

Dell EMC PowerStore 500: 1,500 VMware Horizon VDI Users

February 2022

H18729.2

Reference Architecture

Abstract

This document describes an architecture for deploying VMware Horizon virtual desktops with a Dell EMC PowerStore 500 model appliance.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA February 2022 H18729.2.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

- Executive summary.....4**
- Introduction5**
- Connectivity.....6**
- Storage design10**
- Data encryption16**
- Test results16**
- Summary.....22**
- Appendix: Technical support and resources23**

Executive summary

Overview

This document details the architecture design, configuration, and implementation considerations when deploying a VMware Horizon virtual desktop infrastructure (VDI) with a Dell EMC PowerStore 500 model storage appliance. This document is storage-focused and details the specific configurations of this system. The network and fabrics that are used for connectivity are explored in detail.

PowerStore layers protocols and services on a flexible, container-based architecture, ensuring that features and services can be added to existing installations with minimal impact. This paper discusses the benefits of a unified storage product that combines block and file services into a VDI platform. PowerStore allows expansion of both functionality and capacity to protect storage investments.

The information in this document is based on the guidelines in the paper [Dell EMC PowerStore: VMware vSphere Best Practices](#). This document includes supplemental information regarding VDI-specific details when running VMware Horizon.

As part of this reference architecture, the solution is tested using the industry-standard Login VSI VDI benchmarking tool. The performance results demonstrate that the PowerStore 500 model appliance delivers submillisecond latency to 1,500 virtual desktops running a Login VSI knowledgeworker profile.

Audience

This document is intended for IT administrators, storage architects, partners, and Dell Technologies employees. This audience also includes any individuals who may evaluate, acquire, manage, operate, or design a Dell EMC networked storage environment using PowerStore systems.

Revisions

Date	Description
May 2021	Initial release for PowerStoreOS 2.0
November 2021	Minor updates
February 2022	Template update

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Marty Glaser

Note: For links to other documentation for this topic, see the [PowerStore Info Hub](#).

Introduction

PowerStore overview

PowerStore storage arrays are designed to enable growth, flexibility, and resiliency. PowerStore services are hardware and protocol independent. This architecture allows file services to be layered without restricting block-storage flexibility.

Dell EMC storage products are integrated into a mature monitoring and reporting suite that is designed to minimize management effort. Robust reporting and notification services reduce storage management time. Integrated dashboards enable at-a-glance insight into the health and performance of the platform.

PowerStore achieves new levels of operational simplicity and agility. It uses a container-based microservices architecture, advanced storage technologies, and integrated machine learning to unlock the power of your data. PowerStore is a versatile platform with a performance-centric design that delivers multidimensional scale, always-on data reduction, and support for next-generation media.

PowerStore brings the simplicity of public cloud to on-premises infrastructure, streamlining operations with an integrated machine-learning engine and seamless automation. It also offers predictive analytics to easily monitor, analyze, and troubleshoot the environment. PowerStore is highly adaptable, providing the flexibility to host specialized workloads directly on the appliance and modernize infrastructure without disruption. It also offers investment protection through flexible payment solutions and data-in-place upgrades.

With the addition of integrated file services, there is no requirement to purchase an external storage device for file shares. User access through NFS or SMB is integrated, with the ability to share files between Microsoft Windows and UNIX or Linux users. This ability simplifies administration as a single interface can be used to manage block and file resources.

PowerStore can serve data using iSCSI, Fibre Channel, or both protocols simultaneously to provide seamless integration into most environments. PowerStoreOS 2.0 also includes support for NVMe over Fibre Channel (NVMe/FC).

VMware Horizon overview

VMware Horizon is an industry-standard solution for virtualizing desktops. It provides advanced control and great flexibility for large VDI environments.

For more information about VMware Horizon, see [VMware Horizon Product Overview](#).

Prerequisite reading

The best practices in this document require knowledge from the following resources at [Product Documentation & Videos](#):

- *Dell EMC PowerStore: VMware Best Practices Guide*
- *Dell EMC PowerStore: Host Configuration Guide*
- *Dell EMC PowerStore: Network Configuration Guide*

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 1. Terminology

Term	Definition
PowerStore Manager	The HTML5 web-based user interface (UI) for storage management.
Appliance	A solution containing a base enclosure and attached expansion enclosures. The size of an appliance could include only the base enclosure or the base enclosure plus expansion enclosures.
Node	The component within the base enclosure that contains processors and memory. Each appliance consists of two nodes.
Cluster	One or more PowerStore appliances in a single grouping. Clusters are expandable by adding more appliances (up to four total).
Base enclosure	The enclosure containing both nodes (node A and node B) and the NVMe drive slots.
Expansion enclosure	An enclosure that can be attached to a base enclosure to provide 25 additional drive slots.
Storage volumes	PowerStore volumes that use block storage.
NAS server	Network-attached storage (NAS) servers allow block storage to be used for file storage. The NAS service translates file requests to block storage. NAS supports Microsoft Windows SMB 3.1.1 and NFS v3 or v4.
File system	Allows PowerStore to provide file-level access to end users and applications. It also controls user permissions to files, and folders or directories.

Connectivity

Introduction

PowerStore supports iSCSI and Fibre Channel (FC) hosts. NVMe over Fibre Channel (NVMe/FC) is supported with PowerStoreOS 2.0. These protocols can be run simultaneously, but this configuration can add complexity to routing and failover scenarios. It is recommended to configure a host to use only one protocol. Hosts in a clustered environment should also be configured to use only one protocol.

For this configuration, a converged Ethernet environment is used to provide block storage (iSCSI) and file storage (SMB). FC or NVMe/FC and can also be used to provide block storage. One advantage to using a converged Ethernet configuration is reduced complexity and cost. Small office, branch office, and edge use cases benefit when minimizing complexity and hardware footprints with converged networks.

In this environment, the following connectivity modules are configured in each PowerStore node (node A and node B (see [Figure 1](#)):

- Two 25 Gb/second Ethernet modules for NAS and file (SMB), using the embedded module
- Two 25 Gb/second Ethernet modules for block storage (iSCSI), using I/O Module 1

This configuration provides more than enough bandwidth for the block and file user data load. To isolate and allow traffic prioritization, VLANs are configured for each type of Ethernet traffic.

