

Dell EMC PowerStore: VMware Horizon VDI Best Practices

Abstract

This document provides best practices for deploying VMware® Horizon® virtual desktops with Dell EMC™ PowerStore™. It also includes recommendations for performance, availability, scalability, and integration.

April 2021

Revisions

Date	Description
April 2020	Initial release: PowerStoreOS 1.0
April 2021	Update for PowerStoreOS 2.0

Acknowledgments

Authors: Damon Zaylskie, Marty Glaser

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

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 © 2020–2021 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. Other trademarks may be trademarks of their respective owners. [4/15/2021] [Best Practices] [H18203.1]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	5
Audience	5
1 Introduction.....	6
1.1 PowerStore overview.....	6
1.2 VMware Horizon overview.....	6
1.3 Prerequisite reading	6
1.4 Terminology	7
2 Connectivity.....	8
2.1 Block storage	8
2.1.1 Boot from SAN.....	8
2.1.2 iSCSI.....	8
2.1.3 Fibre Channel	8
2.1.4 NVMe over Fiber Channel	9
2.2 NAS servers.....	9
2.2.1 Availability.....	9
2.2.2 File systems.....	9
3 Storage design	11
3.1 VMware best practices	11
3.1.1 Volume count.....	11
3.2 I/O size.....	11
3.3 Heavy writes	11
3.3.1 Instant clone pools.....	12
3.4 Capacity.....	12
3.5 Guests	12
3.5.1 Windows	12
3.5.2 Linux	13
3.6 Networking.....	13
3.6.1 Management.....	13
3.7 DRE and storage containers	14
3.8 Changing capacity needs	14
3.9 Horizon configuration.....	15

3.10	Replication	15
3.11	Secondary pools	16
4	Data encryption	17
5	Management and monitoring.....	18
5.1	Data reduction	18
5.2	Scripting and automation.....	18
5.2.1	REST API	18
5.2.2	PowerStore CLI	19
5.3	Alerting.....	19
5.4	SupportAssist.....	19
A	Technical support and resources	20
A.1	Related resources	20

Executive summary

This document provides best practices for using Dell EMC™ PowerStore™ with VMware® Horizon® virtual desktops. The PowerStore platform offers optimal flexibility and resiliency. 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 virtual desktop infrastructure (VDI) platform. PowerStore allows expansion of both functionality and capacity to protect storage investments.

Although the guidance in this paper is specific to VMware Horizon, it applies to most VDI environments. Each environment is unique, and all best-practices guidance that is provided in this paper may not apply.

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

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.

1 Introduction

PowerStore storage arrays are designed to enable growth, flexibility, and resiliency. PowerStore services are hardware- and protocol-independent, allowing 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. With robust reporting and notification services, storage management can be as minimal as needed. The integrated dashboards also enable at-a-glance insight into the health and performance of the platform.

1.1 PowerStore overview

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 Microsoft® Windows® SMB is integrated, with the ability to share files between Windows and UNIX® or Linux® users. This ability simplifies the control with a single interface for data management.

PowerStore can serve data using iSCSI, Fibre Channel, or both protocols simultaneously to provide seamless integration into most environments. NVMe over Fibre Channel (NVMe-FC) is a new protocol that is supported with PowerStoreOS 2.0.

1.2 VMware Horizon overview

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

1.3 Prerequisite reading

The best practices in this document require knowledge from the following resources:

- *Dell EMC PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#)
- [VMware Horizon overview](#)

1.4 Terminology

The following terms are used with PowerStore.

PowerStore Manager: The 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).

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. These volumes are displayed in PowerStore Manager under Storage > Volumes.

NAS server: Network-attached storage (NAS) servers allowing block storage to be used for file storage. The NAS service translates file requests to block storage. PowerStore NAS supports Microsoft Windows SMB 3.1.1 and NFS v3 and v4.

File system: A system that allows PowerStore to enable file-level access to end users and applications. It also controls user permissions to files, and folders or directories.

