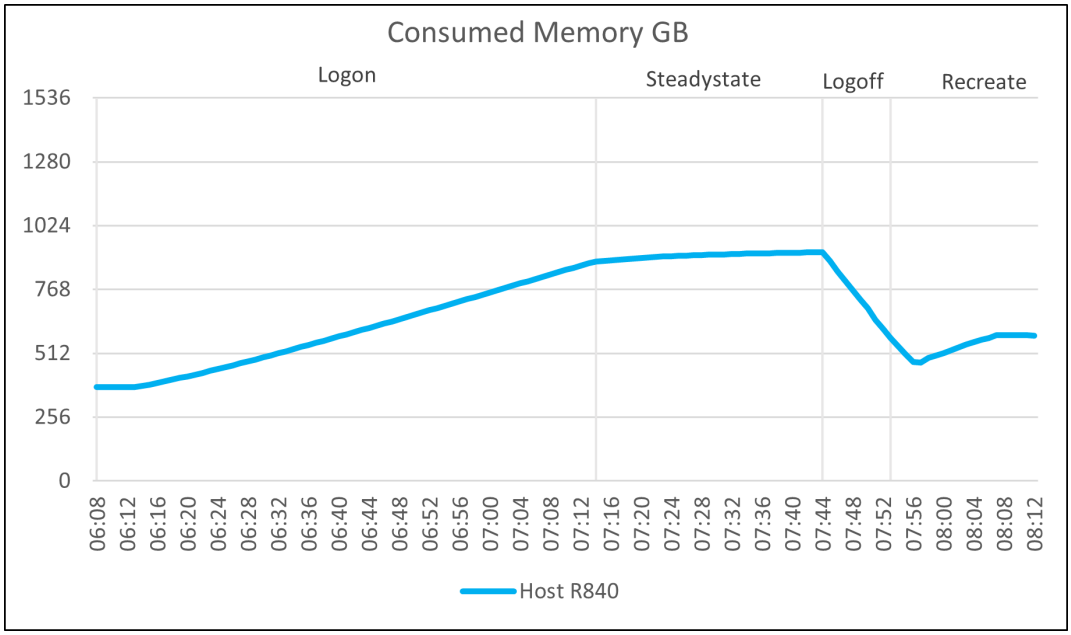

Figure 15. Memory consumed

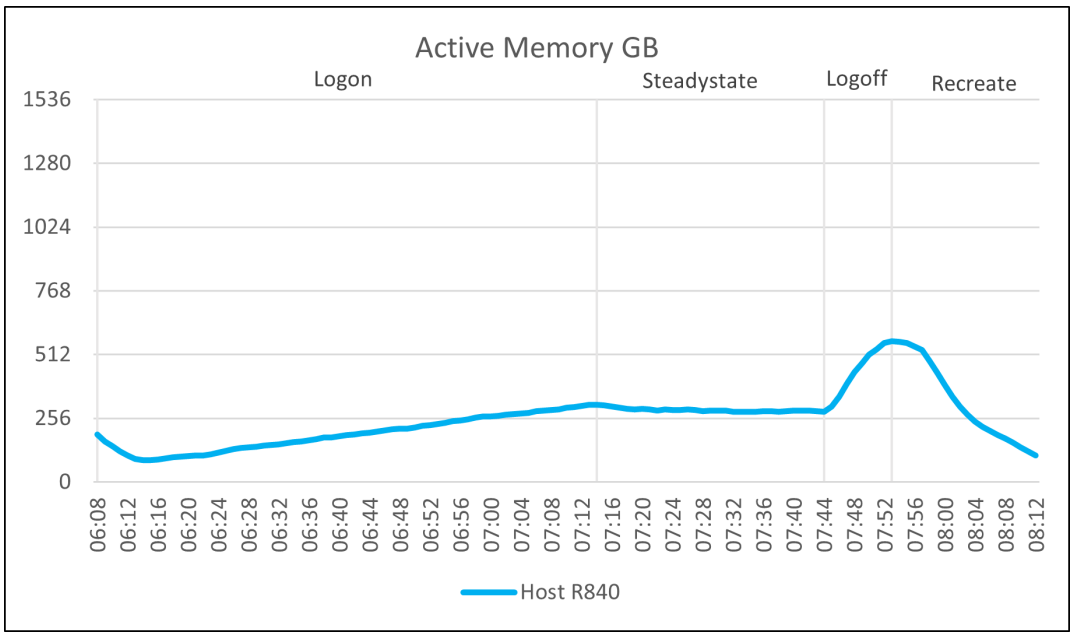

Figure 16. Memory active

Network usage

Network bandwidth was not an issue during testing. The network usage recorded a steady-state average of 1,809 Mbps. The busiest period for network traffic was during the login phase when a peak value of 4,268 