

Dell PowerStore: Microsoft SQL Server Best Practices

July 2022

H18250.7

White Paper

Abstract

This document includes architecture, deployment, and performance guidance for Microsoft SQL Server and Microsoft Data Platform with Dell PowerStore.

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 © 2020-2022 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 July 2022 H18250.7.

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**
- Introduction 5**
- Scale-up and scale-out architecture..... 6**
- Host configuration 7**
- Sizing and performance optimization 8**
- Performance 18**
- Management and monitoring 27**
- Data protection and disaster recovery 33**
- References 36**

Executive summary

Overview

This paper provides best practices for using Dell PowerStore in a Microsoft SQL Server environment. SQL Server is a robust product that can be used in various solutions. The relative priorities of critical design goals such as performance, manageability, and flexibility depend on the environment. This paper provides considerations and recommendations to help meet design goals. For general best practices for using PowerStore systems, see the document [Dell PowerStore: Best Practices Guide](#).

These guidelines are intended to cover most use cases. They are recommended by Dell Technologies, but they are not strictly required. For questions about the applicability of these guidelines in your environment, contact your Dell Technologies representative to discuss the appropriateness of the recommendations.

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 networked storage environment using PowerStore systems.

Revisions

Date	Description
April 2020	Initial release: PowerStoreOS 1.0
July 2020	Updated AppsON host configuration and sizing guidance
February 2021	Legal disclaimer update
April 2021	PowerStoreOS 2.0
May 2021	Clarification on PowerStore X and T
October 2021	Updated NVMe terminology
October 2021	Updated template
July 2022	PowerStoreOS 3.0

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

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: Doug Bernhardt

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

Introduction

PowerStore and SQL Server

Dell PowerStore is a robust and flexible storage and compute option that is well suited to SQL Server. The capabilities of the PowerStore platform provide some unique benefits and exciting architecture options for SQL Server workloads.

PowerStore product 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.

Microsoft SQL Server product overview

Microsoft SQL Server is the industry-leading data platform. A product which once started as a database engine on Microsoft Windows has evolved into an entire data platform. This platform enables cloud, hybrid, and on-premises deployments in several different environments including Windows, Linux, containers, and Kubernetes. Besides the wide number of environments that SQL Server can be deployed in, the use cases have expanded. These uses now range from typical online transactional processing (OLTP) to various analytics options such as data warehousing, artificial intelligence (AI), machine learning (ML), and big data. The adaptable and flexible nature of SQL Server makes PowerStore a perfect fit for SQL Server.

Prerequisite reading

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

- PowerStore documentation on [Dell.com/powerstoredocs](https://www.dell.com/powerstoredocs)
- *Dell PowerStore Host Configuration Guide* at [Dell.com/powerstoredocs](https://www.dell.com/powerstoredocs)
- [Microsoft SQL Server documentation library](https://docs.microsoft.com/en-us/sql/)

Terminology

The following terms are used with PowerStore.

Table 1. Terminology

Term	Definition
Appliance	Term used for solution containing a base enclosure and any attached expansion enclosures. The size of an appliance could be only the base enclosure or the base enclosure plus expansion enclosures.
Base enclosure	Used to reference the enclosure containing both nodes (node A and node B) and the NVMe drive slots.

Term	Definition
Cluster	Multiple PowerStore appliances in a single grouping. Clusters can consist of one appliance or more. Clusters are expandable by adding more appliances. A cluster supports up to four appliances.
Expansion enclosure	Enclosures that can be attached to a base enclosure to provide 25 additional SAS-based drive slots.
Node	The component within the base enclosure that contains processors and memory. Each appliance consists of two nodes.
PowerStore Manager	The web-based user interface (UI) for storage management.
Storage volumes (volumes)	PowerStore volumes using block storage. These volumes are displayed in the Block area of the PowerStore dashboard.
Virtual Volumes (vVols)	PowerStore X model volumes support VMware vSphere Virtual Volumes (vVols). These volumes are displayed in the Virtualization area of PowerStore Manager.
Watchlist	Configurable list of volumes or hosts that are available on the dashboard for quick reference and monitoring.

Scale-up and scale-out architecture

Introduction

Scale-up and scale-out are common architecture models used when accommodating a growing data footprint. Microsoft SQL Server has supported both scale-up and scale-out architecture models for several years. Dell PowerStore aligns with the SQL Server by also providing both scale-up and scale-out architecture options. When designing data architectures, there are tradeoffs in design, functionality, and scalability that must be considered.