Block storage

PowerStore presents storage to external hosts through either block or file interfaces. Block storage is the most common datastore path for virtual machines. The various speeds and protocols offered make it ideal for performance. NFS is also used but is less common. PowerStore supports 25 GbE and 32 Gb FC connections for broad compatibility and performance requirements. For instructions and best practices for configuring hosts, see the *Dell EMC PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#).

Configuring multipath I/O (MPIO) is critical to achieving the best performance from VMware hosts and should be configured properly. For MPIO configuration guidance, see the document *Dell EMC PowerStore: VMware vSphere Best Practices*.

The large number of small I/Os generated from VDI environments benefit greatly from MPIO round robin and rapid-path switching. The use of rapid-path switching can help improve VM performance.

The MPIO configuration used for this configuration is as follows:

Path selection policy: Round Robin (VMware)

Storage array type policy: VMW_SATP_ALUA

This configuration provides optimal pathing with failover protection and path discovery.

See the *Dell EMC PowerStore VMware Best Practices Guide* and *Dell EMC PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#) for guidance on setting MPIO and other VMware best practices.

iSCSI

Using iSCSI for block storage allows convergence of storage and networking infrastructure. It requires careful planning to ensure adequate bandwidth and fault tolerance. Redundant networks are preferred for data availability.

With iSCSI, increasing the maximum transfer unit (MTU) size from a 1,500-byte standard frame to a 9,000-byte jumbo frame is also highly recommended. The use of jumbo frames allows greater packet efficiency for higher bandwidth. Using jumbo frames is not a requirement but should be considered. Combined with the 25 GbE connectivity available in PowerStore, iSCSI can handle demanding workloads and high throughput requirements. All devices in the iSCSI data path in this configuration (hosts, switches, PowerStore) are configured for jumbo frames with an MTU of 9,000. This change reduces network switch loads as a smaller number of large packets are processed.

Another benefit of using Ethernet infrastructure for iSCSI is the ability to share the network with file services. This use requires planning to ensure that there is enough bandwidth for all services, but it also reduces the hardware requirements. The use of different VLANs allows traffic prioritization to ensure that block traffic has the highest priority.

In this environment, iSCSI traffic is configured to use two dedicated 25 GbE ports in I/O Module 1 on each PowerStore node, for four ports total. Jumbo frames are enabled on all

network devices in the data path with MTU = 9000 bytes. PowerStore automatically configures ports in redundant pairs across I/O modules for redundancy and to optimize throughput.

Networking

environment for this configuration is an Ethernet-only solution. The connections are 25 GbE, with trunking for management redundancy and performance. The NAS servers use dual 25 GbE ports on each node for an aggregate bandwidth of 100 Gb/sec.

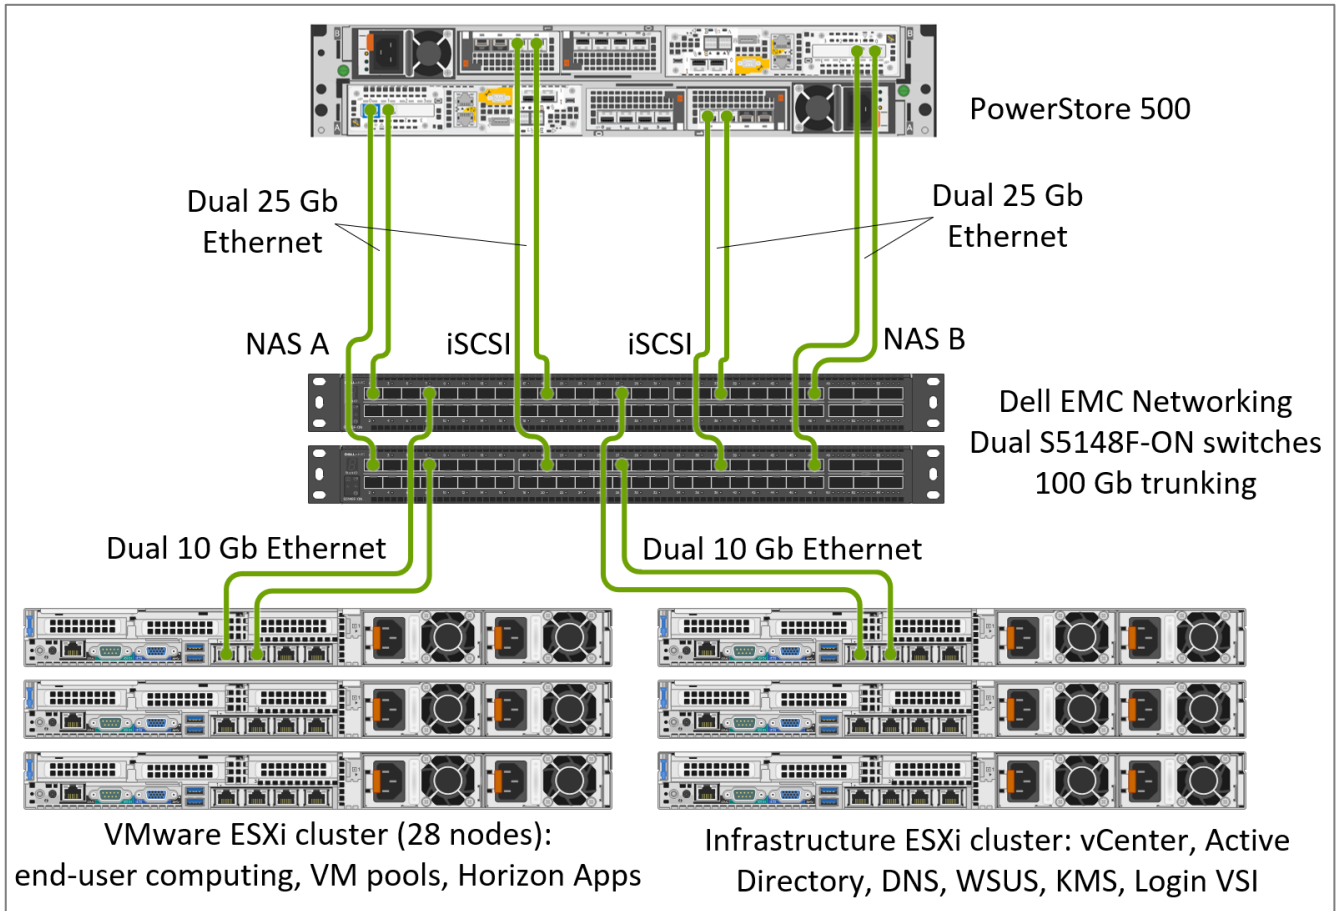


Figure 1. Logical network

NAS servers

A key feature of PowerStore is the ability to support multiple NAS servers on the same array. This support has multiple benefits, including the following:

- NAS servers support multitenancy.
- NAS servers are logically separated from each other.
- Clients of one NAS server do not have access to data on other NAS servers.
- Each NAS server has its own independent configuration (for example, DNS, LDAP, NIS, interfaces, protocols, and so forth).

Each NAS instance is isolated so multiple environments can share a single system. Enforced isolation helps increase security and stability. Because a NAS instance runs on a dedicated PowerStore node for maximum performance, a minimum of two NAS servers

is recommended in a VDI environment. The load can then be split between two NAS instances with a NAS server on each node. The type of data being stored determines the best method of balancing the load.

For VDI user profiles, a tool like [VMware Dynamic Environment Manager \(DEM\)](#) can be used to balance the data across multiple file systems. This tool allows setting dynamic rules to place data based on user-defined criteria. This use automates data redirection for greater control and centralized management.

In this environment, two NAS servers are configured. PowerStore Manager automatically assigns the first NAS server created to node A. PowerStore Manager automatically assigns the second NAS server created to node B. This configuration helps keep the appliance balanced to optimize performance. As each NAS server is an isolated process, each NAS server has its own file system (FS) provisioned from the appliance storage pool. Multiple file systems can be created for each NAS instance if necessary.

See the [Dell EMC PowerStore: VDI Best Practices Guide](#) for more information about the advantages of using file redirection with VDI.

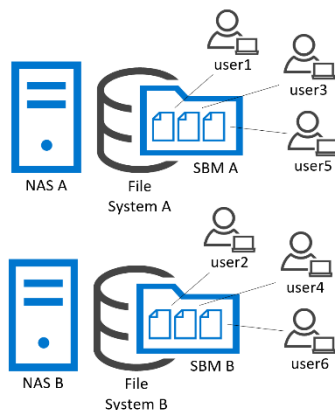


Figure 2. NAS, file, and SMB configuration used for VDI profile redirection

File systems

PowerStore arrays support multiple file systems for flexibility, security, and availability. This support allows for creating file shares that are based on user type or operating environment. A single NAS and file-system instance can support both SMB and NFS simultaneously.

For this testing, DEM is used for round-robin placement of user profiles as shown in [Figure 2](#). This configuration load-balances the file activity and reduces the number of files on each share for backup purposes. The scenario is designed to simulate Windows desktop users accessing files from a Windows server.


Storage design

Introduction The size and type of traffic that is generated from a VDI environment are noticeably different than most business applications. The difference in the size and pattern of VDI traffic is consistent but requires a few changes to the storage design for proper management.

Storage array The PowerStore 500 model was added to the PowerStore portfolio with the PowerStoreOS 2.0 release. The PowerStore 500 provides a lower-cost entry point and provides customers with additional options when determining the optimal storage design in a smaller environment. Branch office or similar edge computing scenarios are ideal use cases for the PowerStore 500. The PowerStore 500 provides excellent performance and functionality with advanced capabilities with a breadth of features and monitoring to serve the needs of a broad spectrum of customers.

PowerStore supports many front-end connectivity options to support a wide variety of environments and workloads. The PowerStore 500 model used is in this testing is configured as shown here.

- Dual controller nodes (node A and node B) for load balancing, fault tolerance, and high availability
- 25 flash based NVMe drives
- Dynamic Resiliency Engine (DRE) for optimized disk performance and fault tolerance
- Two 25Gb Ethernet ports on each node are configured for iSCSI



Management and NAS Ports

iSCSI ports

Figure 3. PowerStore 500 model hardware front and rear views

VMware best practices

Most VMware best practices apply to VDI environments, but there are a few changes required due to the nature of high-density VM configurations. For VMware best practices and how to apply them, see these documents:

- [Dell EMC PowerStore: VMware vSphere Best Practices](#)
- [Dell EMC PowerStore Host Configuration Guide](#)

With most applications and workloads, storage traffic is typically 70% reads. There are applications that generate all reads or all writes, but the average is a 70/30 read/write ratio. VDI is typically 70-80% writes. The array-sizing changes are based on the higher write ratio.

For best practices for Horizon on PowerStore, see the document *Dell EMC PowerStore: VMware Horizon Best Practices Guide*.

Volume count

There are several factors that determine the optimal volume count with Horizon on PowerStore. The volume count will vary based on the scenario. Performance and management overhead are typically the two most important factors when determining the number of volumes.

For performance reasons, a minimum of 16 volumes is recommended for a VDI environment. This configuration ensures good queue balancing and reduces I/O bottlenecks.

Balance this volume recommendation with the need to keep the volume count low to reduce management overhead. The VMware guidance for volume count should also be considered. Internal testing shows that 16 volumes can support 3,600 users from a performance perspective, but it may have an impact on VM recovery from a backup. Using array-based snapshots reduces recovery time and minimizes the impact of higher user-per-volume counts.

This configuration consists of 16 volumes for VM storage.

I/O size

The average I/O size of storage traffic in a VDI environment typically ranges from 18k to 32k. This size can vary if host-based caching is enabled or allowed. Some VM configurations do not allow host-based caching.

The average I/O size from Horizon VMs during the workload test was 20k at the array level.

Heavy writes

VDI environments are typically heavily write-based. After a VM is booted and the user logs in and begins work tasks, most of the I/O traffic consists of changes from the VM. The traffic consists of file changes, swap file writes, memory paging, and updates to user preferences.