2 Connectivity

PowerStore supports iSCSI and Fibre Channel hosts. In addition, NVMe over Fibre Channel (NVMe-FC) is supported with PowerStoreOS 2.0. While these protocols can be run simultaneously, it adds complexity to routing and failover scenarios. It is recommended to use only one protocol per host.

2.1 Block storage

PowerStore presents storage to external hosts through either block or file interfaces. Block storage is the most commonly used datastore path for virtual machines due to the various speeds and protocols that are offered, making 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. Instructions and best practices for configuring hosts are in the *Dell EMC PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#).

Multipathing is critical to achieving the best performance from VMware hosts and should be configured properly. For configuration guidance, see the document *Dell EMC PowerStore: VMware vSphere Best Practices* on the [PowerStore Info Hub](#). The large number of small I/Os generated from VDI environments benefit greatly from round robin and rapid path switching. Consider using and perform testing with a low I/O-per-path change, which can help improve VM performance.

2.1.1 Boot from SAN

PowerStore supports boot from SAN for environments that need to further virtualize storage resources. Besides the storage-virtualization benefits, boot from SAN can also be used to protect the operating-system configuration and allow for quick recovery. It is beneficial in environments where the recovery time objective (RTO) is strict. Boot from SAN can also help deploy a standardized-server configuration quickly, or roll back from a failed update.

Boot from SAN can be used with any HBA that has a boot ROM that supports boot from SAN.

2.1.2 iSCSI

Using iSCSI for block storage allows convergence of storage and networking infrastructure though it does require careful planning for bandwidth requirements and fault isolation. Redundant networks are preferred for data availability.

With iSCSI, Jumbo frames are also recommended. They allow 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.

Another benefit of using the 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 also allows traffic prioritization to ensure that block traffic has the highest priority.

2.1.3 Fibre Channel

Using up to 32 Gb FC is also supported. There are many options depending on the throughput requirements, from 4 Gb to 32 Gb. VDI tends to involve a small-packet workload, so the throughput tends to be less of an issue. Dual fabrics are recommended at a minimum for availability reasons.

2.1.4 NVMe over Fiber Channel

NVMe-FC is supported with PowerStoreOS 2.0. See the most recent PowerStore documentation to verify supported operating systems and HBAs. Setup of NVMe is similar to FC. NVMe-FC offers better performance and lower latency than FC.

2.2 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).

Because each NAS instance is isolated, multiple environments can share a single system. Enforced isolation helps increase security and stability.

Because a NAS instance runs on a single node to get maximum performance, a minimum of two NAS servers is recommended. The load can then be split between the two instances. The type of data determines the best method of balancing the load.

For 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.

2.2.1 Availability

NAS servers are highly available by design. If there is a single-node failure or an in-place upgrade, the service is designed to fail over to the other node automatically. However, there are some points to be aware of.

Failback is a manual process. If a node failure does occur at some point, a maintenance event should be scheduled to move the NAS instance back to the preferred node.

Failover generally completes within 30 s to avoid host timeouts. A maintenance event should be scheduled if possible but is not required.

NAS servers are automatically moved to the peer node during upgrades. This movement should be performed outside of business hours if the IO or CPU demand on the appliance is high during the business day. Once the upgrade is complete, NAS servers should be moved back to their preferred nodes.

2.2.2 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.

2.2.2.1 SMB

The SMB file system on PowerStore supports SMB versions up to SMB3.1.1. This support allows essentially any version of Microsoft® Windows® to connect to file shares hosted on PowerStore.

The following Windows features are supported at the file-system level:

- Sync writes enabled
- Oplocks enabled (default)
- Notify on write enabled
- Notify on access enabled

At the share level, the following features are available:

- Offline availability of files
- Access-based enumeration
- Branch cache
- Continuous availability

2.2.2.2 NFS

With support of NFS v3 and v4.1, and Secure NFS, files can be shared across operating systems. With LDAP service integration, user permissions can be set across an organization.