SQL Server scale-up

When you scale up an instance of SQL Server, you can increase the amount of memory, compute, and storage by upgrading to a higher software version or increasing the hardware capabilities. You get the same or more functionality and retain the same logical data design and therefore the data architecture is unaffected. This approach is common because it is low impact.

SQL Server scale-out

Scale-out architectures on SQL Server enable a data footprint to scale past what can be achieved by a single instance. The trade-off is that these architectures require considerably more planning and consideration and often require applications to be designed accordingly. Concepts such as data placement, load balancing, and data sharding may need to be addressed, and the impact on features such as replication or availability groups.

PowerStore scale-up

PowerStore models can scale up with several different upgrade paths. You can add drives or drive expansion shelves to an appliance, add storage network paths, upgrade to next-generation hardware as it becomes available, or upgrade to more powerful models within the same family. Like SQL Server scale-up, this approach is also low impact. PowerStore

redundancy features and the Anytime Upgrades program offered by Dell Technologies ensure a frictionless upgrade experience.

PowerStore scale-out

PowerStore can also scale-out by adding nodes to the PowerStore cluster. PowerStore can add up to four appliances in a cluster to provide extra capacity and performance. Like scale-up or scale-out architecture decisions for SQL Server, there are similar considerations for PowerStore. See the document [Dell PowerStore: Clustering and High Availability](#) for complete details.

Host configuration

Introduction

PowerStore provides the option of running workloads directly on the PowerStore appliance as virtual machines or running workloads on external hosts and presenting storage through block or file interfaces.

Before you create objects or deploy workloads on PowerStore, we recommend reviewing the following guides for best practices and tuning recommendations. We recommend applying these changes before deploying workloads since some changes require a reboot of the PowerStore appliance nodes. For resiliency, ensure that external hosts are connected through multiple paths, or multipath, to the PowerStore appliance.

External SQL Server hosts

PowerStore presents storage to external hosts through either block or file interfaces. Block storage is commonly used for SQL Server workloads due to the various speeds and protocols that are offered which make it ideal for performance. With PowerStoreOS 2.0, PowerStore supports 25 GbE and 32 Gb Fibre Channel (FC) and NVMe over Fibre Channel (NVMe/FC). These options are recommended for SQL Server high-bandwidth storage workloads such as analytics that can generate a large amount of large block I/O. Instructions and best practices for configuring hosts can be found in the *Dell PowerStore Host Configuration Guide* at Dell.com/powerstoredocs. File storage (SMB) can also be used in environments running SQL Server on Windows.

For PowerStore T model appliances, follow recommendations for configuring external hosts in the *Dell PowerStore Host Configuration Guide*. In addition, when running SQL Server on virtual machines with VMware ESXi hosts, review the guidance in the documents [Dell PowerStore Virtualization Infrastructure Guide](#) and [Dell PowerStore: VMware vSphere Best Practices](#).

Boot from SAN

PowerStore supports boot-from-SAN for environments that want to further virtualize storage resources. In addition to the storage virtualization benefits, Boot-from-SAN can also be used to protect the operating-system configuration and allow for quick recovery. This configuration can be beneficial in environments where the recovery time objective (RTO) is strict.

AppsON

Integration of the PowerStore container-based architecture with integrated VMware ESXi results in a new level of consolidation for enterprise storage. This integration combines the benefits of a local on-array application environment with unmatched integration with the vSphere management environment and server resources. This ability allows users to bring applications closer to storage, by running applications as virtual machines directly on PowerStore.

Benefits of the AppsON capability include a new level of agility for application deployments, with seamless movement between the PowerStore appliances and VMware ESXi servers. AppsON also provides the ability to shrink the stack by eliminating server and networking footprint for space-efficient edge and remote deployments. AppsON is ideal for data-intensive applications that require low latency or a storage-heavy imbalance of compute and storage.

AppsON can be an attractive option whether the goal is to optimize I/O performance by reducing components in the data path, add performance to an existing vSphere cluster, offload a noisy neighbor, or other reasons.