Mbps was recorded. The following figure shows network usage:

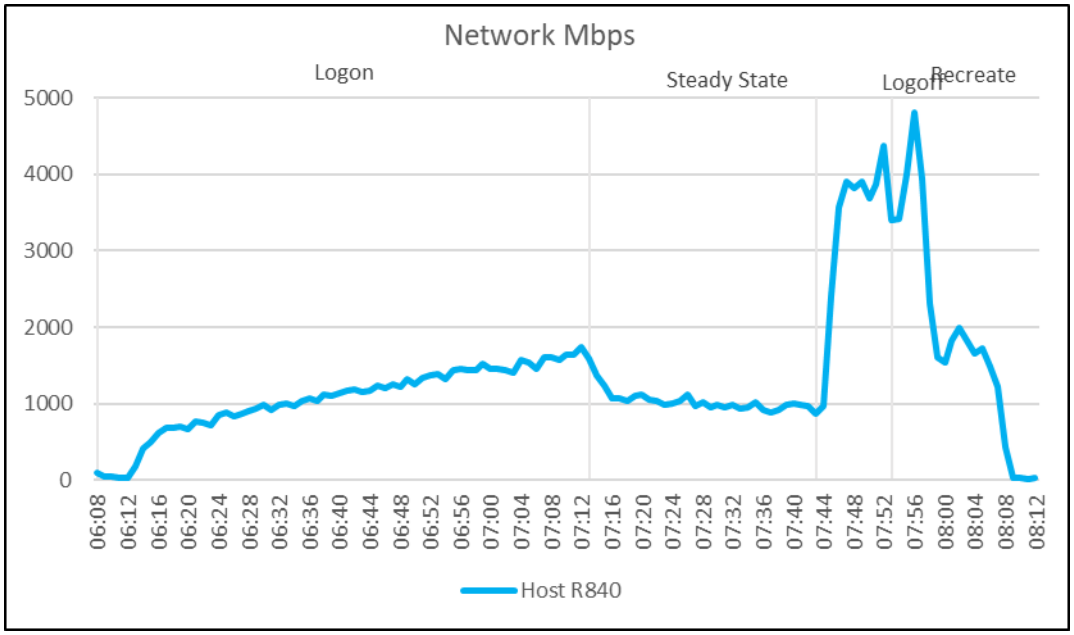

Figure 17. Network usage

Storage IOPS

IOPS reached a maximum value of 7,541 for read IOPS and 9,273 for write IOPS during the steady-state phase. The average steady-state IOPS was 16,814 combined read and write IOPS. The following figure shows storage IOPS:

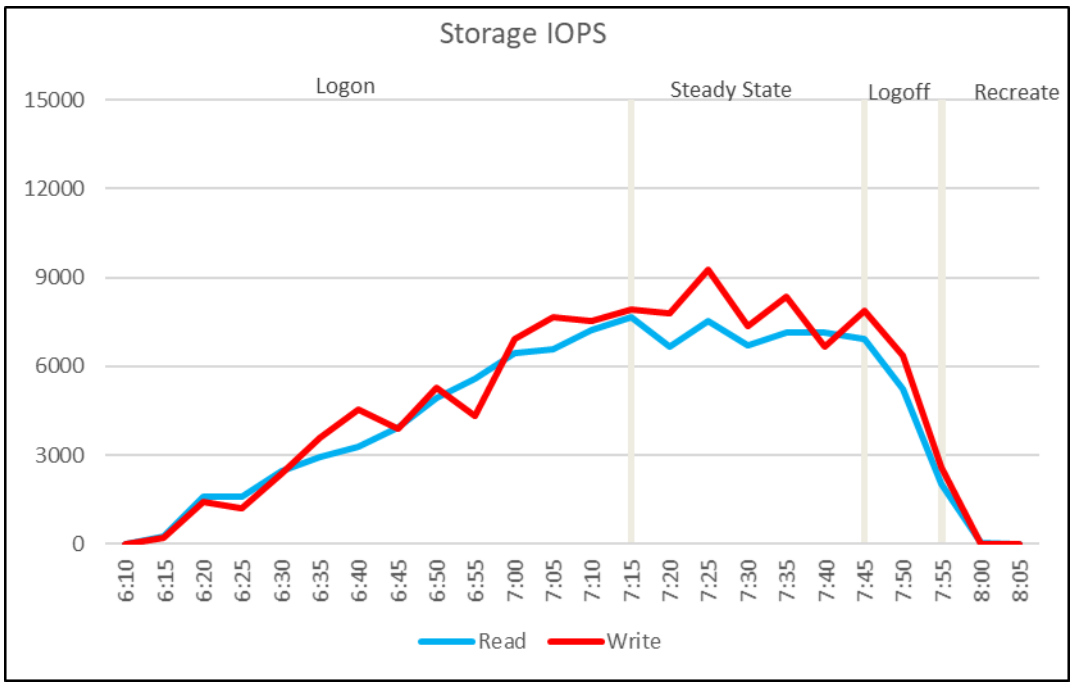

Figure 18. Storage IOPS

Storage I/O latency

Storage I/O latency reached a maximum read latency of 0.03 ms and a maximum write latency of 0.71 ms during the steady-state phase. The average steady-state latency was 0.72 ms. The following figure shows storage I/O latency:

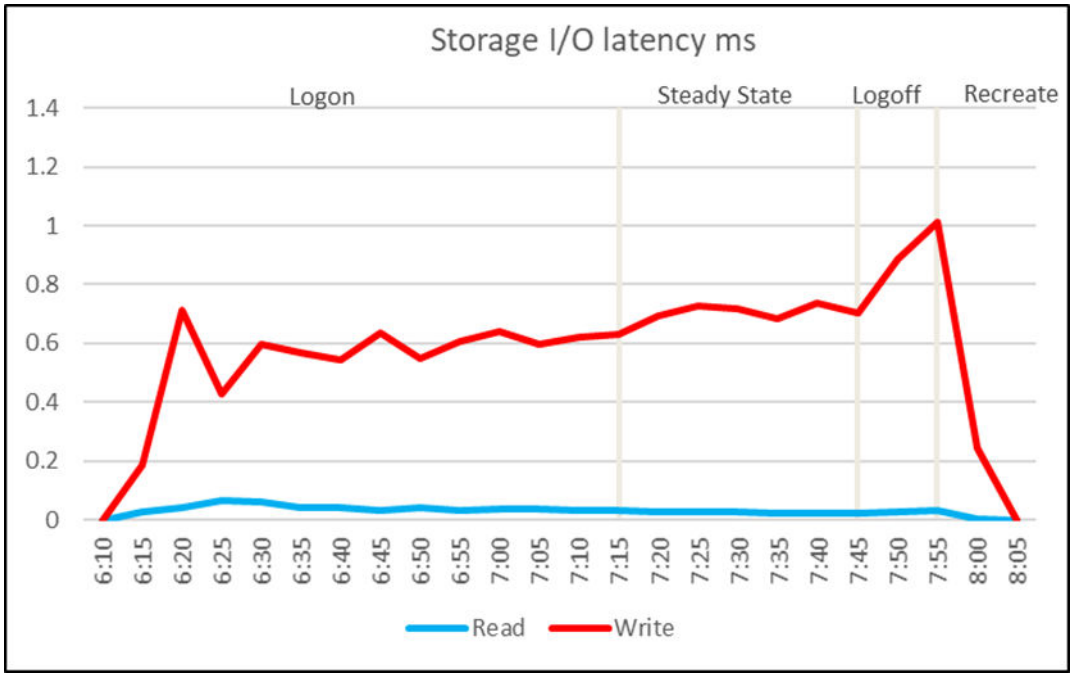

Figure 19. Storage I/O latency

User experience

The baseline score for the Login VSI test was 1,237. This score falls in the 1,199 to 1,599 range rated as "Fair" by Login VSI. For more information about Login VSI baseline ratings and baseline calculations, see VSImax baseline scores. As indicated by the blue line in the following figure, the system reached a VSImax average score of 1,665 when 225 sessions were loaded. This

value is well below the VSI threshold score of 2,237 set by the Login VSI tool. During testing, VSImax was never reached, which typically indicates a stable system and a better user experience.

The Login VSImax user experience score for this test was not reached. When manually interacting with the sessions during the steady-state phase, the mouse and window movement were responsive, and video playback was good. No "stuck sessions" were reported during the testing, indicating that the system was not overloaded at any point. See Appendix A, which explains the Login VSI metrics.

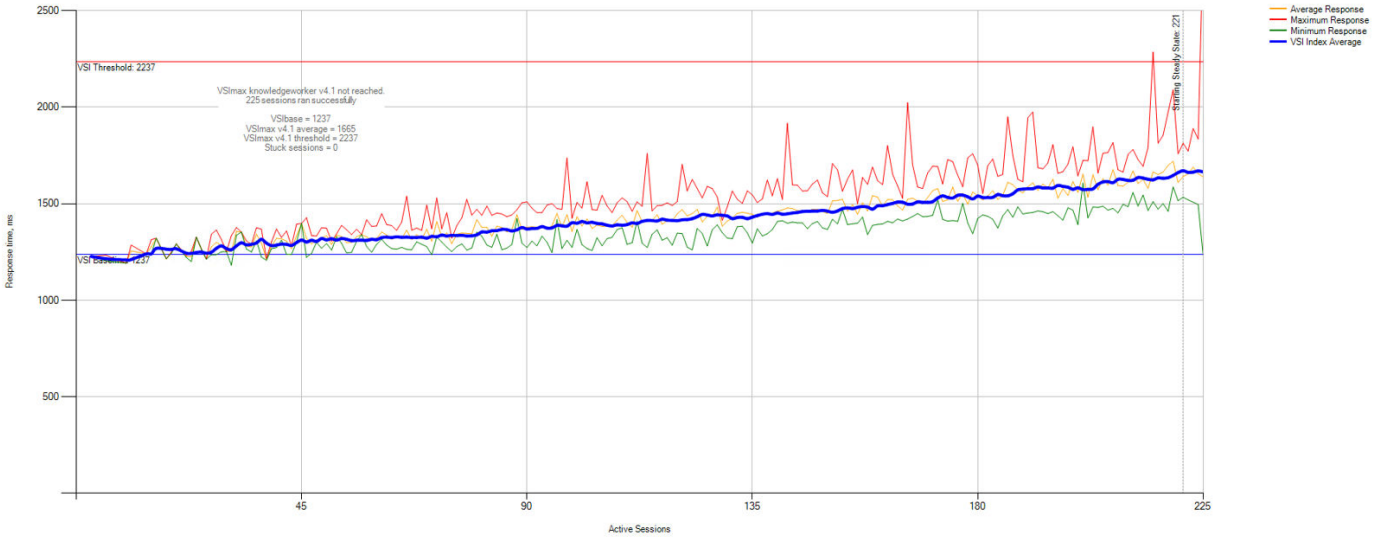

Figure 20. User experience

4

Sizing the Solution

Topics:

Sizing and scaling overview

VDI solutions based on 3-tier architectures using PowerEdge servers, Dell storage, and Dell PowerSwitch networking provide flexibility as you scale, reducing the initial and future cost of ownership. Add physical and virtual servers to the server pools to scale horizontally. Add virtual resources to the infrastructure to scale vertically.

