Virtual Desktop Infrastructure with VMware Horizon on VxRail and vSAN Ready Nodes
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 Dell infrastructure, including VxRail and vSAN Ready Nodes.

Dell Technologies Solutions

Dell
Validated Design

DELLTechnologies

Notes, cautions, and warnings

(i) **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 vSAN-based infrastructure provide a quick and easy way to simplify and extend your VMware environment. These solutions are ideal for VDI applications because they combine compute, storage, networking, virtualization, and management, these solutions are ideal for VDI applications.

The Dell Validated Designs for VDI are built on Dell VxRail or Dell vSAN Ready Nodes. These true hyperconverged infrastructure (HCI) platforms provide performance, flexibility, and scale for VDI environments.

Dell Technologies recommends VxRail for an enhanced VDI solution that uses a wide range of software, tools, and resources co-developed by Dell Technologies and VMware. The VMware hyperconverged software is vSphere-ready and based on vSAN software-defined storage (SDS). Dell Technologies deployment and support tools integrate the software management within VxRail Manager. Data protection and replication are included and can support either hybrid or all-flash storage configurations.

vSAN Ready Nodes do not include the full automation suite that is available in VxRail, but they provide more flexibility in platform choices. vSAN Ready Nodes offer the confidence that your pre-validated configuration will work with vSAN technology as well as the VMware Horizon 8 software suite.

Installing VMware Horizon 8 with its VDI components on VxRail or vSAN Ready Nodes 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 VxRail and vSAN Ready Nodes 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 VxRail or vSAN Ready Nodes.

What's new?

Differentiated security hardening of the layers of the VDI solution stack. This is demonstrated through a combination of scripts and manually modified settings based on a combination of the following:

- Defense Information Systems Agency Security Technical Implementation Guides (STIGs)
- Microsoft Security Compliance Toolkit (SCT) Baselines

For more information, see the STIG and Microsoft Security Baseline-Based Hardening of a VMware Horizon on a VxRail-Based VDI Environment Implementation Guide.

Dell VxRail hyperconverged infrastructure on 15th Generation PowerEdge Nodes with 3rd Generation Intel Xeon Scalable Processors provides:

- Memory—Support for Intel Optane Memory 200 series (Optane configurations are not compatible with double-wide GPUs, including the A16 and the A40).
- GPU—NVIDIA A16 GPUs with support for up to 128 vGPU users per node
- Storage—NVMe vSAN Cache Tier
- Networking—Support for 4-port networking design

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 choose from two vSAN-based appliance options. While VxRail and vSAN Ready Nodes have the same underlying technology (VMware vSAN), each provides unique benefits. VxRail provide a turnkey solution for deploying and managing infrastructure with additional software features to assist throughout the product life cycle. vSAN Ready Nodes require similar architectural design elements, but do not include the software components and elements that are found in the VxRail product. However, vSAN Ready Nodes offer a greater array of platforms and additional peripheral device choices.

Dell Validated Designs for VDI offer several key benefits:

- **VxRail Lifecycle Management (VxRail LCM)**—Integrated VxRail software that keeps updates and patches current. Updates are easy to install and safe to implement because all software components are engineered, tested, and released as a bundle. The VxRail Manager software notifies administrators when updates are available and can intelligently deploy important updates in your production environment with no downtime. VxRail LCM reduces complexity and risk, and makes the infrastructure more secure.

  (i) **NOTE:** VxRail Manager is available only on VxRail.

- **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 64 nodes per cluster, supporting thousands of virtual machines (VMs). 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.
- **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 Horizon and vSAN-based solutions.
- **User experience**—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

This section provides an architecture overview and guidance on managing and scaling a VMware Horizon environment on Dell VxRail.

Solution architecture

The following figure depicts the architecture that we tested. It shows the validated solution, including the network, compute and graphics, management, and storage layers. This architecture aligns with the VMware Horizon pod and block design. A pod is divided into multiple blocks. Each block is made up of one or more vSphere clusters and a vCenter Server.
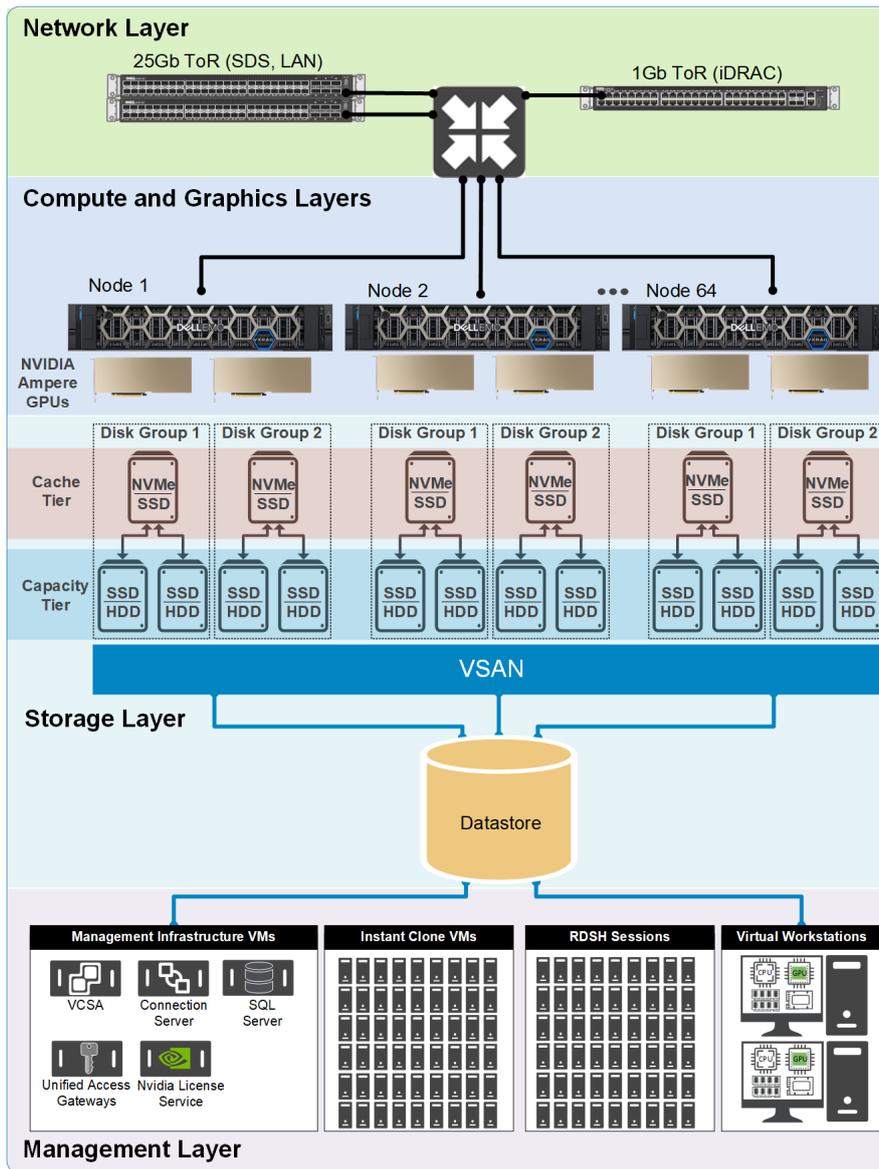
Figure 1. Validated solution architecture

Physical Architecture

There are a variety of VxRail configurations available, but this section focuses primarily on the physical architecture of the VxRail V670F, which is the recommended product for VDI deployments and is leveraged for both "Density Optimized" and "Virtual Workstation" configurations.

VxRail

Dell VxRail is available in 1U or 2U rack building blocks. It is built on VMware vSAN technology and Dell Technologies software. Add VMware Horizon Universal Subscription or Horizon Enterprise Edition (TERM) to license your VxRail environment for a full VDI deployment.

The following figure shows the VxRail components:
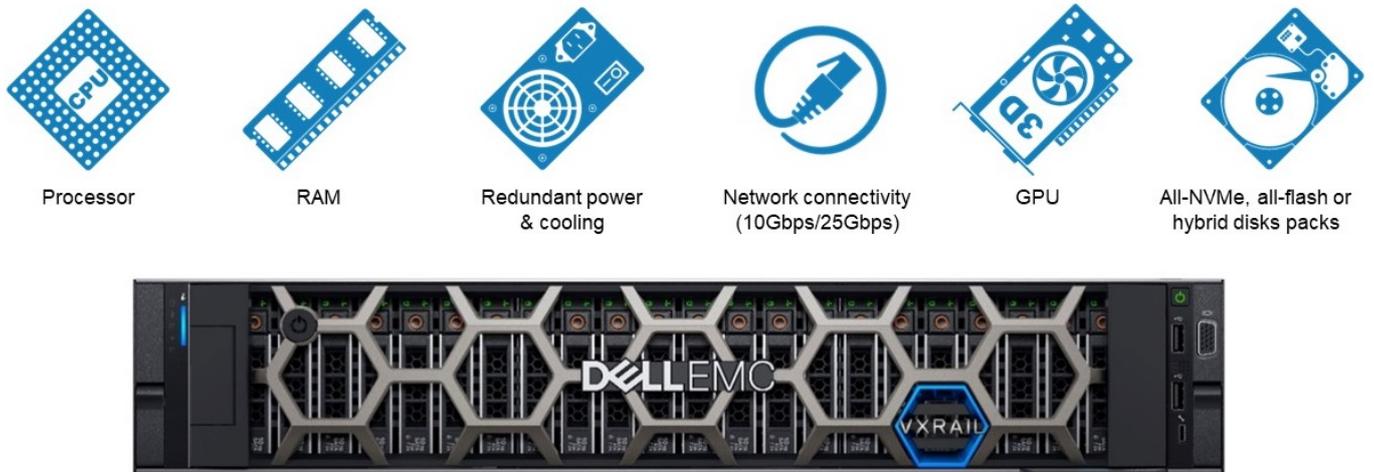
Figure 2. Dell VxRail

VxRail platforms are equipped with 3rd Generation Intel Xeon Scalable processors. You can deploy a cluster with as few as three nodes, providing an ideal environment for small deployments. To achieve full vSAN high availability (HA), the recommended starting block is four nodes. VxRail can support workloads with high storage I/O requirements using storage-dense nodes, graphics-heavy VDI workloads with GPU hardware coupled with virtual GPU software, and entry-level nodes for remote and branch office environments.

With VxRail you can start small and scale as your requirements increase. Single-node scaling and low-cost entry point options give you the freedom to buy the right amount of storage and compute resources to start and then add capacity to support growth. A single-node VxRail V670F can be configured with 16 to 80 CPU cores per node, up to 4 TB of memory (or 8 TB of tiered memory), and supports a maximum of 161 TB of storage. A 64-node all-flash cluster delivers a maximum of 5,120 cores, 256 TB of memory (or 512 TB of tiered memory), and 10,304 TB of raw storage. The following table shows the platforms that are recommended for VDI:

Table 1. VxRail platforms

Platform	Description	Configurations	Form factor
E Series	Everything platform	NVMe, all-flash, and hybrid	1U
V Series	Optimized for acceleration	All-flash only	2U
P Series	Performance focused	NVMe, all-flash	2U
S Series	Storage dense	Hybrid only	2U
D Series	Durable and rugged	All-flash and hybrid	1U

VxRail Manager

VxRail Manager, which is available on VxRail only, is the primary deployment and element manager interface for VxRail. VxRail Manager simplifies the entire life cycle from deployment through management, scaling, and maintenance. It also enables single-click upgrades and dashboard monitoring for health, events, and physical views.