For PowerStore X model appliances, the [PowerStore X Performance Best Practice Tuning](#) paper provides several tips that increase performance for SQL Server workloads. These tips are especially helpful if you are running workloads that use block sizes greater than 128k such as analytic workloads and backup/restore operations. Besides the recommendations in this paper, when running workloads larger than 128k, you can find other performance gains by increasing the iSCSI maximum I/O size, which is 128k by default. To enable this result, run the following command on each host:

```
esxcli system settings advanced set -o /ISCSI/MaxIoSizeKB -i 512
```

Instructions for running esxcli commands on PowerStore X nodes are in the [PowerStore X Performance Best Practice Tuning](#) paper.

Sizing and performance optimization

Introduction

There are several best practices for provisioning compute and storage to SQL Server. The size of power and flexibility of SQL Server allow it to be used in several different deployment scenarios. This section covers some considerations for optimizing PowerStore.

SQL Server design considerations

The I/O storage system is a critical component of any SQL Server environment. Sizing and configuring a storage system without understanding the I/O requirements can have unfavorable results. Analyzing performance in an existing environment using a tool like [Live Optics](#) can help define the I/O requirements. For best results, capture performance statistics for a period of at least 24 hours that includes the system peak workload.

Demanding SQL Server workloads that require low latency, high bandwidth, or both can benefit from moving the compute to PowerStore X model arrays. Moving these intensive workloads onto the storage platform where compute and storage are located together can deliver higher I/O performance for performance-sensitive applications. Also, virtual machines running other application components can be on PowerStore X models for

further performance gains. Using PowerStore X models in a larger vSphere cluster and using VMware vSphere vMotion capabilities can be a powerful solution to balance SQL Server workloads and application workloads across the entire environment.

Database design

When creating SQL Server databases, the default is to have one database file, one filegroup, and one log file. If at any point the database will have high performance requirements, it is recommended to use at least two data files per file group. SQL Server will write data evenly across all files in the filegroup and therefore the underlying files. Using multiple files allows the database workload to be spread across multiple volumes or vVols. Even if they are placed on the same volume in the beginning, it is advantageous to create this configuration from the start. More files can be added later, but the data must be restriped to achieve balanced access through re-indexing or some other technique. This reason is because SQL Server uses a proportional fill strategy when writing to data files such that it will write more data to files that have a greater amount of free space.

More filegroups can be created and may be beneficial to use some SQL Server features. However, the number of filegroups has no impact on performance from a storage perspective.

SQL Server log files are accessed sequentially and do not use proportional fill. Therefore, there is no performance benefit to creating extra log files.

OLTP workloads

While every environment is unique, an online transaction processing (OLTP) workload typically consists of small, random reads and writes. A storage system that services this type of workload is primarily sized based on capacity and the number of IOPS required. PowerStore models allow the flexibility to be configured for block-optimized, unified (block and file), or AppsON to meet unique OLTP requirements.

Analytic workloads

An online analytic processing (OLAP), decision support system (DSS), or big data workload is typically dominated by large, sequential reads. A storage system that services this type of workload is primarily sized based on throughput. When designing for throughput, the performance of the entire path between the server and the drives in PowerStore needs consideration. This is handled automatically when using virtual volumes on PowerStore X models. When mapping volumes to hosts, consider using 32 Gb Fibre Channel (FC) or 25 Gbps iSCSI connectivity to the array. For high-throughput requirements to be met, multiple HBAs may also be used in the server, the array, or both.

Mixed workloads

The most common scenario is a mixed workload environment. Typically, SQL Server I/O patterns do not strictly fall into an OLTP or analytics pattern. This factor is what can make SQL Server workloads challenging because no two workloads behave the same. In addition, the same SQL Server host or instance may be servicing multiple applications or transaction workloads. Mixed workload can also imply that multiple applications (in addition to SQL Server) are residing on the same host or accessing the same storage. The combined workload of these applications invalidates any typical application I/O usage

pattern. For these reasons, it is important to gather performance metrics for best sizing results.

SQL Server deployment platforms

SQL Server is now available and supported on several different platforms. PowerStore presents several flexible options for various SQL Server deployments.

Windows

When deploying SQL Server on Windows on PowerStore, virtual volumes (vVols) and storage volumes (volumes) are available on PowerStore X models. In addition to vVols and volumes, PowerStore T also contains file support through SMB. All these options are acceptable for use with SQL Server and rely on the configuration of the volumes through Windows. When using SMB storage for SQL Server, ensure the Sync Writes Enabled and Oplocks Enabled are turned on for the file system. See the Windows section of the *PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#) for instructions about configuring host access on Windows.