Integrating NFS file systems with Windows users requires careful group-membership and permissions planning to ensure that users can properly share files.

Minimum security: Minimum security is allowed when connecting to an NFS export.

- Sys
- Kerberos
- Kerberos with integrity
- Kerberos with encryption

Access levels: Levels can be set for default access or individual hosts.

- No access
- Read/write
- Read-only
- Read/write, root only
- Read-only, root only

If access is shared across Windows and UNIX or Linux clients, access levels and permissions should be clearly defined and applied across exports and shares. Because of the differing permissions models, take caution when files and directories are shared among these platforms.

3 Storage design

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.

3.1 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, see the document *Dell EMC PowerStore: VMware vSphere Best Practices* on the [PowerStore Info Hub](#).

With most applications, storage traffic is typically 70% reads. There are applications that generate all reads or all writes, but the average tends to be a 70% read to 30% write ratio.

3.1.1 Volume count

There are several factors that determine the optimal volume count with Horizon on PowerStore. There is not a single volume count that fits all scenarios. Performance and management overhead are typically the two most important factors.

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

Balance this volume recommendation with the goal of keeping the volume count low for minimized management. The VMware guidance for volume count should also be considered. Internal testing shows that 16 volumes can support 3,000 users from a performance perspective but 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.

3.2 I/O size

The average I/O size of storage traffic in a VDI environment typically ranges from 24k to 32k. This size can vary if host-based caching is enabled or allowed. Some VM configurations do not allow host-based caching.

3.3 Heavy writes

VDI environments are typically heavily biased towards writes. This bias occurs because once the VMs are booted, the I/O traffic includes mostly changes from each VM. The traffic consists of file changes, swap file writes, memory paging, and updates to user preferences. This traffic pattern makes VDI one of the more demanding applications.

The heavy write ratio of VDI requires more attention to be focused on the number of writes that are generated and the rate of change to the environment. For environments that are persistent, these factors should be accounted for when snapshots are taken. The average snapshot size is larger because of the rate of change occurring.

If persistence is maintained at the VM level and profile redirection technology is not used, the VM volumes must be larger. The snapshot growth is larger as well.

3.3.1 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 VMFork technology creates the machines quickly, the volume of traffic is significant. The I/O required to complete the creation is brief but large. Plan for bursts of high traffic when creating or refreshing instant clone machines. These bursts of high traffic may impact other workloads if a PowerStore appliance is not dedicated to VDI. For a large VDI implementation, a PowerStore appliance should be dedicated to VDI.

Instant-clone machines reset on logout. As users log off, their virtual desktops refresh automatically. This action occurs throughout the workday and should be considered in the design. Virtual desktop refreshes cause a brief spike in IO that is proportional to the number of desktops refreshing at any given time. If many desktops refresh simultaneously, the I/O demand will spike as they refresh.

3.4 Capacity

The density of data on VDI tends to be high because many Horizon VDI environments use data-reduced clones. With profile-redirection technology like VMware Dynamic Environment Manager (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 advantages of small virtual desktops include faster boot times, reduced capacity requirements, reduced pool creation time, and fewer updates required. With instant clones, the administrative overhead is reduced even further with the automatic refresh of virtual desktops on logout. Any misconfiguration in a VM is cleared when the user logs out and the machine is recloneed.

Drive counts for block-only VDI are a function of IOP requirements rather than capacity. Since the VMs generate many IOPs per TB of storage that is consumed, performance is more of a concern than capacity.

In a unified-file-storage product, capacity may increase drive counts beyond performance requirements. To see the NAS limits of the PowerStore platform, see the document *Dell EMC PowerStore: Best Practices* on the [PowerStore Info Hub](#).

The number and size of the file shares or exports determine the capacity that must be added, based on performance guidance on [Dell.com/Support](#).

3.5 Guests

VMware Horizon supports two guest operating systems for desktop pools. Not all features are available with both operating systems because of the architecture of the individual operating systems. While only one image can be used per pool, multiple pools can be created, one for each virtual machine role.