VxRail V670F Front View

The following diagram shows the VxRail V670F front view with cache and capacity storage tiers:
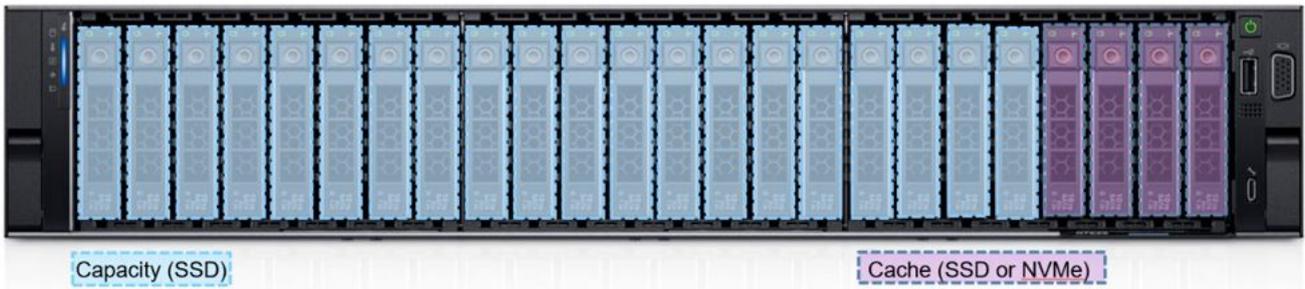
Capacity (SSD)   Cache (SSD or NVMe)

Figure 3. VxRail V670F Front View

VxRail V670F Rear View

The following diagram shows the VxRail V670F rear view with network, I/O, and GPU options:

Boot: 2 x M.2 (BOSS-S2 RAID 1)   Network Adapters   Nvidia A40 or A16 GPU

Slot 1 | Slot 2 | Slot 3 | OCP 3.0 | Slot 7 | Slot 8

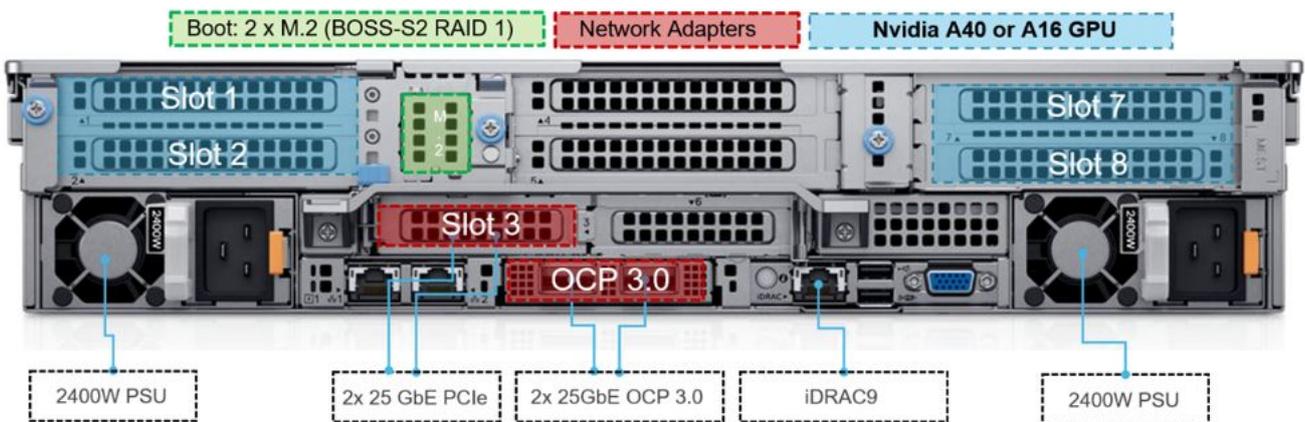2400W PSU | 2x 25 GbE PCIe | 2x 25GbE OCP 3.0 | iDRAC9 | 2400W PSU

Figure 4. VxRail V670F Rear View

Dell vSAN Ready Nodes

Dell vSAN Ready Nodes are pre-validated Dell server configurations that reduce the complexity of deploying vSAN technology. vSAN Ready Node configurations are listed in the vSAN Ready Node compatibility guide on the VMware website. vSAN Ready Nodes provide more customized hardware and platforms to meet your data center needs.

Because vSAN Ready Nodes do not offer life cycle management and additional bundled software, Dell Technologies recommends customized deployment services and at least three years of ProSupport Plus. Add VMware Horizon Universal Subscription or Horizon Enterprise Edition (TERM) to license your Dell vSAN Ready Nodes for a full VDI deployment.

VDI-optimized configurations

Dell Technologies recommends VDI-optimized 2U/1 Nodes that support GPU hardware for graphics-intensive desktop deployments.

The VxRail V Series and vSAN Ready Node R750 server can be configured with or without GPUs. Dell Technologies also offers similar configurations in 1U/1 Nodes, although graphics configurations are limited on these platforms.

We have designated common configurations as "Management-Optimized," "Density-Optimized," and "Virtual Workstation." These designations are referenced throughout the documentation. The following table shows the common configurations:

Table 2. Common configurations

Configuration	CPU	RAM	Disk	GPU (optional)	Description
Management-Optimized	1 x Intel Xeon Gold 6330 (28-core @ 2.0 GHz)	256 GB (8 x 32 GB @ 2,933 MHz)	4 TB + (capacity)	None	Offers density and value to provide a dedicated

Table 2. Common configurations (continued)

Configuration	CPU	RAM	Disk	GPU (optional)	Description
					environment to deploy virtualized management infrastructure.
Density-Optimized	2 x Intel Xeon Gold 6348 (28-core @ 2.6 GHz)	1024 GB (16 x 64 GB @ 3,200 MHz)	8 TB + (capacity)	Up to 2 x full length, dual width (FLDW) 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	2 x Intel Xeon Gold 6354 (18-core @ 3.0 GHz)	512 GB (16 x 32 GB @ 3,200 MHz)	6 TB + (capacity)	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 workloads.

NVIDIA GPUs

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

- NVIDIA A40—NVIDIA A40 GPUs provide a leap 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—NVIDIA A16 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.

Physical network components

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)—The S5248F-ON 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.
- Dell Networking S4048-ON (10 GbE ToR switch)—The S4048-ON 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.

See Dell PowerSwitch S Series 10GbE Switches and Dell PowerSwitch S Series 25/40/50/100 GbE Switches for more information about these switches.

Networking

VxRail 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 based on 40 GbE and non-blocking switches. This architecture maintains consistent performance without any throughput reduction.

File workload

The increased growth in the amount of data that is stored in file shares and user home directories across IT environments in recent years has resulted in an increased focus on the need to better manage 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 storage

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/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 file storage

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. Thus, 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.

Data center infrastructure

Enterprise equipment requires power to operate, racks for streamlined management, and cooling to maintain reliable operations.

You must carefully select infrastructure solutions that provide these capabilities to ensure uptime, scalability, energy efficiency, and ease of management. Dell Technologies provides a wide range of data center infrastructure solutions such as:

- Dell Netshelter SX racks—Deploy server, storage, and networking equipment and other IT hardware while optimizing power, cooling, cabling, and systems management.
- Dell Keyboard Video Mouse (KVM) and Keyboard Monitor Mouse (KMM) solutions—Manage 8 to 1,024 local and remote servers running various operating systems across the enterprise.
- Dell Smart-UPS—Deliver reliable power and protect IT equipment, including servers, storage, networking, point-of- sale, and medical equipment.
- APC Rack Power Distribution Units (PDUs)—Provides reliable power distribution that is designed to increase manageability and efficiency in your data center.

Client components

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

Table 3. Recommended client components

Component	Description	Recommended use	More information
Latitude laptops and 2-and-1s	- Biggest screens in a smaller footprint with a wide array of	- Mobility and space-saving devices	www.delltechnologies.com/ Latitude

Table 3. Recommended client components (continued)

Component	Description	Recommended use	More information
	- ports to connect peripherals and enjoy speakerphone experience - More responsive apps with Dell Optimizer and intelligent audio for 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	- Allows users to be productive and stay connected with versatile, space-saving mobile solutions - Offers a modern portfolio built to prioritize customer experience and keep employees productive wherever they work with a selection of laptops, 2-in-1s, and ecosystem products	
OptiPlex business desktops and All-in-Ones	- Intel 9th Gen 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
Precision workstations	- The most complete workstation portfolio with towers, racks, and mobile form factors - Extremely 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

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 clone technology

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 completely removed in Horizon 8 2012.

vSAN software-defined storage

vSAN is available in hybrid or all-flash configurations depending on the platform.

After vSAN is enabled on a cluster, all disk devices that are presented to the hosts are pooled to create a shared data store that is accessible by all hosts in the VMware vSAN cluster. You can then create VMs with storage policies assigned to them. The storage policy determines availability, performance, and sizing.

vSAN provides these configuration options:

- All-flash configuration—Uses flash for both the cache tier and capacity tier to deliver enterprise performance and a resilient storage platform. In this configuration, the cache tier is fully dedicated to writes, allowing all reads to come directly from the capacity tier. This model allows the cache device to protect the endurance of the capacity tier. All-flash configured solutions enable data deduplication features to extend the capacity tier.
- Hybrid configuration—Uses flash-based devices for the cache tier and magnetic disks for the capacity tier. Hybrid configurations are ideal for clients looking for higher volume in the capacity tier. The performance of SSD and magnetic spinning disks is comparable in VDI applications.

NVIDIA vGPU

NVIDIA vGPU is the industry's most advanced technology for virtualizing true 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 three 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.

Validation

Topics:

Validation

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.

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 4. 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
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 1 hour.
- All desktops were started before users were logged in.

Login VSI workloads

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

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.

Desktop VM test configurations

The following table summarizes the desktop VM configurations used for the Login VSI workloads that we tested:

Table 5. Desktop VM test configurations

Login VSI workload	vCPUs	ESXi configured memory	ESXi reserved memory	Screen resolution	Operating system
Task Worker	8	32 GB	32 GB	1280 x 720	Windows Server 2019 64-bit
Knowledge Worker	2	4 GB	2 GB	1920 x 1080	Windows 10 Enterprise 64-bit

Summary of test 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 6. Summary of test results

Server configuration	Login VSI workload	Operating system	User density	Average CPU	Average GPU	Average active memory	Average IOPS per user	Average network Mbps per user
Density Optimized	Task Worker	Server 2019, 1809	323	86%	N/A	106 GB	2.1	0.5 Mbps
Density Optimized	Knowledge Worker	Windows 10, 21h1	230	84%	N/A	271 GB	8	1.8 Mbps

> (i) **NOTE:** VM densities in excess of 200 were used for benchmarking purposes only. A limit of 200 VMs per vSAN-enabled host should be used in a production environment. For additional information, see VMware Configuration Maximums.

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 period. 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.

Overview

We conducted PAAC testing on this solution using the Login VSI load- generation tool. Login VSI is an industry-standard tool for benchmarking VDI workloads. It uses a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load testing.

Login VSI Task Worker

We performed this test with the Login VSI Task Worker workload on a 3-node VxRail cluster (see Table 2). We created the RDSH VMs using VMware Horizon instant clone technology and used the VMware Horizon Blast Extreme display protocol. We populated each compute host with 12 RDSH VMs.