SQL Server on Windows deployment

SQL Server on Windows is deployed with an executable (exe) file which can be found on the Microsoft SQL Server website. Deployment is done by downloading the SQL Server exe file for the intended version and running the executable. The SQL Server deployment wizard will walk through several available options. Storage layout can either be customized during the install or after install once the SQL Server instance is running. When performing large enterprise deployments, there is an unattended install option available to automate the installation.

Linux

When deploying SQL Server on Linux on PowerStore, vVols and volumes are available on both PowerStore models. Although PowerStore T models contain file capabilities, they do not support NFS v 4.2 which is the minimum that is required for SQL Server on Linux. Therefore, NFS is not supported.

Be sure to follow [Microsoft Installation guidance for deploying SQL Server on Linux](#) for supported Linux versions, minimum requirements, and general recommendations. Currently, Red Hat Enterprise Linux, SUSE, and Ubuntu Linux distributions are supported. The partnership of Dell Technologies, Microsoft, and RedHat provides a platform with full end-to-end enterprise support. Therefore, Red Hat Enterprise Linux has quickly become a platform of choice for deploying SQL Server on Linux.

SQL Server on Linux deployment

Once the Linux host is presented to the array, storage volumes can be presented to the host. Since Linux is a relatively new platform for SQL Server, an example is provided below for configuring a Red Hat Enterprise Linux host to consume PowerStore block volumes. This example uses multipath, which is a best practice for resiliency. For complete instructions about configuring Linux hosts and multipath, see the *PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#).

After the storage volumes have been provisioned to the host, run the following command to make the volumes visible to the host:

```
#/usr/bin/rescan-scsi-bus.sh -a
```

Identify the volume name in the mapper folder by listing the multipath volumes and searching for the volume WWN from PowerStore. The volume WWN is visible in the PowerStore volumes list under Storage > Volumes. For example, if the WWN for the new volume is **68ccf0980018e6c92364df1970e08d33**, the following command will identify the mapper folder to use:

```
# multipath -ll | grep 68ccf0980018e6c92364df1970e08d33
```

The volume will need to be formatted before use. See the Microsoft SQL Server documentation for supported file systems. Currently, XFS and EXT4 are supported. In this example, XFS is being used. For XFS, the `mkfs.xfs` command is used to format the volume passing the correct mapper folder as the parameter.

```
# mkfs.xfs /dev/mapper/mpathdb
```

A mount folder needs to be created to represent the new volume. Create a folder to represent the mount:

```
# mkdir /var/opt/mssql/data5
```

To make the mount persistent, the mapping needs to be stored in `/etc/fstab` so it will be created at boot time. The entry in the file will consist of six fields contained on a single line separated by a space: device, mount point, file system type, options, backup operation, and file system check order. When mounting volumes for SQL Server, it is a Microsoft best practice to add the `noatime` attribute to the list of options to disable updates to the volume timestamp for better performance. The fields *backup operation* and *file system check order* will be left at the default of 0 and can be changed as needed. Based on our example, the new entry in the `/etc/fstab` file would look like this:

```
/dev/mapper/mpathdb /var/opt/mssql/data5 xfs defaults,noatime 0 0
```

Once the entry for the new mount is created, running the `mount` command will read the file and refresh the mounted devices. Even if the mount was manually created, it is a good idea to run the `mount` command after updating `/etc/fstab` to ensure there are no errors. If there are errors in the file, it could prevent Linux from booting properly.

```
# mount -a
```

Finally, before SQL Server can access the new mount, permissions need to be set. In a production environment, follow the security best practices for your organization when assigning permissions and ownerships to mount points. This example uses the same default permissions and ownership from the SQL Server installation by copying the permissions from the `/var/opt/mssql/data` folder and applying them to the new mount point:

```
# chmod --reference /var/opt/mssql/data /var/opt/mssql/data5
```

At this point, the new volume is ready for SQL Server to use. See the Linux section of the *PowerStore Host Configuration Guide* on the [PowerStore Info Hub](#) for further instructions on configuring storage volume access on Linux.

Containers

Running SQL Server inside Docker containers is a deployment option that was introduced with SQL Server 2017. By default, the database storage is inside the container and is ephemeral in nature. If databases need to be persisted beyond the lifetime of the container, external storage must be presented. External storage is presented to a container through a bind mount which is a local file system folder or mount from the operating system that is presented to the container at startup. Follow the storage guidance above for either Windows or Linux to create file system folders or mount points to present as a volume bind mounts to containers.