3.5.1 Windows

The behavior of Windows desktops in a VDI environment can vary greatly. Because of the number of services and process in Windows, the host load can vary based on customizations that are applied to the guest operating system.

One of the best ways to reduce CPU and disk load from 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 of optimizing Windows, but it requires testing of the configuration for all user applications.

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 of applications that does not need protection is placed on volumes with no data protection enabled. This approach reduces replication traffic and the overall storage load.

3.5.2 Linux

Horizon supports virtual desktops on Linux with some caveats and restrictions on versions and distributions. For details, see the VMware product documentation about the [System Requirements for Horizon for Linux](#). The storage requirements are determined more by the applications that are supported than the base operating system. This document does not address the variations in Linux configurations.

3.6 Networking

PowerStore has options to support diverse network designs and protocols. The abstraction between the storage and front-end connectivity enables great flexibility in design.

3.6.1 Management

PowerStore is managed through the first two ports on the integrated four-port card (see Figure 1). These ports are automatically bonded for availability.




Figure 1 PowerStore management ports

These ports provide connectivity through several protocols and storage methods such as web interface (HTTP), REST API, and PowerStore CLI (PSTCLI).

An integrated web-based interface is included for most management needs (see Figure 2). The interface addresses all day-to-day management tasks and storage alerting. Automated notification can be configured for remote alerting as well.