During the steady state phase, the write ratio is typically 70% or higher. VDI has a higher write ratio than most other workloads. This write-heavy traffic pattern means VDI is a more demanding application for a storage array to support. PowerStore is well suited to provide optimal performance for the most demanding heavy-write VDI workloads.

Instant clone pools

Instant clones generate a large volume of traffic during the provisioning process due to the method used to create the VMs. Since the in-memory VM Fork technology creates the machines quickly, the volume of traffic is significant. The I/O required to complete the creation is brief but large. The VDI administrator must plan for bursts of high traffic on PowerStore when creating or refreshing pools of instant clone VMs. It is desirable to dedicate a PowerStore array to VDI if it supports a large VDI implementation so other workloads are not impacted.

Since instant clone machines reset on logout, as users log off, their virtual desktops refresh. Logoffs occur throughout the workday and the I/O and CPU demand to refresh desktops should be considered when sizing the environment. If many (or all) VDI users log off simultaneously, there is also a brief but significant burst in I/O and CPU demand as desktops are refreshed.

For this design, two instant clone pools are configured, with a dedicated Horizon Connection server for each pool. This configuration ensures load balancing across multiple connection servers. A randomization algorithm is applied to the login order to spread the login load evenly across the connection servers to ensure best performance and to better simulate a production environment.

Capacity

The density of data on VDI tends to be high because Horizon VDI environments typically use data-reduced clones. With profile-redirection technology like VMware DEM, clone pools appear as persistent to users. DEM enables users to personalize their desktop, but pools can still employ the benefits of nonpersistent VMs. Keeping the VMs small and clean improves performance and recovery time.

The system is configured with 16 1TB datastores for virtual desktops and two 1TB volumes for file storage. Since any volume can be expanded at any time, these size recommendations are starting points. All volumes are also thin provisioned and have space reclaim enabled. Thin provisioning and on-demand volume resizing maximize storage efficiency.

Virtual machine types

In a virtual desktop environment for end users, Horizon supports several different types of virtual workspaces that are referred to collectively as virtual desktop infrastructure (VDI). However, the various types of virtual desktops have different characteristics.

Full clones

Full clones were the first type of clone used. The use of full clones requires no special tools. It is as simple as cloning a VM in vCenter. If the clone is a Windows VM, it is typically provisioned from a system-prepared (Sysprep) source image. Changing the name and IP address of the clone before joining it to a domain also works. If Horizon is used, it brokers the creation of unique full clones.

Although PowerStore data reduction can help reduce the disk footprint, full clones require the most disk storage because they do not share a common base image. An important consideration with full clones is management overhead. Each full clone must be managed as an individual desktop and each full clone requires individual patch management and updating.

The I/O profile of full clones is that of full machines, with reads and writes spread across the entire dataset. Full clones typically save user data to the local profile resulting in a persistent virtual experience. This VM type typically has the lowest overall storage I/O and CPU impact since data is saved between sessions. The VM is not refreshed at logoff. The downside is management overhead as each full clone must be managed and updated individually like a physical desktop.

The impact on network bandwidth and storage can be significant, however, if thousands of full clone machines attempt to do Windows Updates all at once. Activities such as patch management should be planned carefully. Leverage a local Windows software update server (WSUS) to avoid consuming Internet bandwidth needlessly. Maintenance windows should be defined and scheduled so WSUS updates a subset of machines that finish updating before the next group updates. This strategy helps to spread out the resource demands. The impact of each round of monthly updates can vary based on the number and type of updates released for that cycle, making it difficult to predict the resource impact.

Linked clones

Linked clones for a time were the most common type of clone used in VMware Horizon environments.

However, with the rapid adoption and improvements offered with instant clone technology, support for linked clones was deprecated with the release of VMware Horizon 8. In addition, linked clones are not supported in cloud or hybrid cloud environments.

Linked clones use a common base image with a disk digest architecture to redirect changes. All changes to the VM are written to the digest. If persistence is required, all changes are saved. Otherwise, all changes are flushed on logout. The disk load is higher in a nonpersistent configuration as a new user profile gets created on each login.

The disk footprint of linked clones is small. The base image is copied to each datastore in the pool and all machines share the base images. This behavior greatly reduces duplication of the base image.

Instant clones

Instant clones have replaced linked clones as the preferred VDI method with Horizon. The improvements in instant clone technology have enabled this clone type to replace other clone types in most configurations. However, instant clone implementations are I/O intensive during pool creation, and at user logoff as each VM gets refreshed based on the current base image.

A major advantage of instant clones is each machine is refreshed on every user logoff so it matches the current base image. This behavior greatly simplifies the update process: update the base image, and each VM is refreshed to use the new base image at the next user logoff. The base image can be updated frequently without user interruption. The instant-clone deployment and update process reduces management overhead and minimizes end-user impact.

It is necessary to anticipate the significant burst of I/O and CPU demand on the storage array when using instant clones:

- When a new instant clone VM pool is created.

Tip: Create pools outside of peak hours to avoid impacting active VDI users.

- When a user logs off and their VM is refreshed to match the current base image.

Most user logoffs occur toward the end of the workday or work shift. As the number of active VDI users shrinks, logoffs are less likely to impact active VDI users.

When user logoffs are sporadic and spread out, the I/O and CPU bursts are spread out and less impactful. However, if a high percentage off simultaneously, the burst in I/O and CPU demand will stress the storage array. A correctly sized environment can accommodate bursts in I/O and CPU due to pool creation and user logoffs. When using instant clones, the extra demand on the storage array from a CPU and disk I/O perspective must be anticipated and factored in when sizing the environment. The architecture in this document accounts for the bursts in I/O and CPU demand for these activities.

Dedicate a PowerStore array to VDI if the user count is high enough that bursts caused by pool creation and image refreshes at user logoff will impact other workloads.

Windows I/O patterns

The evolution of Windows has brought new features to users, but these features require additional host resources. For example, a Windows 10 desktop requires significantly more resources than Windows 7, increasing the load dramatically. A large portion of the increase is due to added services. Encryption and memory-protection schemes that are built into Windows serve to reduce the attack surface but also generate additional load.

As Windows 10 has evolved, more security and features have been added that have increased the load even further. This constant increase in I/O load with newer Windows 10 version requires matching increases storage performance.