Sizing guidelines

This section provides best practices for sizing your VDI deployment.

Platform configurations

With several configurations to choose from, consider these basic differences:

- The Density Optimized configurations provide a good balance of performance and scalability for various general-purpose VDI workloads.
- The Virtual Workstation configurations provide the highest levels of performance for more specialized VDI workloads, which means you can use it with ISV and high-end computing workloads.

CPU

User density and graphics considerations include:

- For architectures with 2nd Generation Intel Xeon Scalable processors:
  - **Knowledge Workers**—2.8 users per core. For example, 44 knowledge users with dual eight-core processors
- For architectures with 3rd Generation Intel Xeon Scalable processors:
  - See the VDI Design Guide—VMware Horizon on VxRail and vSAN Ready Nodes for approximate scaling guidance.
- For graphics:
  - For high-end graphics configurations with NVIDIA vWS graphics enabled, choose higher clock speeds over higher core counts. Many applications that benefit from high-end graphics are engineered with single-threaded CPU components. Higher clock speeds benefit users more in these workloads.
  - For NVIDIA vPC configurations, use higher core counts over faster clock speeds to reduce oversubscription.
  - Most graphics configurations do not experience high CPU oversubscription because vGPU resources are likely to be the resource constraint in the appliance.

NOTE: vSAN based solutions provide approximately the same density as 3-tier architectures as long as the processor and memory configurations are the same, data efficiency features are not enabled within the vSAN configuration, and there are no other bottlenecks in place in either type of environment.

Memory

Best practices for memory allocation and configuration include:

- Do not overcommit memory when sizing, because memory is often not the constraining resource. Overcommitting memory increases the possibility of performance degradation if contention for memory resources occurs (for example, swapping and ballooning of memory). Overcommitted memory can also impact storage performance when swap-files are created.
- **3rd Generation Intel Xeon Scalable Processors**—Populate memory in units of eight DIMMs per CPU to yield the highest performance. Dell PowerEdge servers using 3rd Generation Intel Xeon Scalable processors have eight memory channels per CPU, which are controlled by four internal memory controllers, each handling two memory channels. To ensure that your environment has the optimal memory configuration, use a balanced configuration, where each CPU supports a maximum of 16 DIMMs (or 32 DIMMs for a dual-CPU server). The most effective configuration on the PowerEdge R650 and R750 is 16 DIMMs (8 per processor).
- **2nd Generation Intel Xeon Scalable Processors**—Populate memory in units of six per CPU to yield the highest performance. Dell PowerEdge servers using 2nd Generation Intel Xeon Scalable processors have six memory channels per CPU, which are controlled by two internal memory controllers, each handling three memory channels. To ensure that your environment has the optimal memory configuration, use a balanced configuration, where each CPU supports a maximum of 12 DIMMs (or 48 DIMMs for a quad-CPU server). The most effective configuration on the PowerEdge R840 is 24 DIMMS (6 per processor).
- Use Intel Optane Persistent Memory (PMem) for cost savings over traditional DRAM or in situations where high memory capacity is required (1 TB or greater). vSphere 7 update 3 introduced vSphere Memory Monitoring and Remediation (vMMR), which provides visibility of performance statistics for tiered memory. For additional information, see the VMware documentation on vMMR.

NVIDIA vGPU considerations

Best practices for sizing and configuring solutions that require graphics accelerators include:

- vPC licenses that support up to 2 GB of frame buffer and up to two 4K monitors or a single 5K monitor to cover most traditional VDI users. Maximum node density for graphics-accelerated use can typically be calculated as the available frame buffer per node divided by the frame buffer size.
- The addition of GPU cards does not necessarily reduce CPU utilization. Instead, it enhances the user experience and offloads specific operations best performed by the GPU.
- Dell Technologies recommends using the BLAST protocol for vGPU enabled desktops. NVIDIA GPUs are equipped with encoders that support BLAST.
- Virtual Workstations are typically configured with at least 2 GB video buffer.