CPU usage

The following graphs show the CPU utilization across the three hosts during testing. CPU usage with all VMs powered on was approximately 1.8 percent before the test started. The CPU usage steadily increased during the login phase, as shown in the following figure.
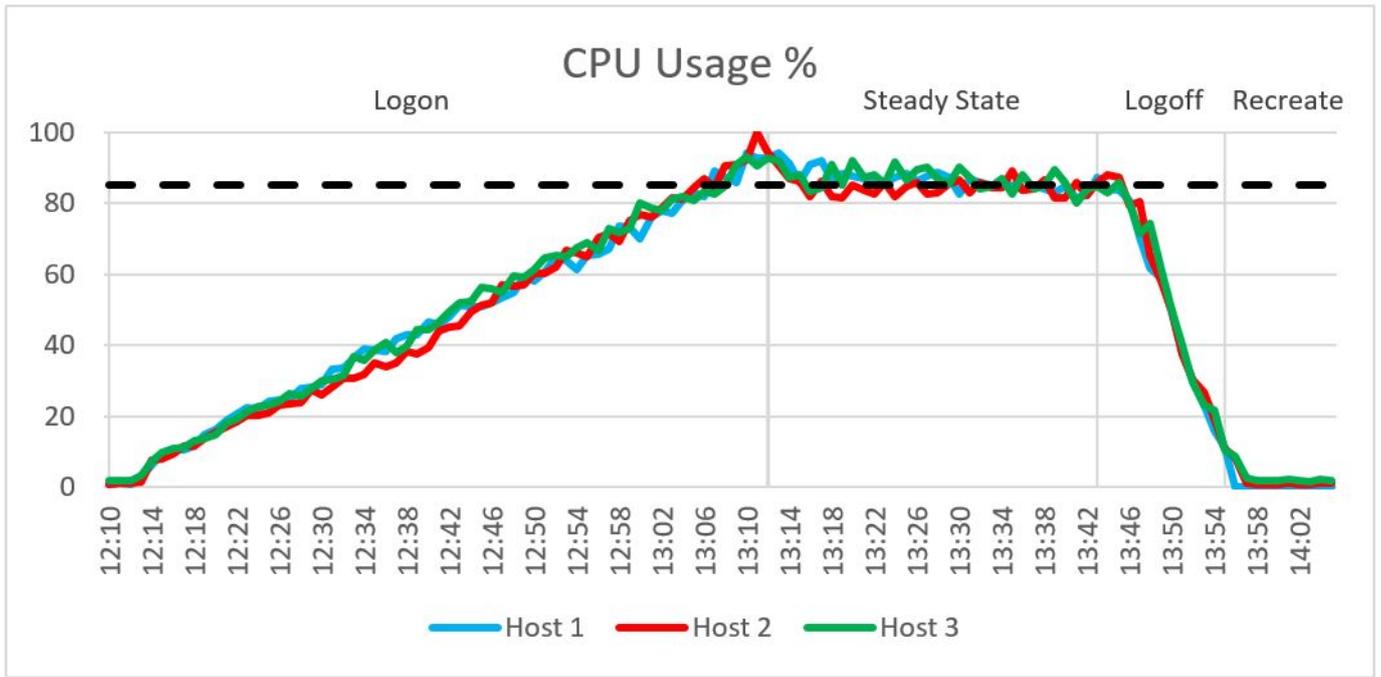
Figure 5. CPU usage

During the steady state phase, an average CPU utilization of 86 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. The average steady state CPU core utilization across the four hosts was 69 percent, as shown in the second figure.
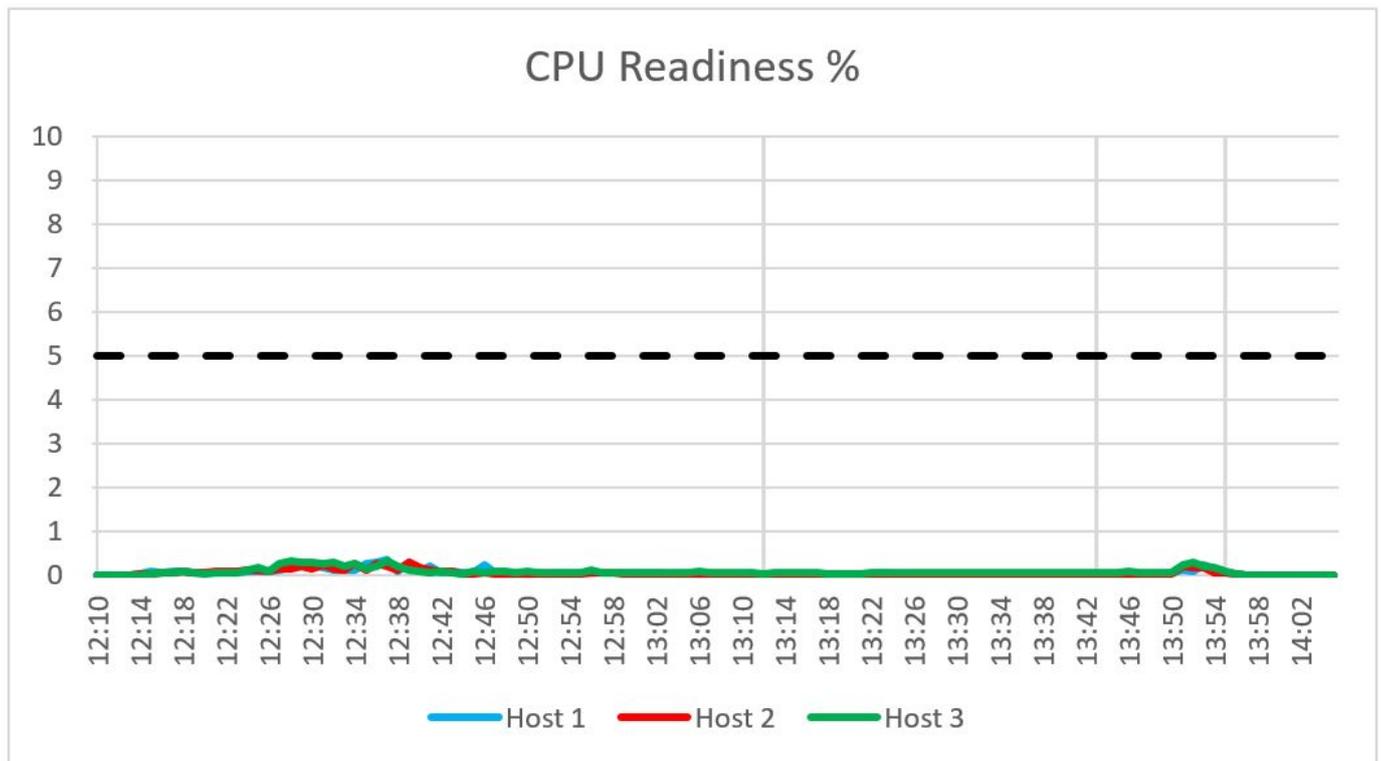
Figure 6. CPU readiness
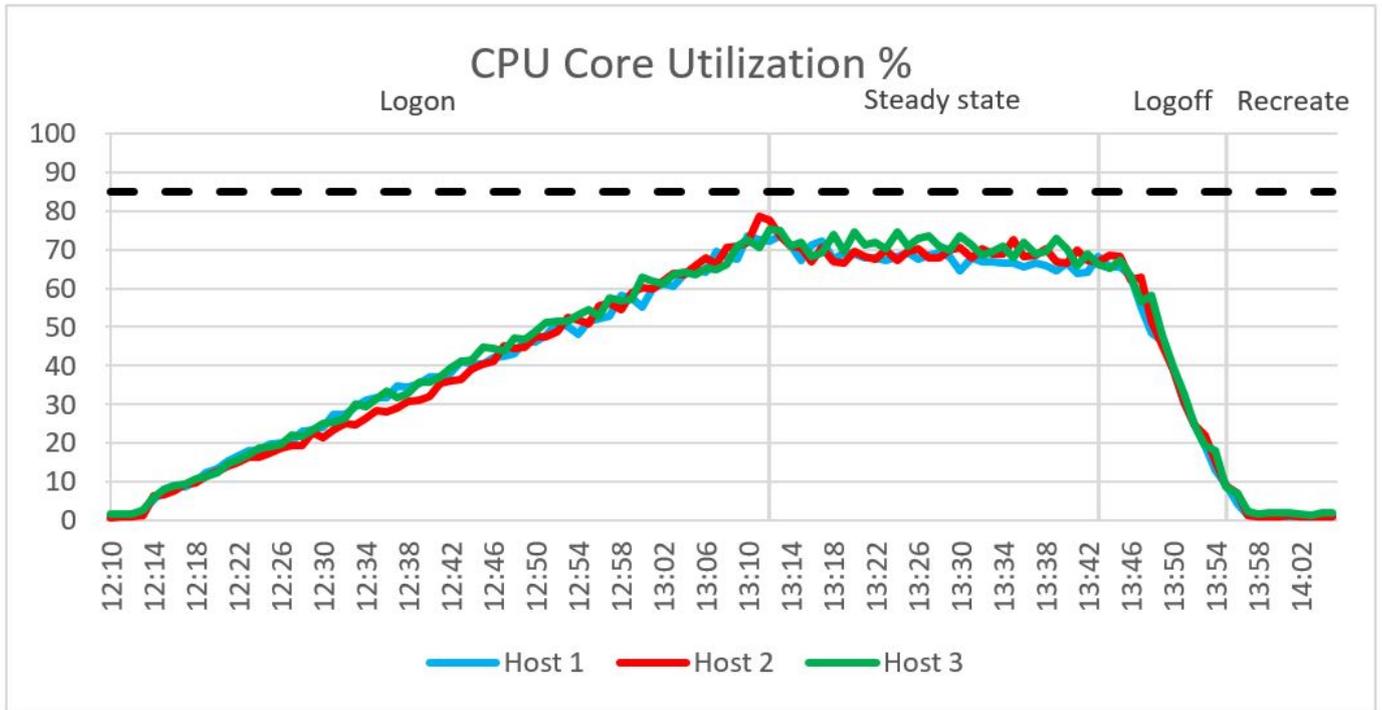
Figure 7. CPU core utilization