Consider standardizing on an older long-term servicing channel (LTSC) version Windows 10 that is still fully supported and fully updated. This strategy can help avoid the undesirable impact of increased I/O, CPU, and memory requirements with newer Windows 10 releases.

As stated previously, the typical I/O pattern for a VDI workload is 70-80% writes during steady state when users are logged in.

Guests

Operating system

Windows 10 build 1809 was used for this test. Two processor cores and 3 GB of RAM are assigned to each VM.

The processor requirements for Windows 10 can change based on the build number and patch level. Processing and memory requirements for Windows 10 have increased over time with newer builds. One of the best ways to reduce the CPU and disk load for a Windows virtual machine is to use the [VMware OS Optimization Tool](#). This tool is a consolidated interface that can be used to change the behavior of virtual machines to reduce their hardware requirements. Be careful if choosing the most-restrictive

configurations, which can cause applications or even Windows features to stop working. This tool is designed to reduce the effort it takes to optimize Windows. It requires testing of the configuration to ensure sure all user applications still work correctly with the optimizations applied.

User profiles

Another tool that was mentioned previously in this paper is [VMware DEM](#). This tool redirects user data to one or more file shares based on certain conditions. This capability allows directing different types of data to different locations. User-profile data and documents can be pointed to home directories that are replicated. Temporary data such as application temp files that do not need to be saved are placed on volumes with no data protection. This approach reduces replication traffic and the overall storage load.

All user profiles are redirected to file shares to capture the user load for this test. The ability of the PowerStore array to host file shares consolidates storage management and reduces backup complexity. A single PowerStore appliance can host all end-user data.

Applications

The applications can be layered using VMware AppVolumes. It greatly improves application management and updating by removing the applications from the base image and managing them independently. When an update is required, the layered application is updated externally and the updated version is applied to the environment.

Layering applications enables rapid application provision and updating. The individual applications are managed independently with the ability to present groups of applications to user groups. The advantage is quick deployment of layered applications to new users or groups of users.

PowerStore DRE The PowerStore dynamic resiliency engine (DRE) automatically manages the underlying storage for maximum performance and capacity, eliminating the need for administrators to configure the storage pool. Manually setting or configuring these options is unnecessary in PowerStore. The underlying storage configuration and drive management are automatic and require no management.

Horizon configuration Horizon does not require any special configuration to work with PowerStore. The storage is presented as block storage volumes discoverable through vCenter. When the pool is created, the volumes are visible as usable disk space.

For this test, Horizon is configured to use all 16 datastores for VM creation across both pools. This configuration spreads the virtual desktops evenly from both pools across all datastores. The goal is to load balance all storage for best performance and consistency.

vCenter provisioning settings

Based on performance testing under load, the Horizon **max concurrent instant clone operations** setting is set to 10. This value is derived from observing storage latency during instant clone provisioning.

Secondary pools One simple way to help users quickly return to work is to keep additional VMs ready in a recovery pool. This pool can be assigned to end users if the primary systems are down.

This pool may not have the same performance characteristics but can aid in business continuity. Because of the rapid deployment nature of virtual clone VMs, it is also possible to have the VMs created on demand in a failover scenario. This strategy reduces the storage requirement. Since only one environment would be generating load, the IOPS requirement would be the same once all the users are logged in. The load would be higher during virtual machine creation.

Data encryption

Many applications have data-encryption requirements for data at rest. Data at Rest Encryption (D@RE) is enabled by default with PowerStore. No user configuration steps are necessary to protect the drives.

Test results

Login VSI is used to run workload tests that closely simulate a real VDI environment. It manages the process of automatically logging users on, running daily worker tasks, and logging users off. Metrics are captured during the test to evaluate the results.