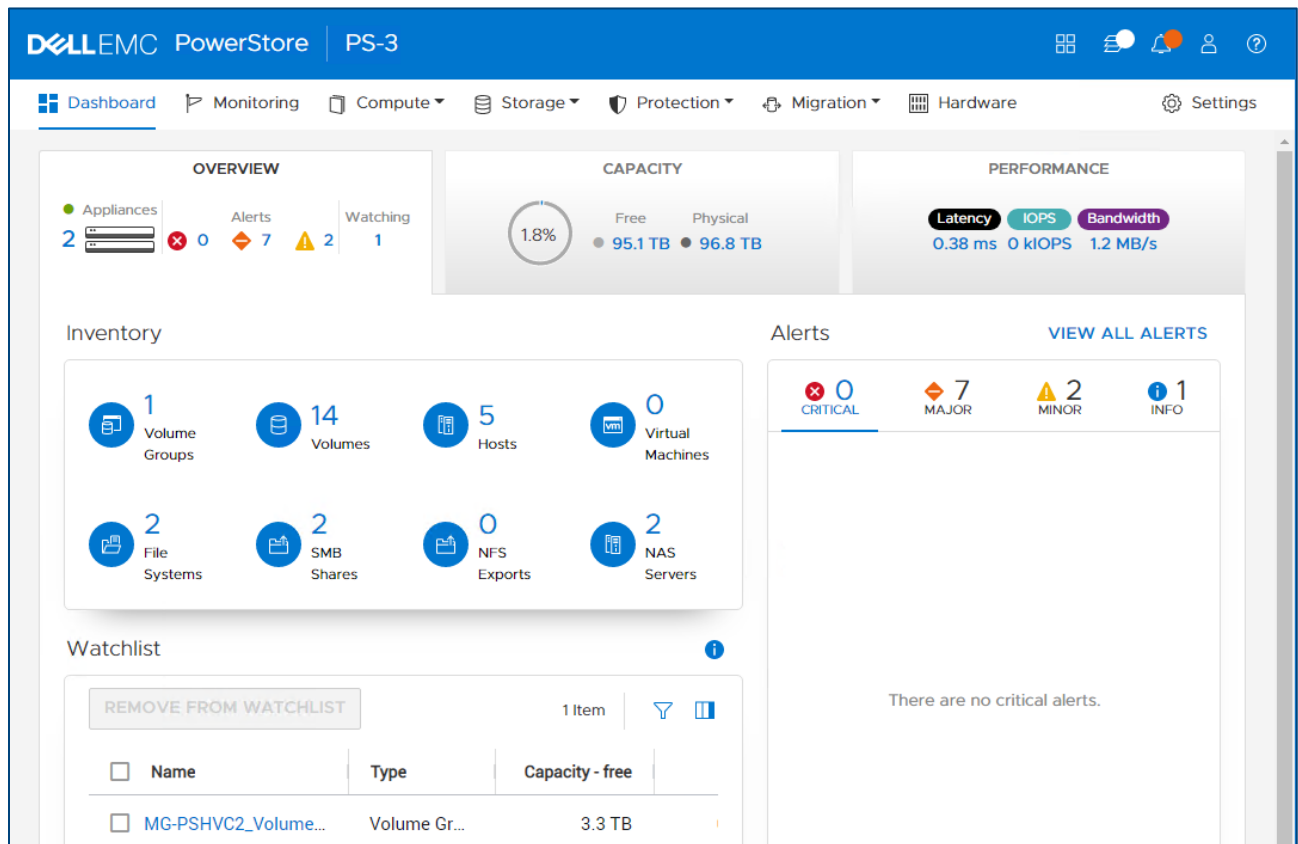


Figure 2 PowerStore management web-based interface

3.7 DRE and storage containers

PowerStore uses a dynamic resiliency engine (DRE) to automatically manage the underlying storage for maximum performance and capacity. DRE eliminates the need for administrators to configure protection settings for the storage pool. Manually setting or configuring these options is unnecessary in PowerStore. The underlying configuration and drive management are automatic.

3.8 Changing capacity needs

Planning for the future is simple with PowerStore. Dynamically scaling as drives are added is as simple as adding drives or enclosures of disks. When new drives are detected, they are automatically added to the available capacity.

Adding capacity to an existing file system is easy. Choose the new size, and the file system expands instantly. No downtime is required and there is no user impact.

Reducing the size of a file system is also simplified. If excess space is allocated to a file system, it can be reduced in size if sufficient capacity remains for existing data. With thin provisioning, no space is wasted, and file systems do not need to be reduced unless hard growth boundaries must be set.

Quotas are supported as well for growth management. This feature can prevent user directories or shared file structures from growing beyond the capacity of the system.

3.9 Horizon configuration

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

Volumes that are presented through PowerStore can grow dynamically. If the original volumes must grow as the virtual machine requirements change, they can be expanded using the standard management methods. This expansion is done using a PowerCLI script (RESTAPI) that also expands the datastores. PowerStore Manager is used to expand a single volume.

3.10 Replication

The nature of VDI environments is to create pools of VMs that are not necessarily tied to a specific user. This ability reduces the administration of the environment, and also reduces the need for backups, disaster recovery, and data-loss mitigation.

This design also reduces the need for replication. Since data is typically not stored in the user VMs, the VMs do not need to be copied offsite. A VM pool can quickly be created offsite if the infrastructure is available, including the template VMs.

User data from file shares can be replicated in several ways. PowerStore uses protection policies to take snapshots and replicate data (see Figure 3).