Memory

We observed no memory constraints during the testing on the compute hosts. Out of 1024 GB of available memory per node, the compute host reached a maximum consumed memory of 303 GB and a steady state average of 303 GB. Active memory usage reached a maximum active memory of 170 GB and recorded a steady state average memory of 109 GB. There was no memory ballooning or swapping on the hosts. The following figures show consumed and active memory.
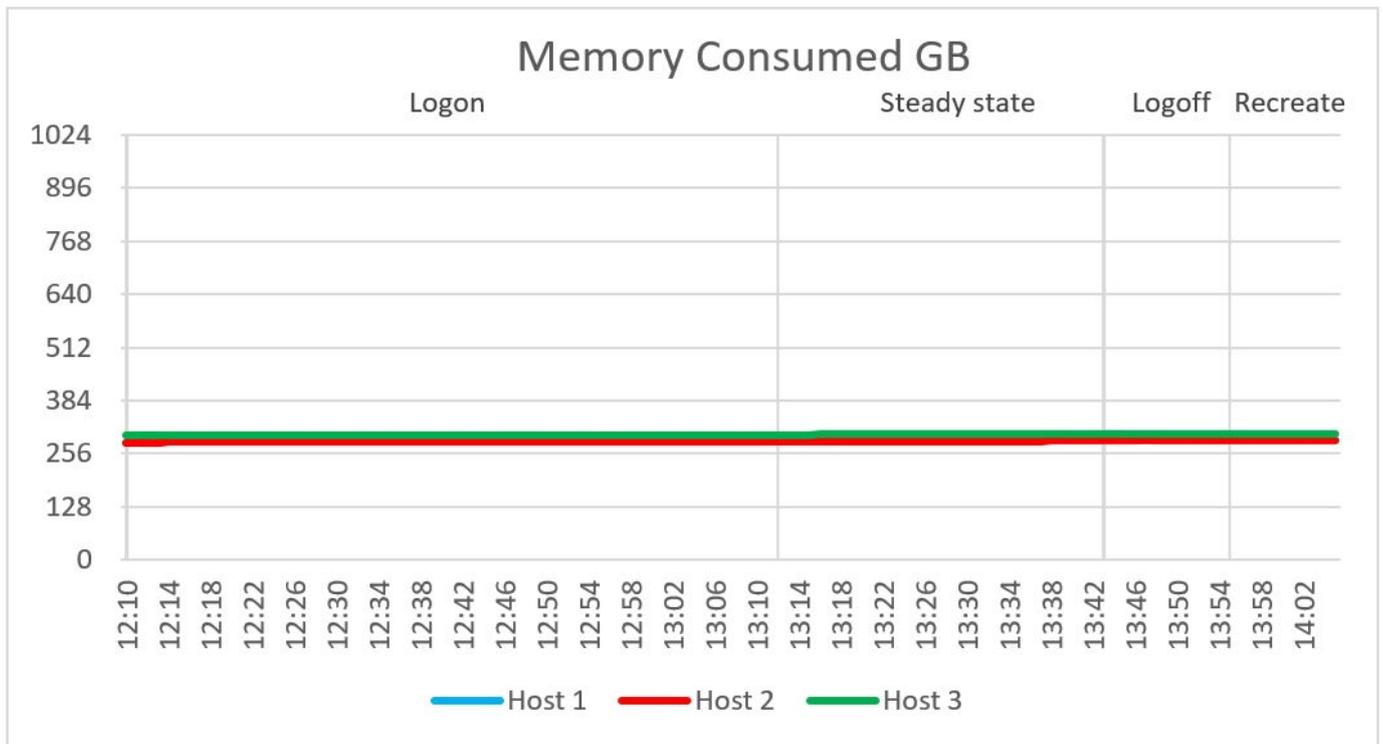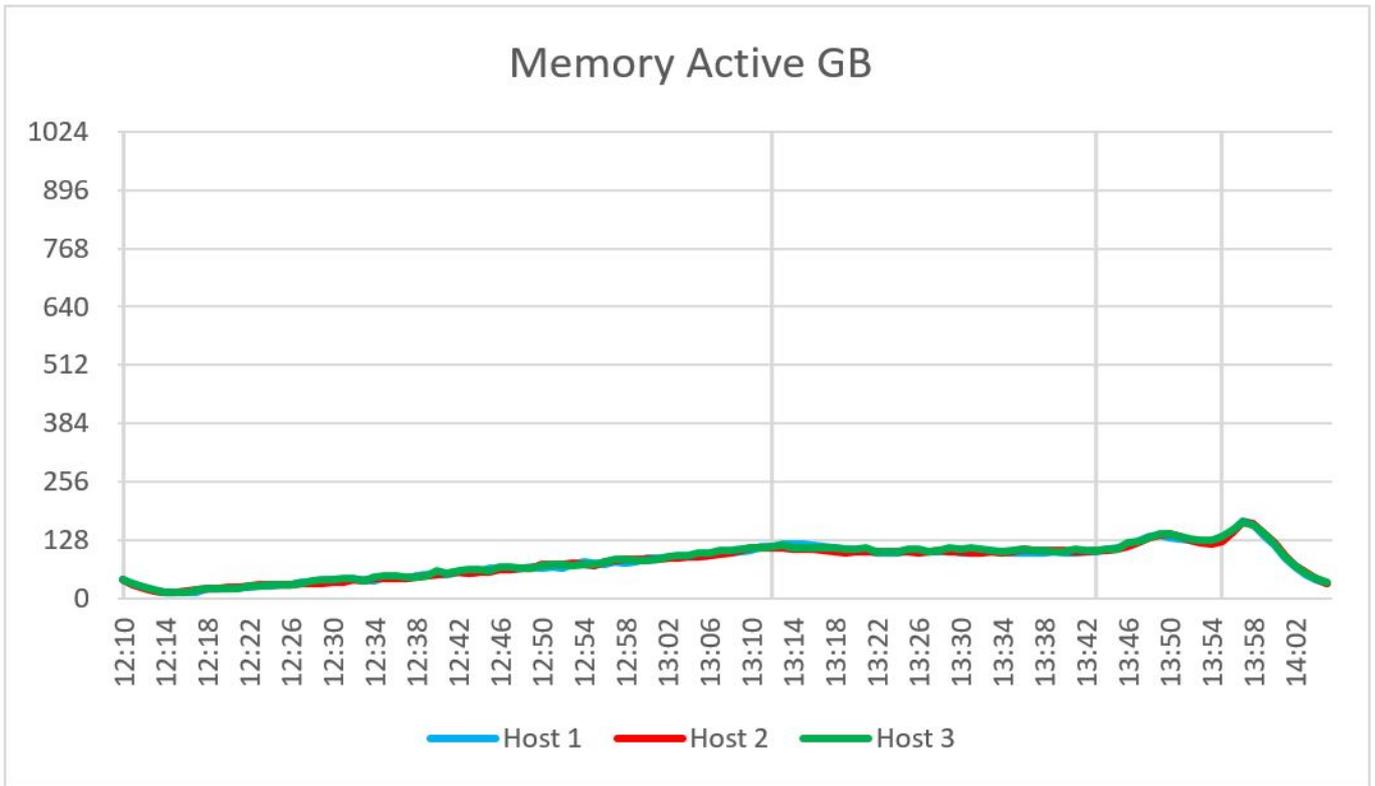

Figure 8. Memory consumed

Figure 9. Memory active

Network usage

Network bandwidth was not an issue during testing. The network usage recorded a steady state average of 563 Mbps. The busiest period for network traffic was immediately after all users had logged in when a peak value of 1,333 Mbps was recorded. The following figure shows network usage:
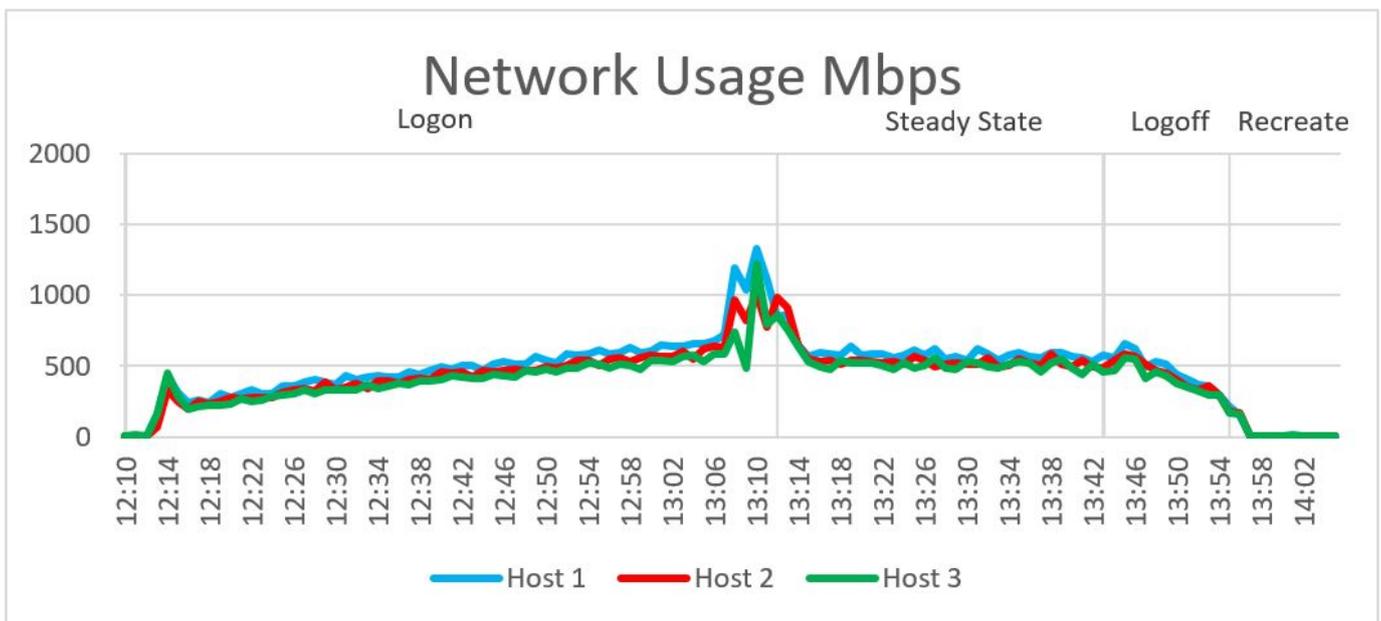

Figure 10. Network usage

VxRail cluster IOPS

Cluster IOPS reached a maximum value of 186 for read IOPS and 1,550 for write IOPS during the steady state phase. The average steady average IOPS were 1,406. The following figure shows cluster IOPS:
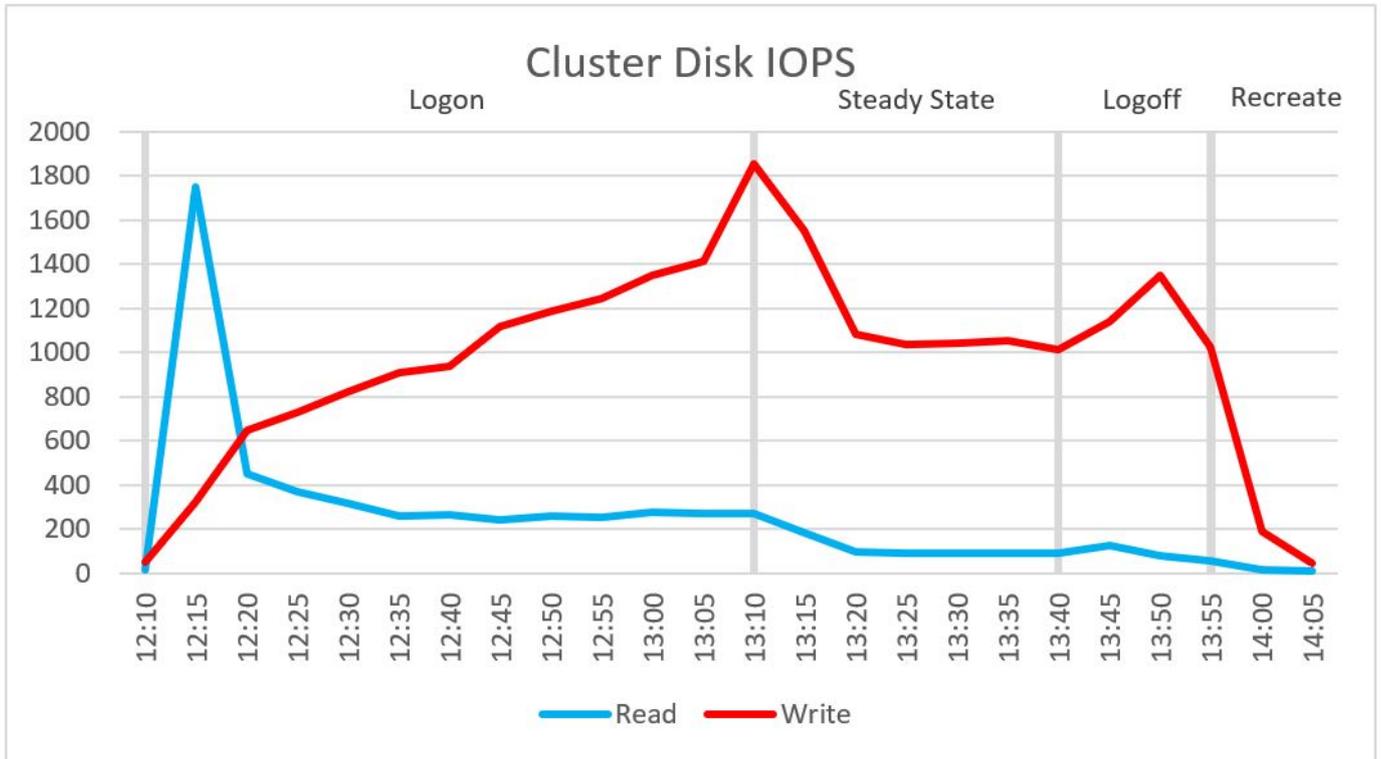