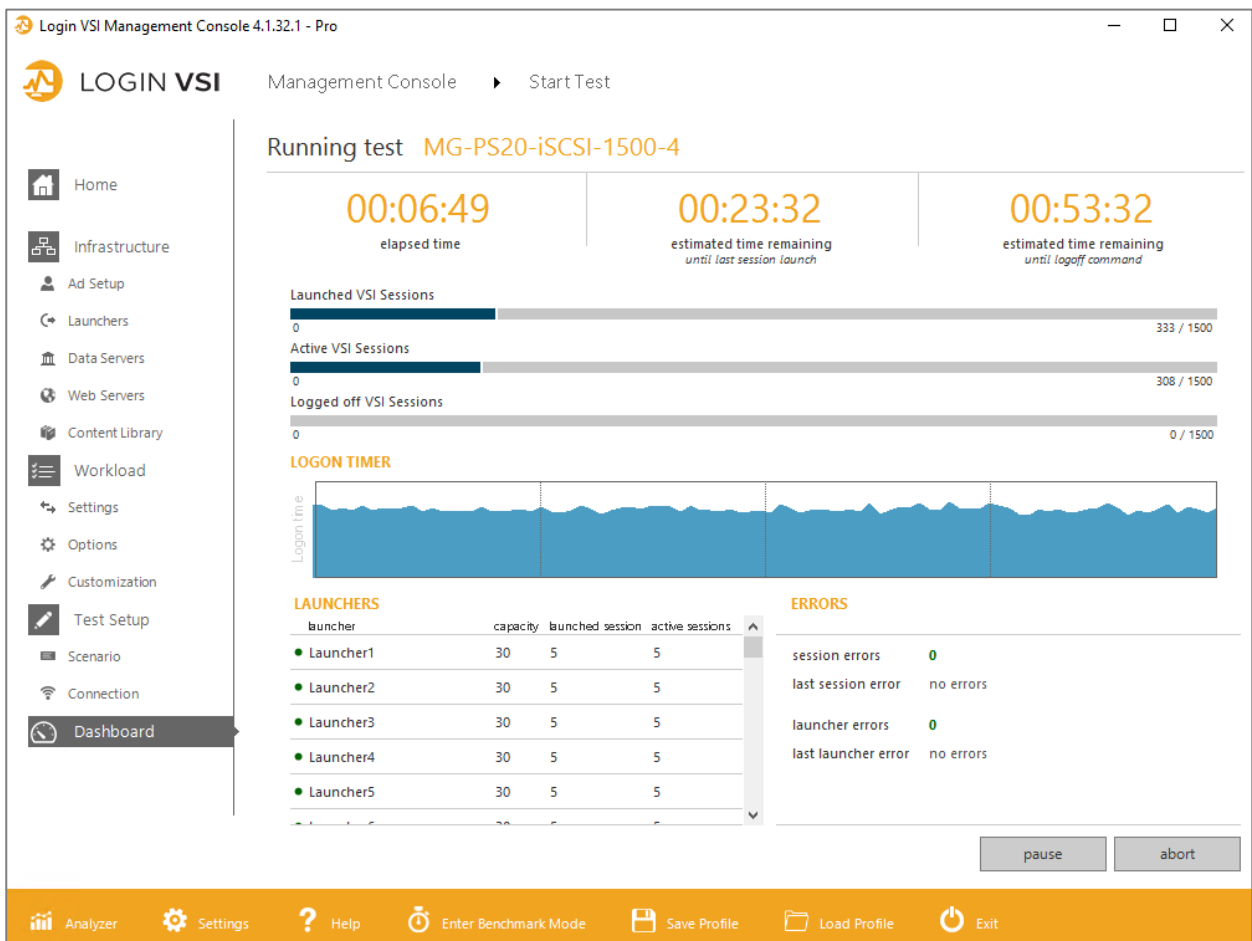


Figure 4. Login VSI management console and dashboard


This test is configured to log in 1,500 Login VSI knowledgeworker users in 30 minutes, at a steady rate of 1.2 users each second. This login rate is more aggressive than might be typical in a user environment. The test is configured this way to show how well the PowerStore 500 model performs when a short time window is specified to log in a large user pool.

VDI user login rates are typically gradual and spread out at the beginning of a workday or work shift. User logoffs are also typically gradual and spread out. However, rapid login and logoff events are possible and should be planned for. For example, a power outage will force all desktops to reset simultaneously. Once power resumes, the array will experience a burst in I/O and CPU demand as instant clone desktops refresh. In addition, all users may attempt to log back into their desktops simultaneously causing a login storm. For this reason, for large VDI deployments, it is desirable to dedicate a PowerStore array to VDI to avoid impacting other workloads when these events occur.

User environments where mass login (login or boot storms) and mass logoff events are typical may support fewer VDI users. Running the most recent version of Windows 10 may also reduce the user count due to higher memory, CPU, and I/O demands. A lower user count will help ensure adequate array performance for any expected scenario where boot or login storms and mass logoffs and refreshes are common. Other factors unique to each customer may result in a lower or higher user count for a PowerStore array than what is shown in this test. Understanding key design elements, user needs, and workload behavior is critical to right-sizing a VDI environment.

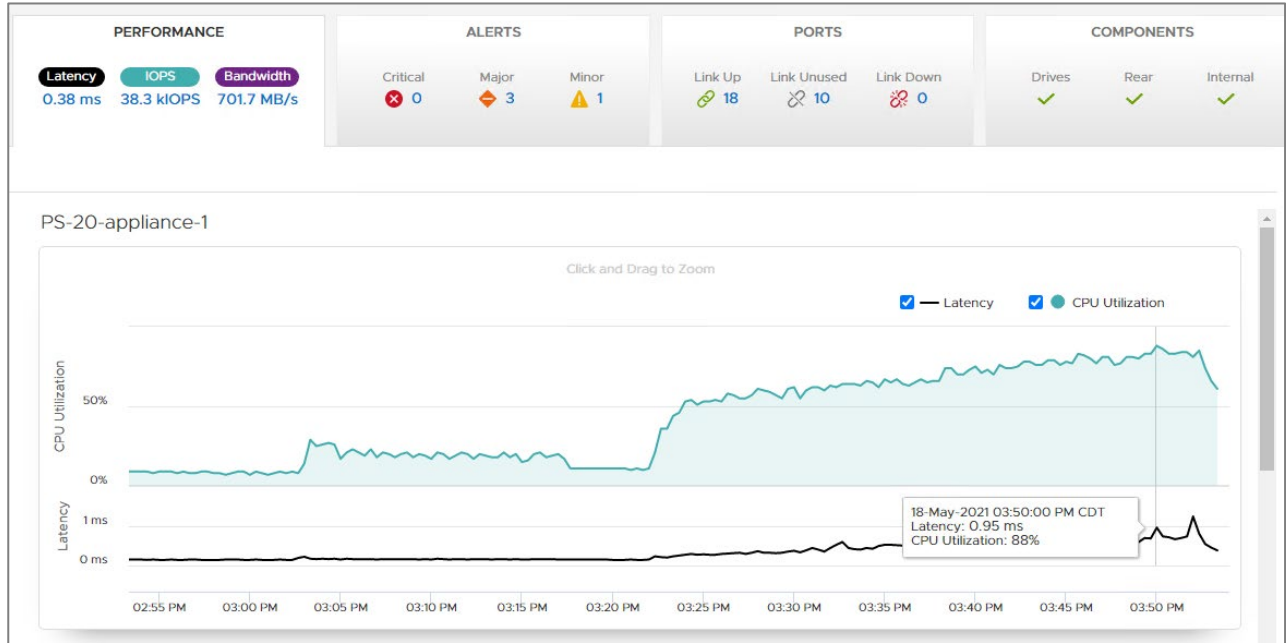
PowerStore test results Login phase

The login phase is where all users are logged into virtual desktops and they start running knowledgeworker tasks. The login phase for this test is set to 30 minutes. The figures below show the key PowerStore performance metrics during the login phase.



Test results

PowerStore shows a steady and gradual increase in IOPS and bandwidth while maintaining submillisecond latency throughout the period.



PowerStore shows a steady and gradual increase in storage CPU utilization that peaked at 88% near the end of the login phase.