Sizing considerations

Best practices for sizing your deployment include:

- **User density**—If concurrency is a concern, calculate how many users will use the environment at peak utilization. For example, if only 80 percent are using the environment at a time, the environment must support only that number of users (plus a failure capacity).
- **Disaster recovery**—For DR planning, Dell Technologies recommends implementing a dual/multi-site solution. The goal is to keep the environment online and, in case of an outage, to perform an environment recovery with minimum disruption to the business.
- **Management and compute clusters**—For small test environments, it is acceptable to use a combined management and compute cluster. For environments deployed at a larger scale, we recommend that you separate the management and compute layers. When creating a management cluster for a large-scale deployment, consider using the PowerEdge R650 platform to reduce the data center footprint. With a larger variety of configuration possibilities, the PowerEdge R840 or R750 platforms are preferred for compute clusters.
- **Network isolation**—When designing for larger-scale deployments, consider physically separating the management and VDI traffic from the vSAN traffic for traffic isolation and to improve network performance and scalability. This design illustrates a four-NIC configuration per server with all the traffic separated physically as well as logically using VLANs.

NOTE: The VMware Workspace ONE and VMware Horizon Reference Architecture provides more details about multi-site design considerations for Horizon.

Scaling guidelines

Overview

VDI solutions based on Dell PowerEdge, Dell storage, and Dell PowerSwitch networking provide flexibility as you scale, reducing the initial and future cost of ownership. Add physical and virtual servers to the server pools to scale horizontally (scaling out). Add additional resources to the individual components of the infrastructure to scale vertically (scaling up).

Scaling out

Each component of the solution architecture scales independently, depending on the required number of supported users. You can add appliance nodes at any time to expand the vSAN SDS pool in a modular fashion. The scaling limit for vSAN is restricted by the limits of the hypervisor at 64 nodes per block. The boundary for a Horizon block is the vCenter. The number of VMs a vCenter can host depends on the type of Horizon 8 VMs in use. The recommended limit of virtual machines per vCenter is 20,000 full-clone or instant-clone VMs.

Sizing recommendations change over time as updates are released and qualifications are performed. See the VMware Configuration Maximums website for the latest recommendations.

This Dell Validated Design for VDI uses instant clones, as shown in the following figures.

VMware recommends a limit of 5,000 instant-clone VMs per block. With these limits in mind, 25 compute nodes with 200 task-user VMs per node would reach the maximum number of VMs for the block.

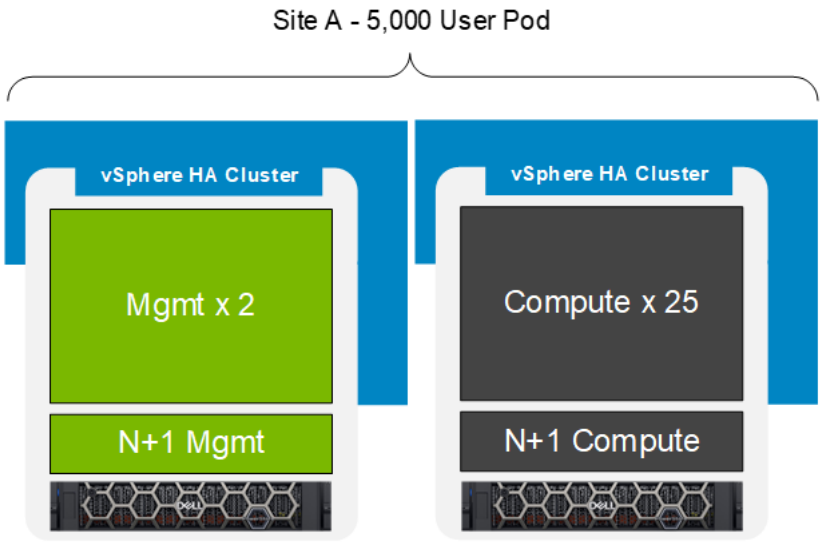

Figure 21. Single 5,000-user block

The following figure shows a scale-out to a 20,000-user Horizon vSAN pod with 5,000 user blocks. Each block contains its own vCenter Server instance and VDI components.