Figure 11. Cluster IOPS

VxRail disk I/O Latency

Cluster disk latency reached a maximum read latency of 0.309 milliseconds and a maximum write latency of 0.604 milliseconds during the steady state phase. The average steady state latency was 0.44 milliseconds. The following figure shows disk I/O latency:
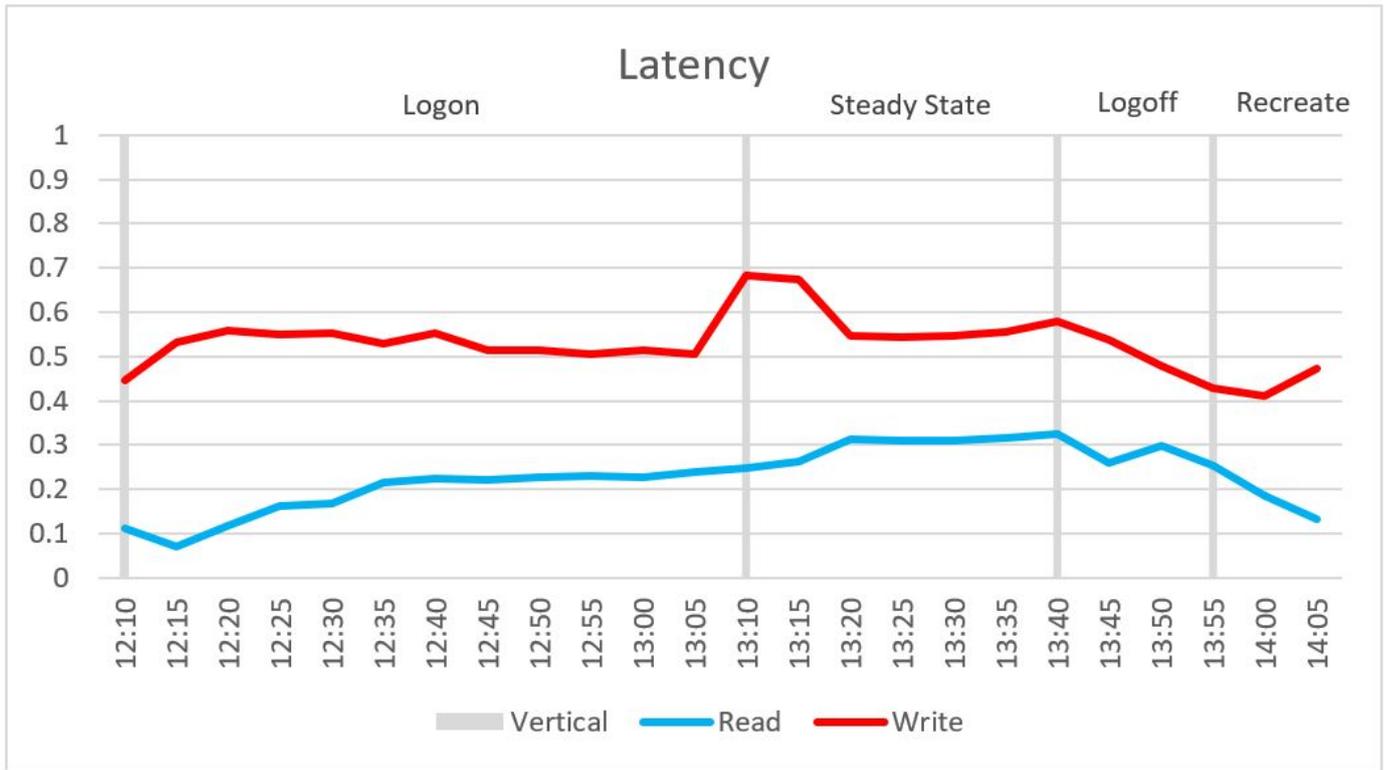
Figure 12. Disk I/O Latency

User experience

The baseline score for the Login VSI test was 587. This score falls in the 0 to 799 range rated as "Very Good" 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 1,292 when 969 sessions were loaded. This value is well below the VSI threshold score of 1,588 set by the Login VSI tool. During the 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.
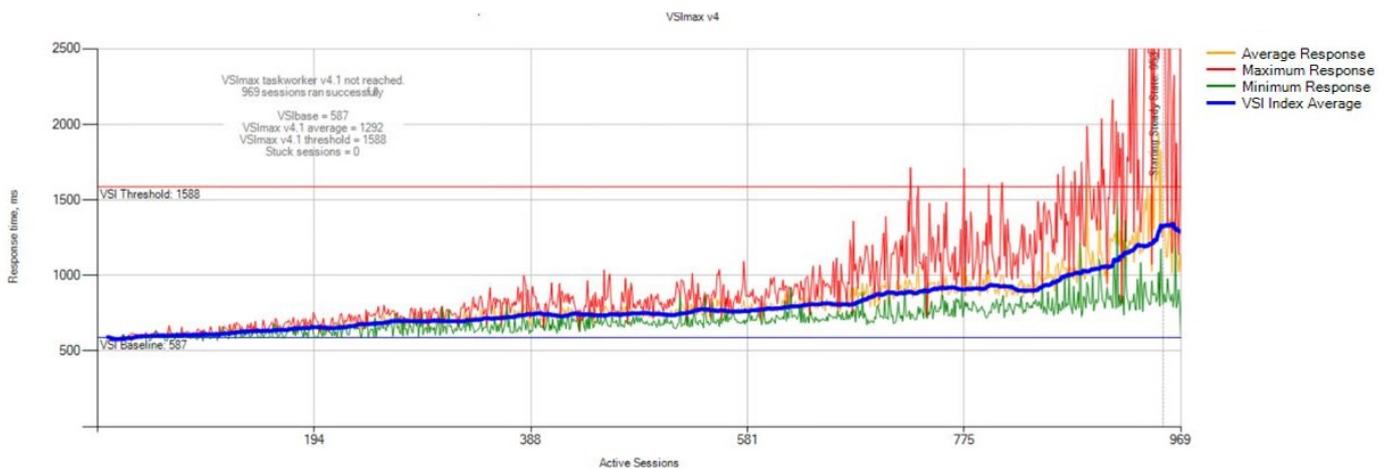

Figure 13.

Login VSI Knowledge Worker

We performed this test with the Login VSI Knowledge Worker workload on a 3-node VxRail cluster (see Table 2). We created the desktop VMs using VMware Horizon instant clone technology and used the VMware Horizon Blast Extreme display protocol. We populated each compute host with 230 desktop VMs.

CPU usage

The following graphs show the CPU utilization across the three hosts during the testing. CPU usage with all VMs powered on was approximately 7 percent before the test started. The CPU usage steadily increased during the login phase, as shown in the following figure.
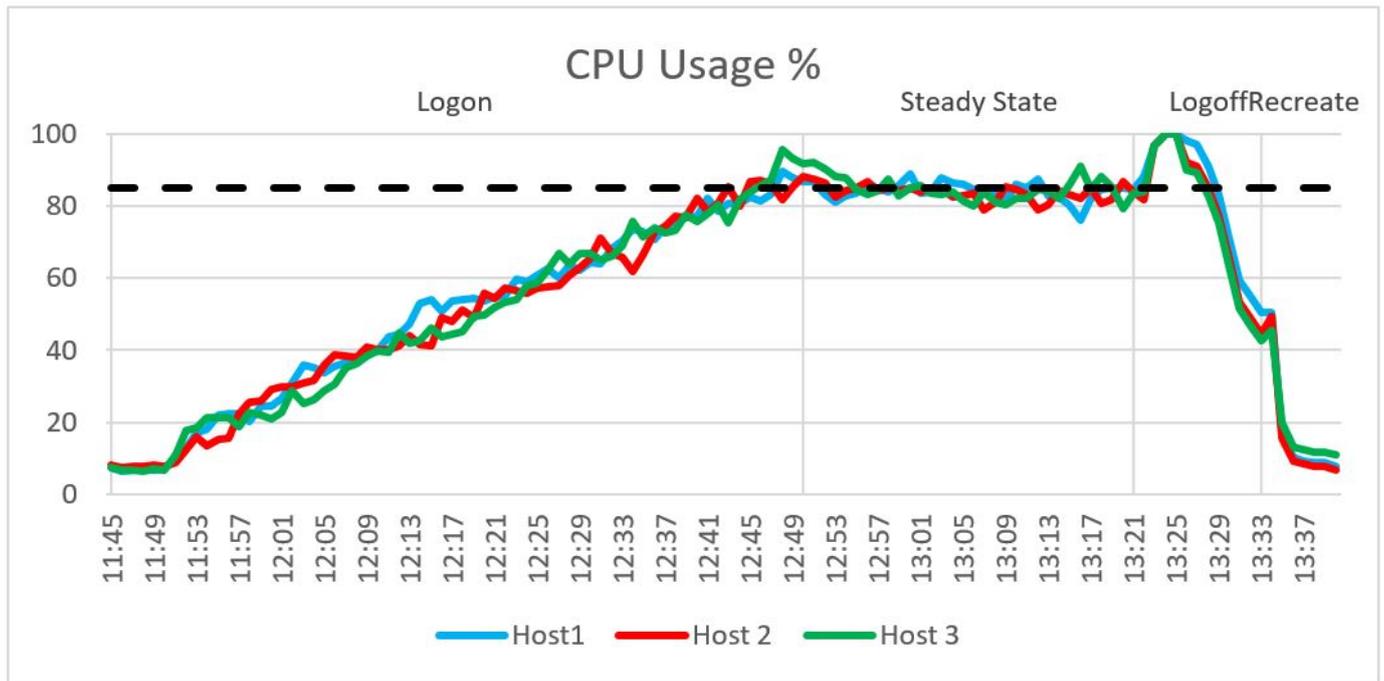

Figure 14. CPU usage

During the steady state phase, an average CPU utilization of 84 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. The average steady state CPU core utilization across the four hosts was 72 percent, as shown in the second figure.
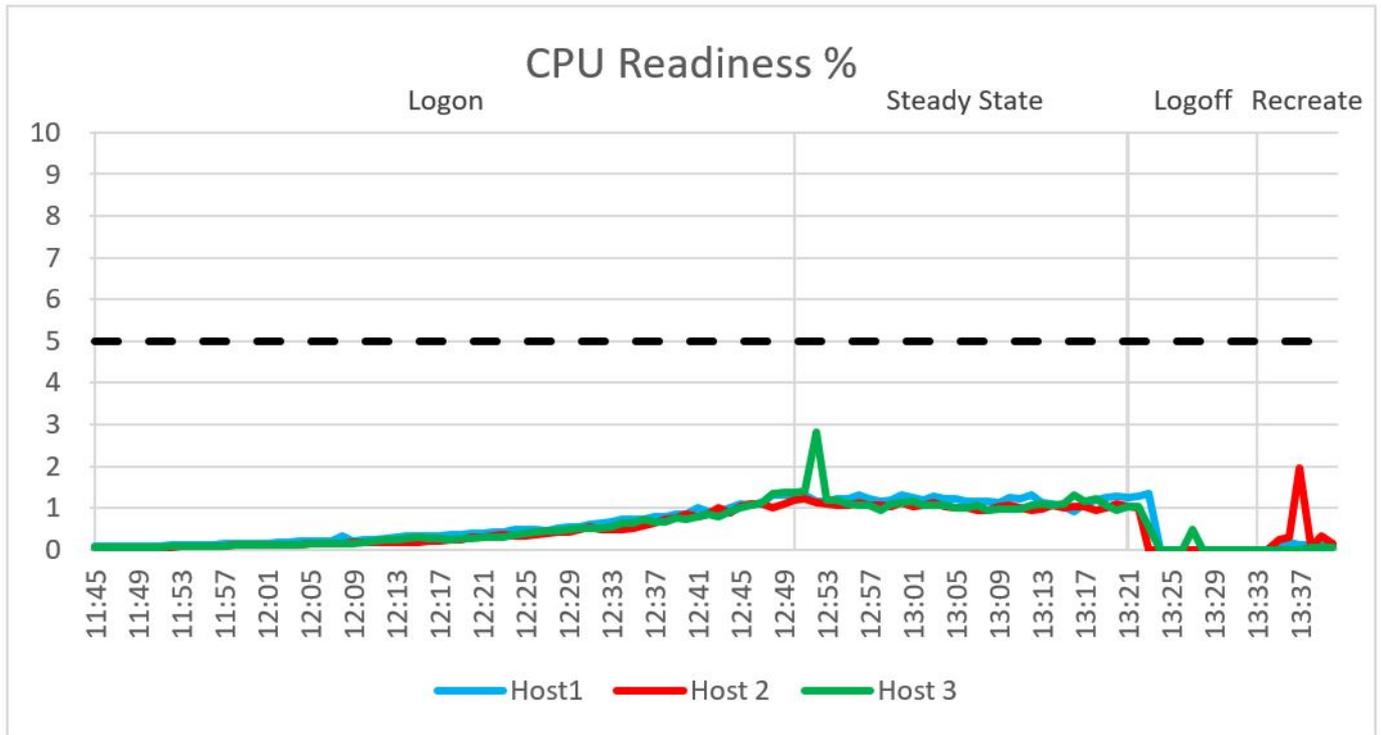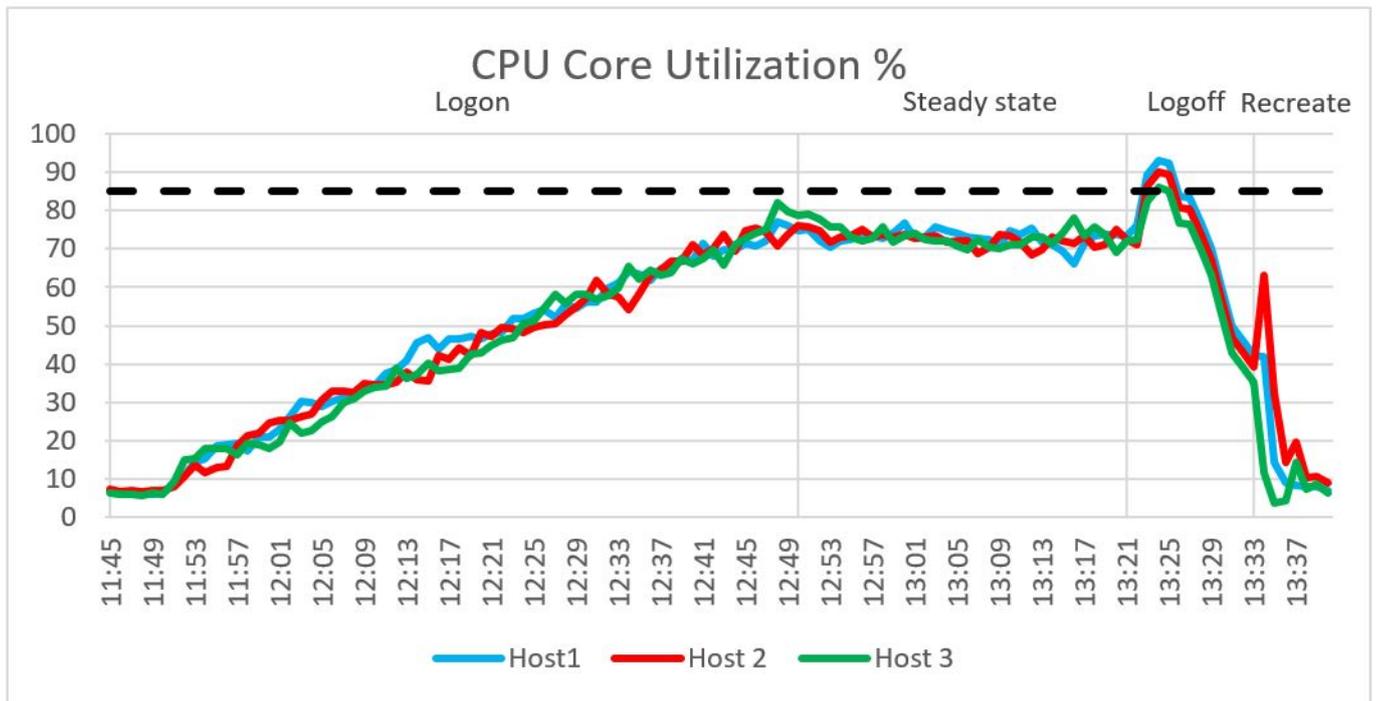
Figure 15. CPU readiness