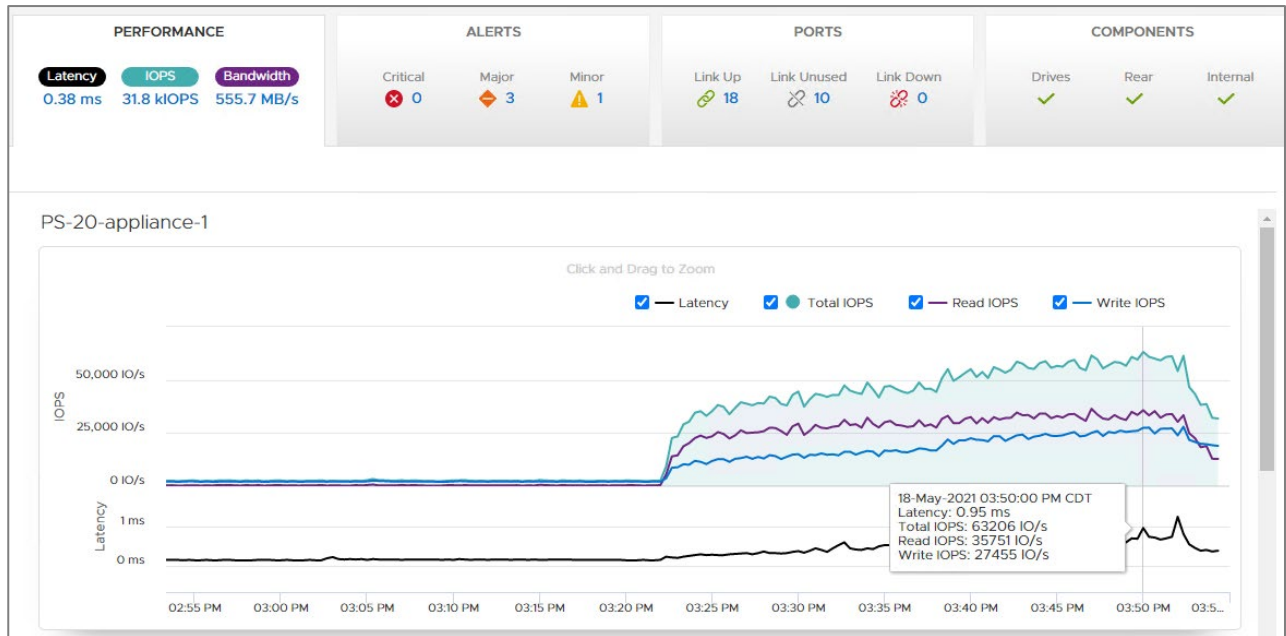


Figure 7. Login phase IOPS detail

IOPS peaked at over 63,000 IO/s during the login phase. As expected, the percentage of read IOPS is higher than write IOPS during the login phase. During steady state, this ratio will reverse as shown below.