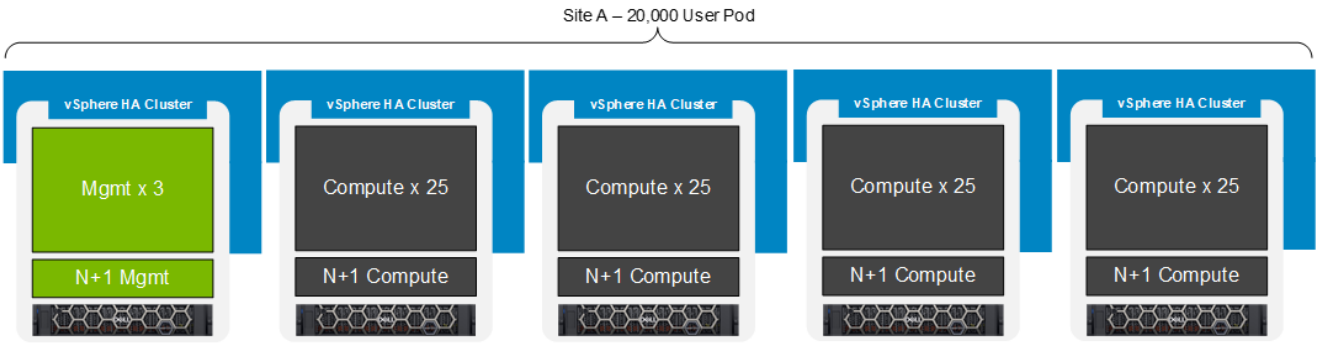

Figure 22. Scale-out for 20,000 users

Scaling up

Each of the components within this design can be scaled up in the following ways:

- Dell PowerEdge servers
  - **Processors**—If more processing power is required, higher core count and clock speed processors can be installed up to the highest model.
  - **Memory**—Additional memory can be installed into each compute node up to the limit of the system architecture.
  - **GPUs**—Additional or more powerful graphics cards can be installed up to the capacity of each compute server.
- Dell storage
  - **Storage Controllers**—Many of the Dell storage systems support scaling up the storage controllers by swapping, clustering, and generational upgrades.
  - Storage capacity and performance can be improved by adding additional disks to storage systems in empty drive slots or by adding expansion shelves.

For more information about Horizon pod and block architecture, and scaling, see the VMware Workspace ONE and VMware Horizon Reference Architecture.

5

Backup and Restore

The growth of VDI adoption has elevated the strategic importance of organizational VDI environments. Users who are critical to business success are increasingly using VDI for their day-to-day productivity tasks. Consequently, the importance of protecting the VDI environment and the business value of its data has also grown as customers seek to ensure their VDI environments meet corporate availability, recovery time objective (RTO), and recovery point objective (RPO) requirements.

For information about data protection of a VMware Horizon environment, see the Data Protection for a VMware Horizon VDI Environment using Dell Data Protection Suite Operations Guide. Dell Technologies provides several data protection solutions for different data protection requirements.

Dell Avamar Virtual Edition

Dell Avamar Virtual Edition is a data protection solution that delivers software-only data protection for virtualized environments and is ideal for the VDI use case. Avamar Virtual Edition is a fully featured data protection solution that is deployed as a virtual appliance. It supports advanced functionality such as cloud-based backup, including VMware Cloud on AWS, change block tracking for fast backup and recovery, and integration with multiple VMware interfaces, such as the vRealize Automation Data Protection Extension. For more information, see Dell Avamar Data Protection Software.