Figure 16. CPU core utilization

Memory

We observed no memory constraints during the testing on the compute hosts. Out of 1024 GB of available memory per node, the compute host reached a maximum consumed memory of 977 GB and a steady state average of 969 GB. Active memory

usage reached a maximum active memory of 540 GB and recorded a steady state average memory of 260 GB. There was no memory ballooning or swapping on the hosts. The following figures show consumed and active memory.
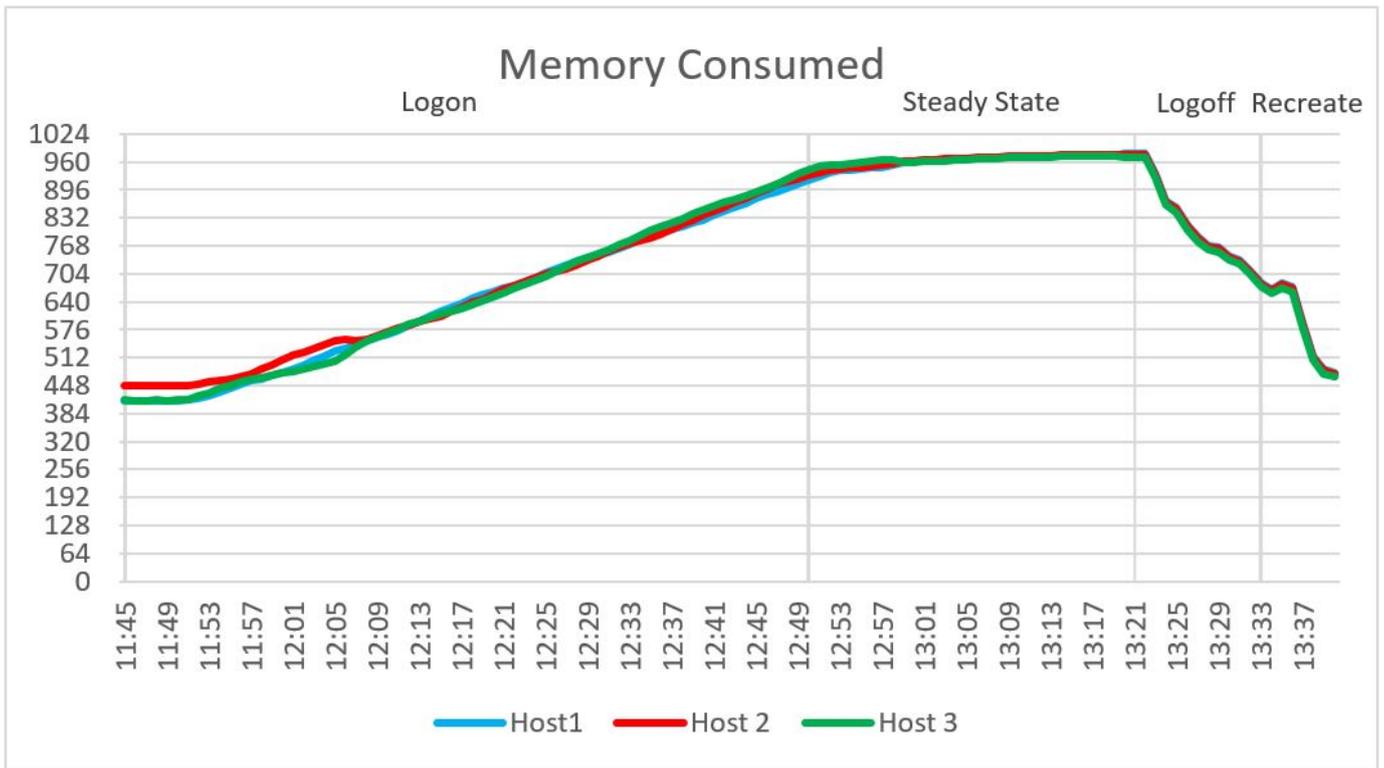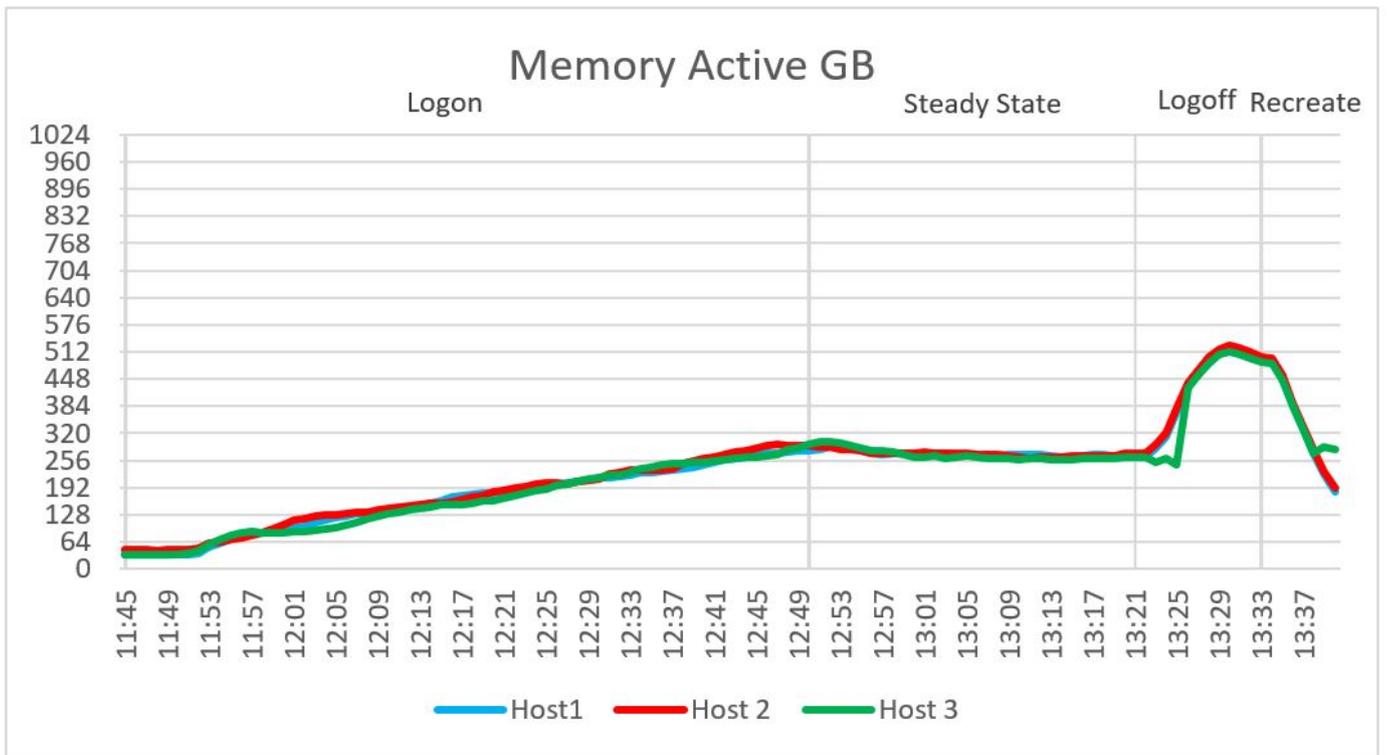