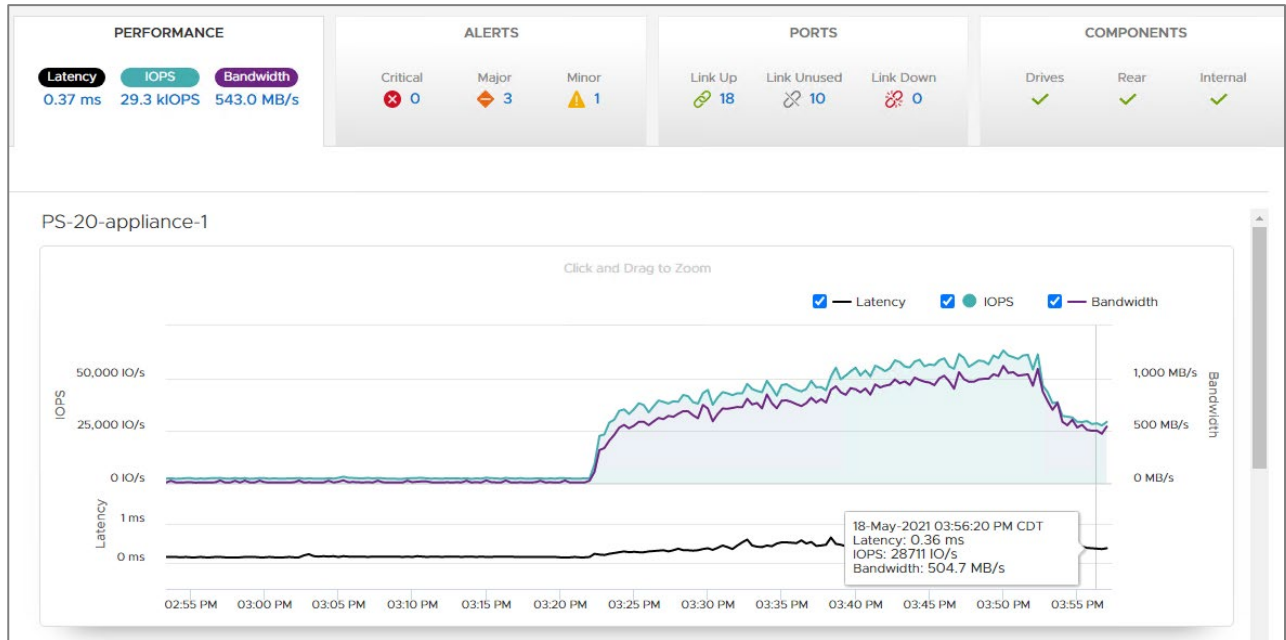


Figure 8. Steady state IOPS, bandwidth, and latency

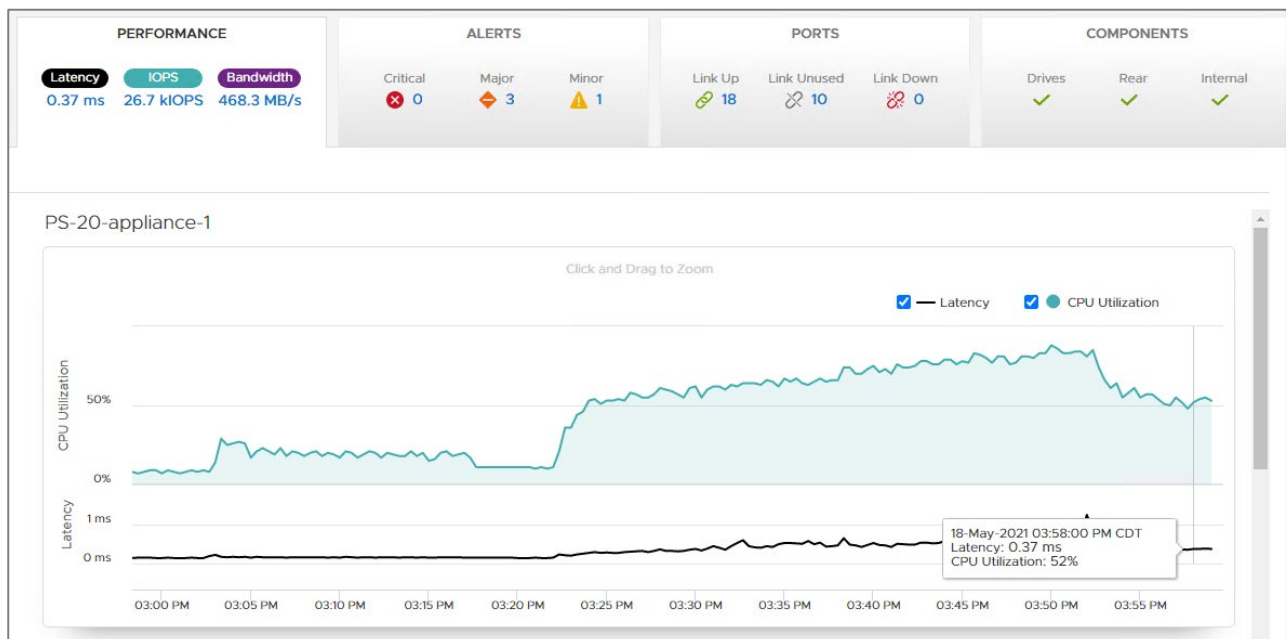


Figure 9. Steady state CPU

CPU demand is also less during steady state, averaging about 50%, with latency performance averaging less than 0.4 milliseconds.

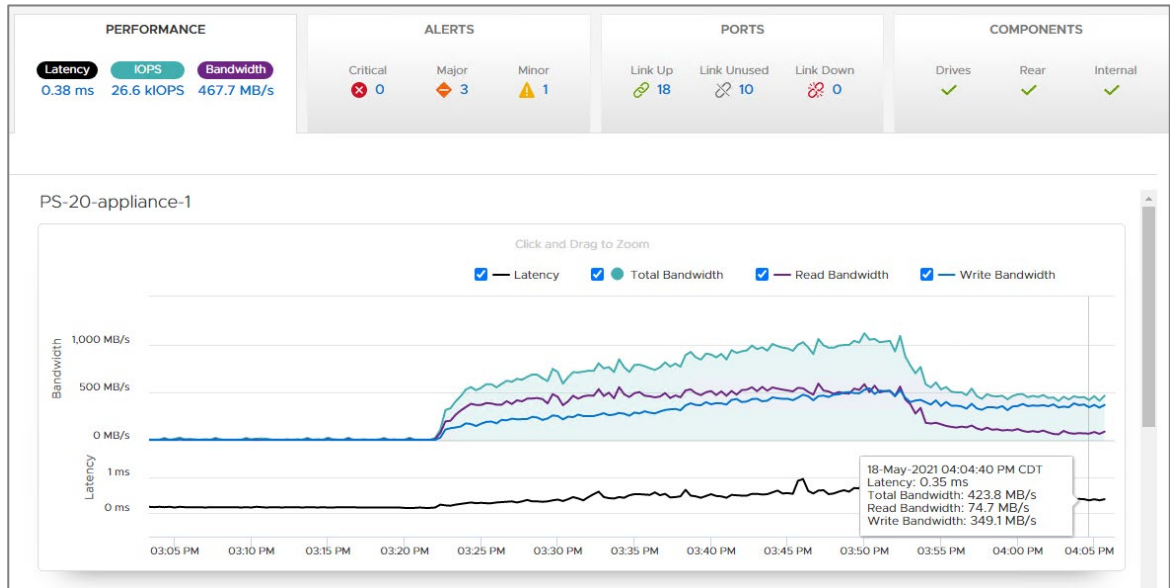


Figure 10. Steady state bandwidth

The read/write ratio shows a higher percentage of writes during steady state. In this test, the data points in the call-out in Figure 10 show a sustained bandwidth of 74 MB/s reads and 423 MB/s writes. This result translates to an 18% read / 82% write ratio during steady state.




Figure 11. Refresh phase

When a user logs off an instant clone VM, it is refreshed to match the current base image. In this test, Log in VSI logs off all 1,500 users at the end of the steady state in about 10 minutes. Figure 11 shows a period of increased I/O, bandwidth, and CPU demand on the PowerStore appliance as the refresh happens. Provisioning and refresh rates can be

adjusted to lower the I/O and CPU demand. The trade-off is that it may take longer when provisioning new pools and for refreshes at logoff to complete.

Although not shown here, defining a new Horizon instant clone VM pool also causes a similar period of high I/O, bandwidth, and CPU activity. When administrators plan for a successful VDI deployment using instant clones, a great user experience during the login phase and steady state are paramount considerations. However, the resource demands for VM pool creation and VM refresh at logoff can also be significant and must also be understood and factored in when sizing the environment.

Login VSI test results




Figure 12. Login VSI test results

The **VSImax threshold** metric predicts the number of estimated users supported based on storage performance as the workload increases, as factored against response time from applications. The graph shows that resource demands ramp predictably and linearly under load, and that the PowerStore 500 model performed well.

Summary

The test results show how well the PowerStore 500 model storage appliance performed during the various phases of this test. The low latency, high bandwidth, and predictable throughput demonstrate that the PowerStore 500 is an excellent storage platform for hosting VMware Horizon VDI workloads.

Appendix: Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

The [Dell Technologies Info Hub > Storage](#) site provides expertise that helps to ensure customer success with Dell EMC storage platforms.

[Dell.com/powerstoredocs](https://dell.com/powerstoredocs) provides detailed documentation about how to install, configure, and manage Dell EMC PowerStore systems.

For more information, see the following related resources:

- [VMware Documentation](#)
- [Dell Technologies VDI info hub](#)