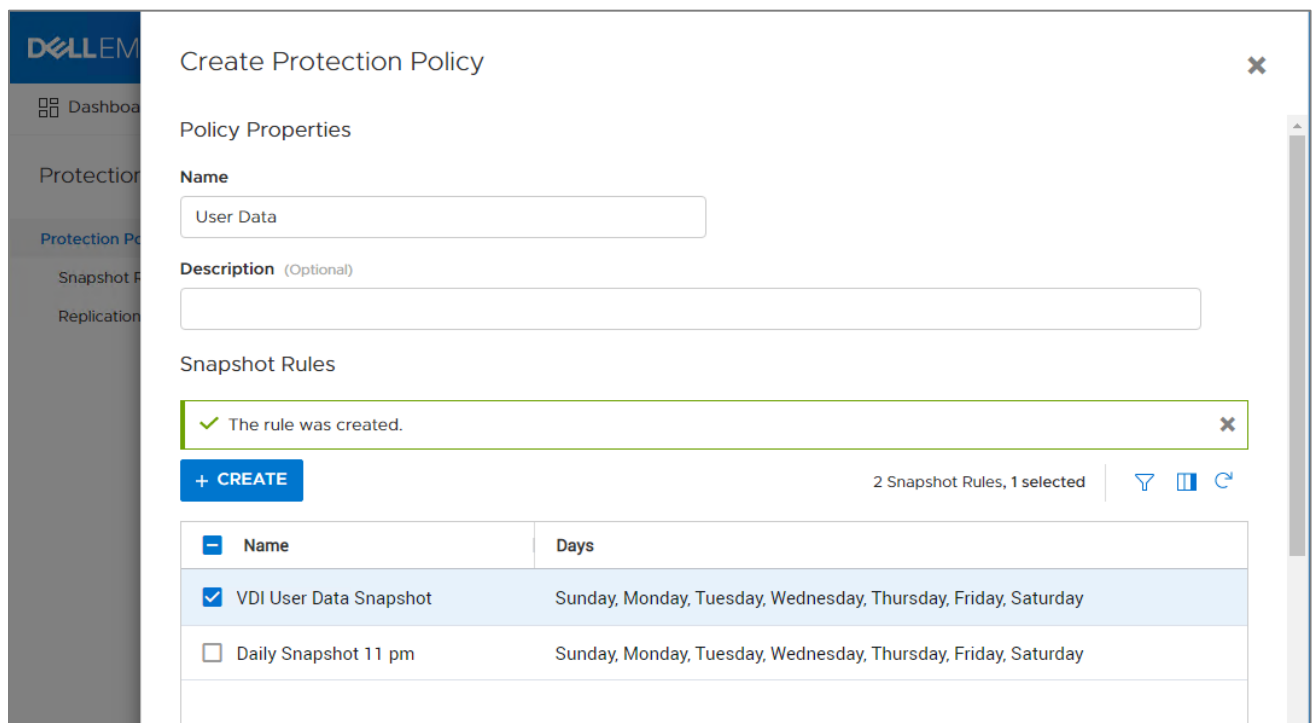


Figure 3 PowerStore protection policies

Protection policies can be created for file or block data. For NAS data, a protocol snapshot policy can be created to ensure that redirected user data is protected. The snapshots can be used in a read-only mode for deleted or damaged file recovery without impacting active file systems.

Once a protection policy has been created to take snapshots at the required interval, it can be replicated to another system.

3.11 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.

4 Data encryption

Many applications have data encryption requirements on data at rest. Data at Rest Encryption (D@RE) can be used as an encryption solution for VDI without requiring VM changes. This ability also avoids potential performance impact to the applications and has no performance impact on the array.

D@RE is enabled by default on the PowerStore array, so no configuration steps are necessary to protect the drives.

5 Management and monitoring

PowerStore offers multiple features for managing and monitoring storage. There are several methods to gather and process alerts and statistics.

5.1 Data reduction

The PowerStore data reduction feature includes zero detection, compression, and deduplication built into the core product. While the amount of reduction varies depending on the type of data, the system automatically selects the best option for reducing the data footprint.

Virtual desktops can vary greatly in data-reduction rates based on the type of VM pool. Linked clones (supported with Horizon 7 and earlier) have the lowest reduction ratio because of their inherent reduction. The pool is already compact because of the linking architecture.

Full clones have the largest reduction ratio of the pool types because they include many identical pages of data. Full clones have the highest impact on storage since they involve a larger data volume and require ongoing operating-system updates to all VMs.

5.2 Scripting and automation

Scripting and automation are often used to improve the quality of repetitive tasks and increase the speed of deployment. The PowerStore platform supports several different tools for administrators who want to automate management tasks or implement Infrastructure as Code (IaC).

5.2.1 REST API

The PowerStore REST API provides a complete interface to perform automation or custom applications for the PowerStore platform. Go to **<https://<mgmt ip>/swaggerui>** in a browser to access an interactive API reference for discovering available commands and viewing sample usage.