Figure 17. Memory consumed

Figure 18. Memory active

Network usage

Network bandwidth was not an issue during testing. The network usage recorded a steady state average of 1286 Mbps. The busiest period for network traffic was during the re-create phase when a peak value of 10,000 Mbps was recorded. The following figure shows network usage:
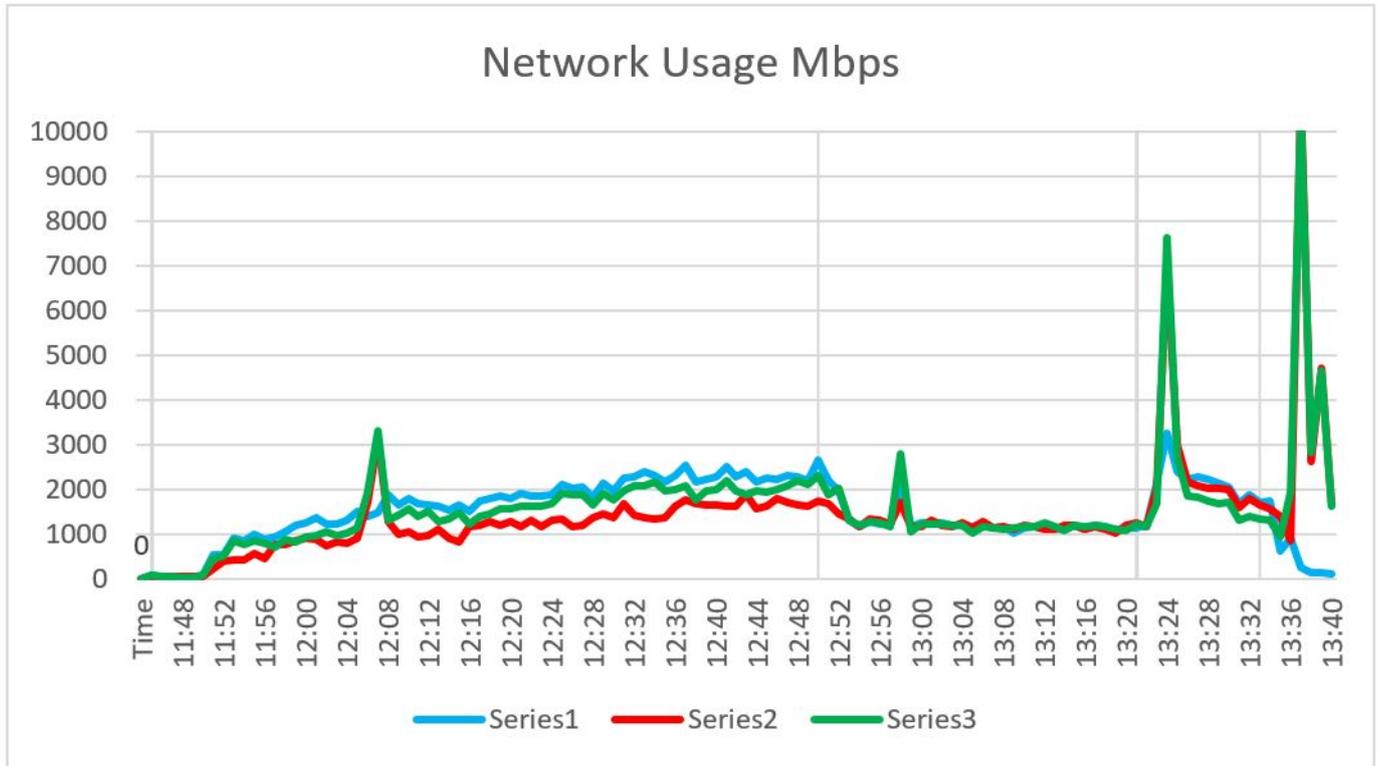

Figure 19. Network usage

VxRail cluster IOPS

Cluster IOPS reached a maximum value of 1,458 for read IOPS and 3,039 for write IOPS during the steady state phase. The average steady average IOPS were 2,109. The following figure shows cluster disk IOPS:
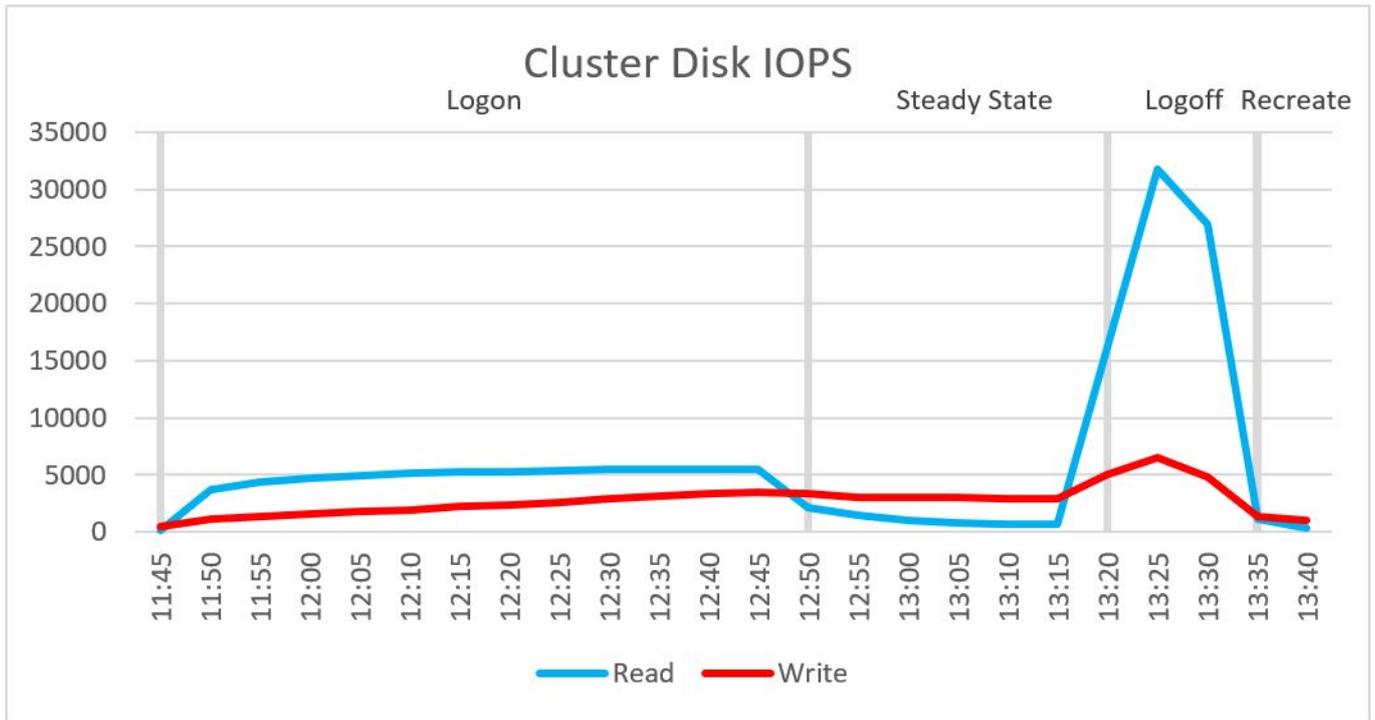
Figure 20. Cluster disk IOPS

VxRail disk I/O latency

Cluster disk latency reached a maximum read latency of 1.187 milliseconds and a maximum write latency of 0.645 milliseconds during the steady state phase. The average steady state latency was 0.39 milliseconds. The following figure shows disk I/O latency:
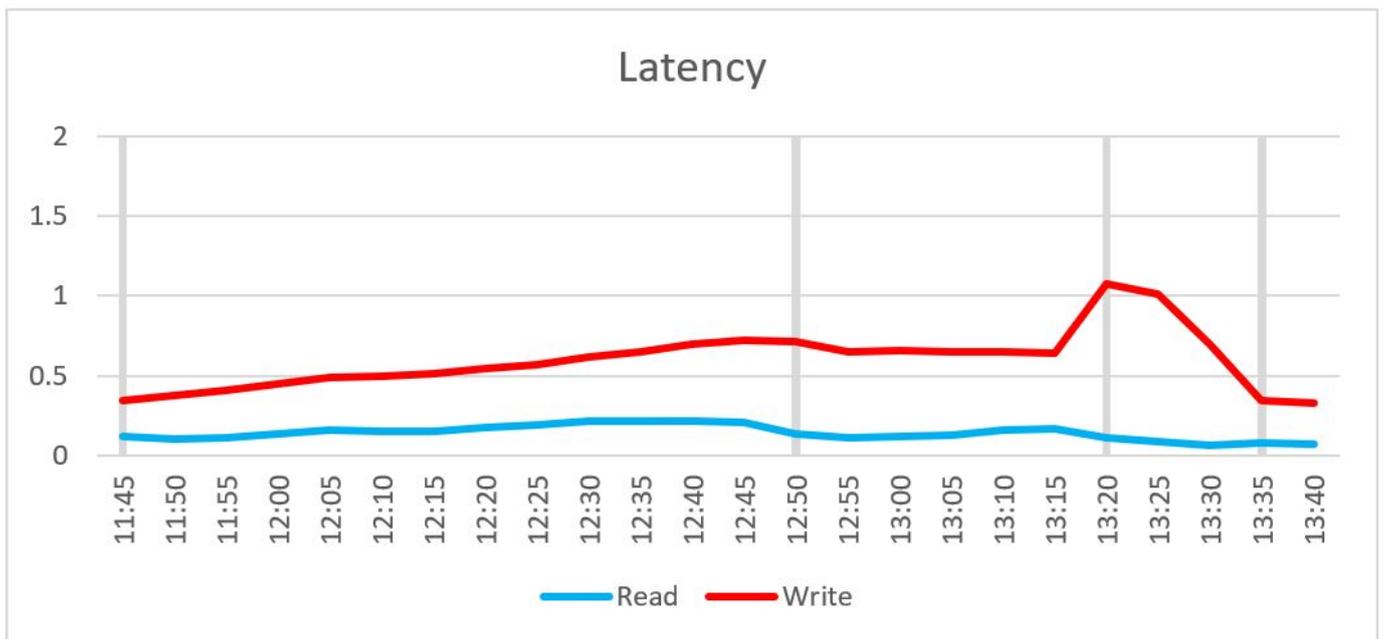

Figure 21. Disk I/O latency

User experience

The baseline score for the Login VSI test was 696. This score falls in the 0 to 799 range rated as "Very Good" 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 970 when 690 sessions were loaded. This value is well below the VSI threshold score of 1,696 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.
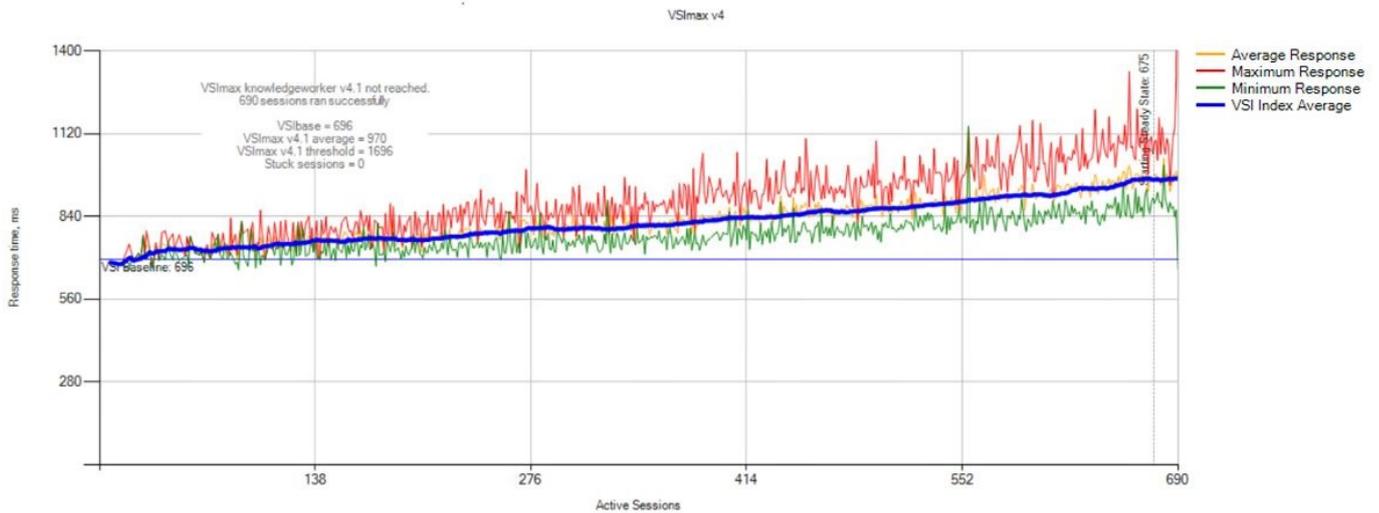

Figure 22. User experience

Summary of findings

Overview

We have carried out extensive performance testing and provided results and guidance based on the PAAC testing carried out with the Login VSI Task worker and Knowledge Worker workloads. The 3rd Generation Intel Xeon Scalable processors in our Density Optimized configuration provide performance, density, and agility for your VDI workloads.

The configurations for VxRail are optimized for VDI. We selected the memory and CPU configurations that provide optimal performance. You can change these configurations to meet your environmental requirements, but changing the memory and CPU configurations from those that have been validated in this document will affect the user density per host. Work with your account team to size the solution and reserve resources for management tools when designing your VDI environment.

All Dell Validated Designs for VDI are configured to produce similar results. You can be sure that the vSAN-based appliances you choose have been designed and optimized for your organization's needs.