Dell PowerProtect DD Virtual Edition

Dell PowerProtect DD Virtual Edition is a data protection storage solution that runs as a virtual appliance on a customer's choice of hardware or on a variety of public cloud options, including VMware Cloud on AWS. For on-premises deployments, PowerProtect DD Virtual Edition is deployed as a virtual appliance on the relevant hardware platform. PowerProtect DD Virtual Edition has a single point of management with PowerProtect Data Domain Management Center and scales up to 96 TB per instance. One of the key features of the PowerProtect DD storage protection solution is DD Boost, which provides advanced integration with data protection applications such as Avamar Virtual Edition to enable client-side deduplication, thus accelerating backup. For more information, see Dell PowerProtect DD Virtual Edition.

Other Dell Technologies data protection products

Dell Technologies provides other data protection products for specific use cases. Products include a range of appliances that reduce data protection complexity. These scalable, preconfigured solutions combine data protection storage with software, search, and analytics. For more information, see Dell Technologies Data Protection and Backup Solutions.

6

Summary

Topics:

Overview

This design guide describes the integration of servers, storage and networking from Dell Technologies and VMware Horizon 8 brokering software to create virtual application and desktop environments. This architecture provides exceptional scalability and an excellent user experience and empowers IT teams to play a proactive strategic role in the organization.

Dell Technologies offers comprehensive, flexible, and efficient VDI solutions that are designed and optimized for the organization's needs. These VDI solutions are easy to plan, deploy, and run.

Dell Validated Designs for VDI offer several key benefits to clients:

- Predictable costs, performance, and scalability to support a growing workforce
- Rapid deployments
- Rapid scaling, ready to serve enterprises of any size
- Dell Technologies support

With VDI solutions from Dell Technologies, you can streamline the design and implementation process, and be assured that you have a solution that is optimized for performance, density, and cost-effectiveness.

To explore more of our Validated Designs for VDI, contact your Dell Technologies account representative. For additional resources and other VDI designs, see the Dell Technologies Solutions Info Hub for VDI.

7

References

Dell Technologies documentation

The following Dell Technologies documentation provides additional information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative. Also see the VDI Info Hub for a complete list of VDI resources.

- Dell Technologies Virtual Desktop Infrastructure
- Dell Technologies Solutions Info Hub for VDI
- Dell VxRail Hyperconverged Infrastructure
- Dell vSAN Ready Nodes
- Dell PowerEdge Servers
- Dell Data Storage
- Dell Networking
- Dell SmartFabric Services
- PowerStore Info Hub
- PowerStore Product Documentation & Videos

VMware documentation

The following VMware documentation provides additional information:

- VMware vSphere documentation
- VMware Horizon documentation
- vSAN Ready Node Configurator
- VMware Compatibility Guide
- vSAN Hardware Quick Reference Guide
- Best Practices for Published Applications and Desktops in VMware Horizon and VMware Horizon Apps
- VMware Workspace ONE and VMware Horizon Reference Architecture

NVIDIA documentation

The following NVIDIA documentation provides additional information:

- NVIDIA Virtual GPU Software Quick Start Guide

A

Appendix A: Definition of performance metrics

The following table explains the performance metrics used during our testing:

Table 11. Definition of performance metrics

Metric	Definition
CPU usage	The average CPU percentage usage over the steady-state period.
CPU core utilization	The CPU utilization percentage of the corresponding core. A core is utilized if one or both of its logical CPUs are used. This figure is averaged across all cores.
CPU readiness	The percentage of time that the virtual machine was ready but could not get a scheduled run on the physical CPU.
Consumed memory	The amount of host physical memory consumed by a virtual machine, host, or cluster.
Active memory	The amount of memory that is actively used, as estimated by the VM kernel based on recently touched memory pages.
Network usage	Network usage per user is the average over the steady-state period divided by the number of users on a host in megabits per second.
Cluster disk IOPS	Read and write Input/Outputs Operations Per Second (IOPS) consumed by all vSAN clients in the cluster, such as virtual machines, stats objects, and so on.
Disk I/O (or cluster) latency	Average read and write latency of I/Os generated by all vSAN clients in the cluster, such as virtual machines, stats objects, and so on.