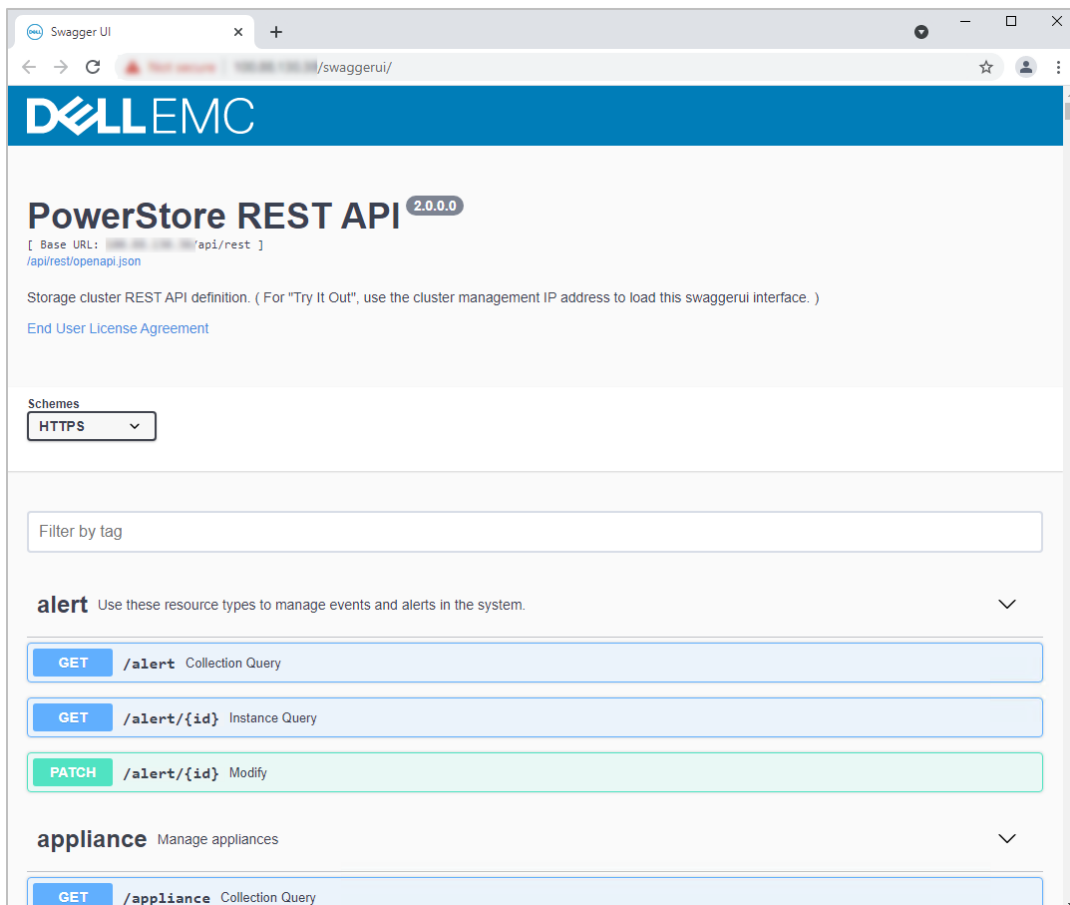


Figure 4 PowerStore REST API

5.2.2 PowerStore CLI

The PowerStore CLI (PSTCLI) is a command-line interface that is provided for managing PowerStore systems. The PSTCLI provides a convenient way to manage PowerStore without the complexity of using a programming interface.

5.3 Alerting

Email alerting is an integrated function for using external notification systems. By configuring email addresses for notification events, exceeding-notification criteria generates an email.

5.4 SupportAssist

Dell SupportAssist can provide an extra level of support by automatically contacting support resources when events occur. This feature can speed up the resolution process by transmitting critical log and event information to escalation resources.

A 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 [PowerStore Info Hub](#) provides detailed documentation on how to install, configure, and manage Dell EMC PowerStore systems.

A.1 Related resources

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