User density recommendations

The recommended user densities in the following table were achieved during performance testing on VxRail. We followed the VMware best practices of FTT = 1 and maintained a reserved capacity of at least 30 percent. Login VSI Task Worker was tested with Server 2019 and Microsoft Office 2019 while all other configurations were tested with Microsoft Windows 10, 64-bit, and Microsoft 365 with Office Apps. We implemented all mitigations to patch known vulnerabilities at the hardware, firmware, and software levels to ensure an improved performance impact, which is reflected in the achieved user densities.

Table 7. Recommended user densities

Server configuration	Workload	Windows version	User density
Density Optimized	Login VSI Task Worker	Server 2019, 1809	323
Density Optimized	Login VSI Knowledge Worker	Windows 10, 21h1	230

Sizing the Solution

Topics:

Sizing and scaling overview

vSAN-based solutions 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 configuration provides a good balance of performance and scalability for various general-purpose VDI workloads.
- The Virtual Workstation configuration provides 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 Ice Lake processors:
  - Task workers—5.8 users per core. For example, 93 power users with dual eight-core processors
  - Knowledge workers—4.1 users per core. For example, 66 knowledge users with dual eight-core processors
- 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.

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, such as swapping and ballooning of memory. Overcommitted memory can also affect storage performance when swap files are created.
- 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 is 16 DIMMs (8 per processor) with Intel Xeon Scalable processors.

● 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 following 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.

External vCenter considerations

When using an external vCenter, the life cycle of the vCenter appliance is not managed by VxRail and must be managed manually. Before upgrading the VxRail clusters, ensure that the vCenter is upgraded to a supported version in accordance with the VxRail and external vCenter interoperability matrix. For additional information about the procedure to update the external vCenter appliance, see the VxRail: How to upgrade external vCenter appliance Knowledge Base article (login required).

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 our small test environment, we used 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 E-Series VxRail or the PowerEdge R650 platform to reduce the data center footprint. With a more easily configured platform, the V-Series VxRail or PowerEdge 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 two-NIC configuration per appliance with all the traffic separated logically using VLAN.

● FTT—Dell Technologies recommends sizing storage with NumberOfFailuresToTolerate (FTT) set to 1, which means that you must double the amount of total storage to accommodate the mirroring of each VMDK.

● Capacity Reserve—With the release of vSAN 7 Update 1, the previous recommendation of reserving 30 percent of slack space has been replaced with a dynamic recommendation that depends on the cluster size, the number of capacity drives, disk groups, and features in use. New features such as "Operations reserve" and "Host rebuild reserve" can be optionally enabled to monitor the reserve capacity threshold, generate alerts when the threshold is reached, and prevent further provisioning. Dell Technologies recommends reviewing VMware's About Reserved Capacity documentation to fully understand the changes and new options available.

● All-Flash compared with hybrid:
   ○ Hybrid and all-flash configurations have similar performance results in the VDI environment under test. Because hybrid configurations uses spinning drives, consider the durability of the disks.
   ○ Only all-flash configurations offer deduplication and compression for vSAN. Dell Technologies recommends all-flash configurations for simplified data management.

○ All-flash configurations need considerably less storage capacity than hybrid configurations to produce similar FTT, as shown in the following table:

Table 8. FTT comparisons

VM size	FTM	FTT	Overhead	Configuration	Capacity required	Hosts required
50 GB	RAID-1 (Mirrored)	1	2 x	Hybrid	100 GB	3
50 GB	RAID-5 (3+1) (Erasure coding)	1	1.33 x	All-flash	66.5 GB	4
50 GB	RAID-1 (Mirrored)	2	3 x	Hybrid	150 GB	4
50 GB	RAID-6 (4+2) (Erasure coding)	2	1.5 x	All-flash	75 GB	6

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

Scaling guidelines

vSAN-based solutions 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.

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.
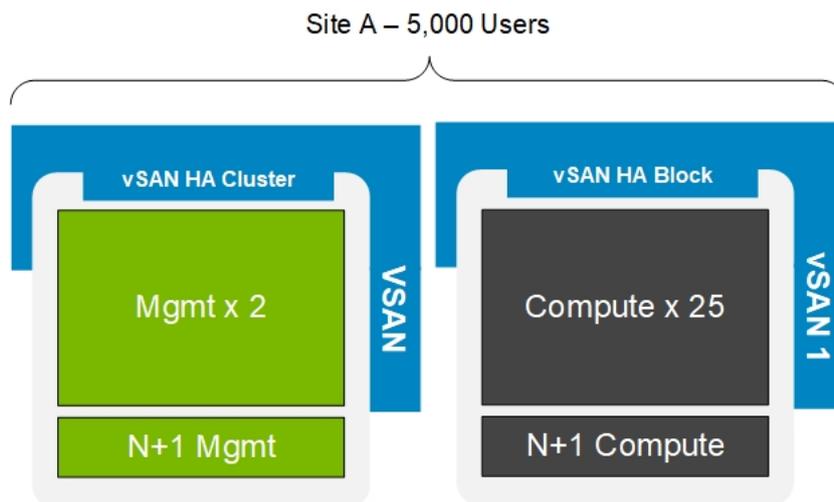
Figure 23. Single 5,000-user block

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.
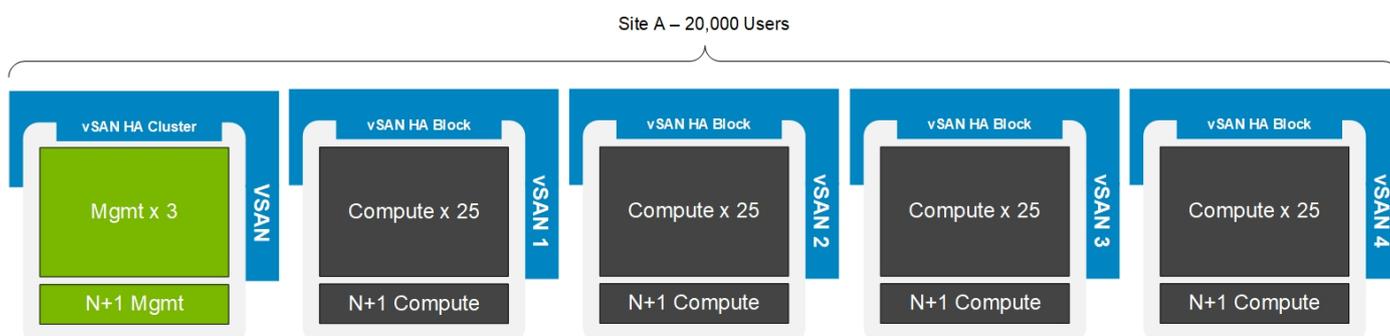
Figure 24. Scale-out for 20,000 users

Scaling up

Dell Technologies recommends a validated disk configuration for general-purpose VDI. These configurations leave drive slots available for future vertical expansion and ensure that you protect your investment as new technology transforms your organization.

ⓘ **NOTE:** These configurations can accept additional or faster processors or memory than the guidance provided in this document.

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

Securing the solution

Importance of security

As customers increasingly deploy VDI environments to meet their "do anything from anywhere" requirements, security is one of the attributes of VDI that is becoming attractive for these customers, since the ability to retain user data in the data center offers significant risk mitigation in relation to data loss.

For more information, see the STIG and Microsoft Security Baseline-Based Hardening of a VMware Horizon on a VxRail-Based VDI Environment Implementation Guide. That document demonstrates how customers can significantly enhance the security posture of their VMware Horizon on VxRail VDI environment by using the STIGs (and associated security approaches such as the creation and enforcement of Microsoft security baselines) available for a VMware Horizon on VxRail technology stack running a Microsoft Windows-based user operating system. This is just one approach and is not meant to cover all possible security requirements.

Components targeted by STIGs and Microsoft security baselines

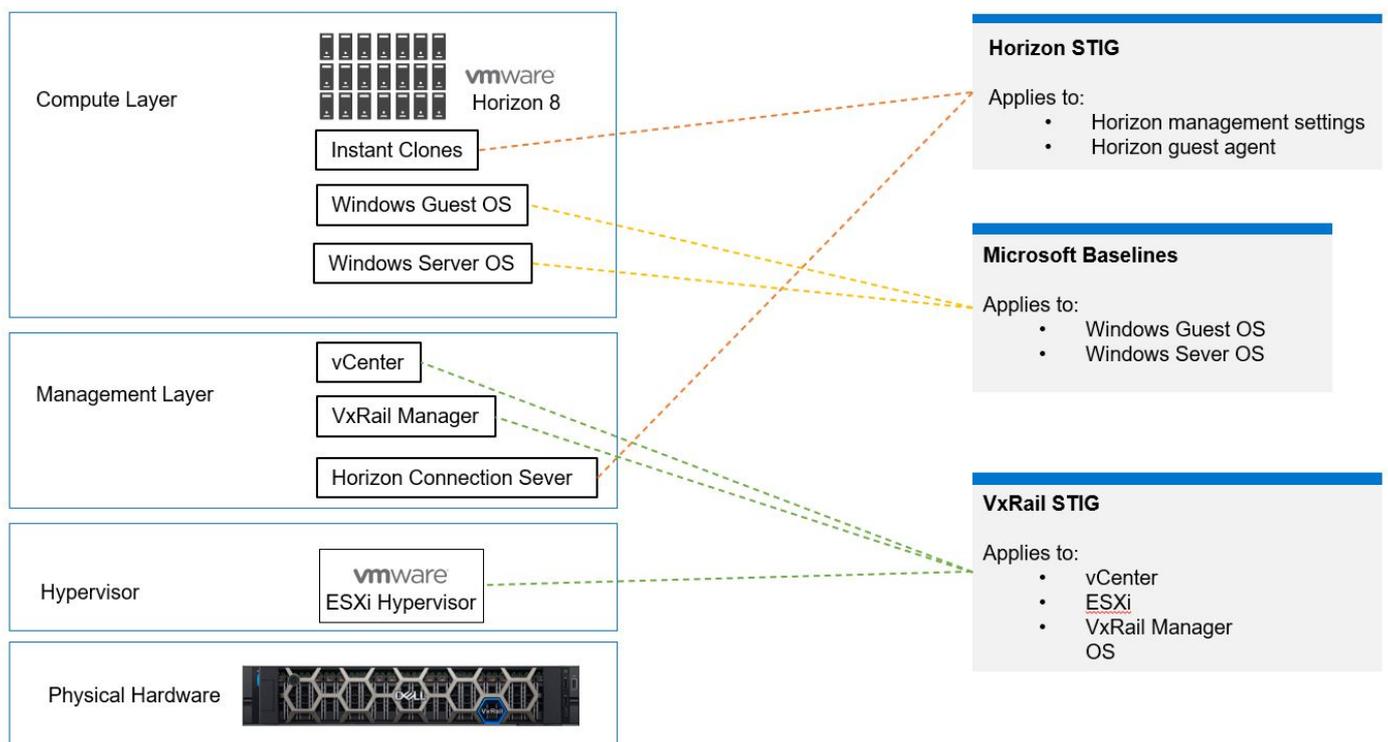The following figure shows the components targeted by the STIGs:

Figure 25. Components targeted by STIGs

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.

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 Dell 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.

Summary

Topics:

- Overview
- Next steps
- We value your feedback

Overview

This design guide describes the integration of vSAN-based appliances 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.

Next steps

Dell Technologies has configurations to fit the needs of any size organization:

- **VxRail E660 or E660F (E Series)**—For small deployments where energy concerns exist or space is limited. Up to two NVIDIA T4 GPUs are supported per node.
- **VxRail V670F (V Series)**—VDI-optimized configuration that offers the highest processor speeds and graphics capability. Up to two NVIDIA A16 or A40 GPUs are supported per node.
- **vSAN Ready Node R650**—This device is a prevalidated configuration in a dense rack platform. Occupying only 1U in the rack, this powerful server supports Density-Optimized configurations for VDI. It supports up to three NVIDIA T4 GPUs per node.

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.

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

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 VxRail Hyperconverged Infrastructure
- Dell vSAN Ready Nodes
- STIG and Microsoft Security Baseline-Based Hardening of a VMware Horizon on a VxRail-Based VDI Environment Implementation Guide

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

Appendix A: Definition of performance metrics

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

Table 9. 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 object, 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 object, and so on.