SQL Server container deployment

Deploying SQL Server inside Docker containers is simple. Below is an example of running SQL Server inside a container followed by a brief explanation of the parameters:

```
# sudo docker run -e "ACCEPT_EULA=Y" -e "SA_PASSWORD=P@ssword1" -p
1401:1433 -v /root/sql1/tpcc_data:/var/opt/mssql/data2 -v
/root/sql1/tpcc_log:/var/opt/mssql/log2 --name sql1 -d
mcr.microsoft.com/mssql/server:2019-latest
```

Table 2. Docker run parameters

Parameter	Description
-e	Sets environment variables, in this case options for licensing and the SA password are set
-p	Publishes the SQL Server port 1433 inside the container as 1401 on the host
-v	Bind mount for a volume. In this example, three are specified. The format is local mount:container mount
--name	Name of the container, otherwise it will be assigned a random identifier
-d	Container image to run. If it does not exist, it will try to locate it from any configured repositories and download it

Kubernetes

Due to the clustered architecture of Kubernetes, there are special storage considerations when running SQL Server instances or SQL Server Big Data Clusters on Kubernetes. Kubernetes has created its own method of dealing with persistent storage in a Kubernetes cluster using persistent volumes.

Dell Container Storage Modules (CSM) for PowerStore enable integration and automation of PowerStore storage resources to Kubernetes. PowerStore storage integrates with Kubernetes through the PowerStore CSI driver. The PowerStore CSI driver supports storage volumes only. Kubernetes clusters running on PowerStore X models can either use the vSphere Cloud Provider for virtual volumes or the PowerStore CSI driver for storage volumes. The PowerStore CSI driver and CSI drivers for other Dell Technologies products can be found on [GitHub](#).

Whenever possible, perform storage operations through Kubernetes. In some cases, modifying CSI volumes directly on the PowerStore appliance can cause CSI operations to

fail or volumes to be orphaned. For example, CSI volumes added to a volume group need to be removed from the volume group before the CSI driver can delete them.

In addition to the CSI driver support for PowerStore, there are PowerStore CSM modules available for observability and replication. The PowerStore CSM Observability module provides storage performance metrics to Kubernetes that are stored in Prometheus and can be viewed with tools such as Grafana. This allows full end-to-end visibility of storage metrics without the involvement of the storage administrator.

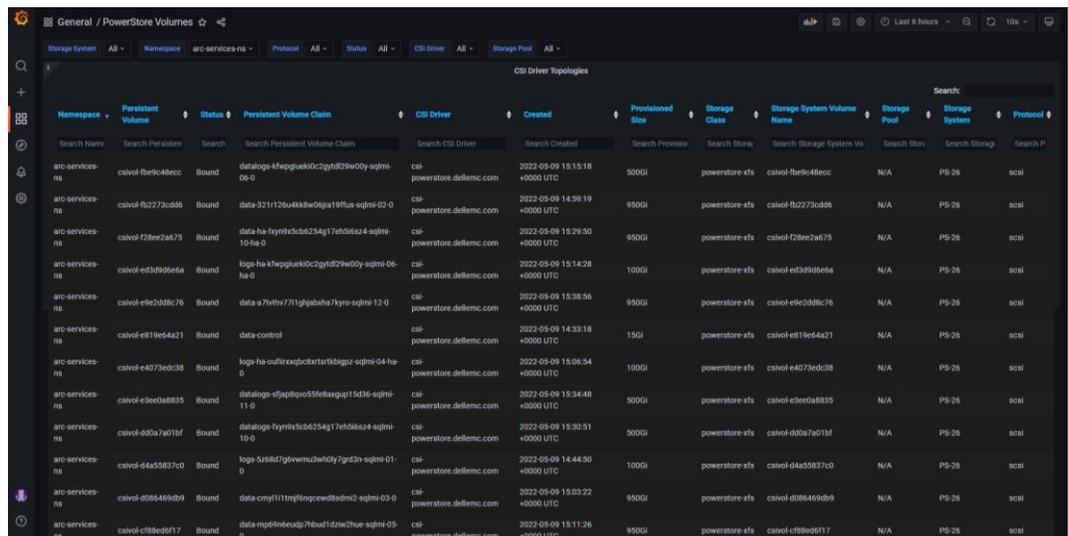


Figure 1. PowerStore volume topology in Grafana



Figure 2. PowerStore volume performance metrics in Grafana

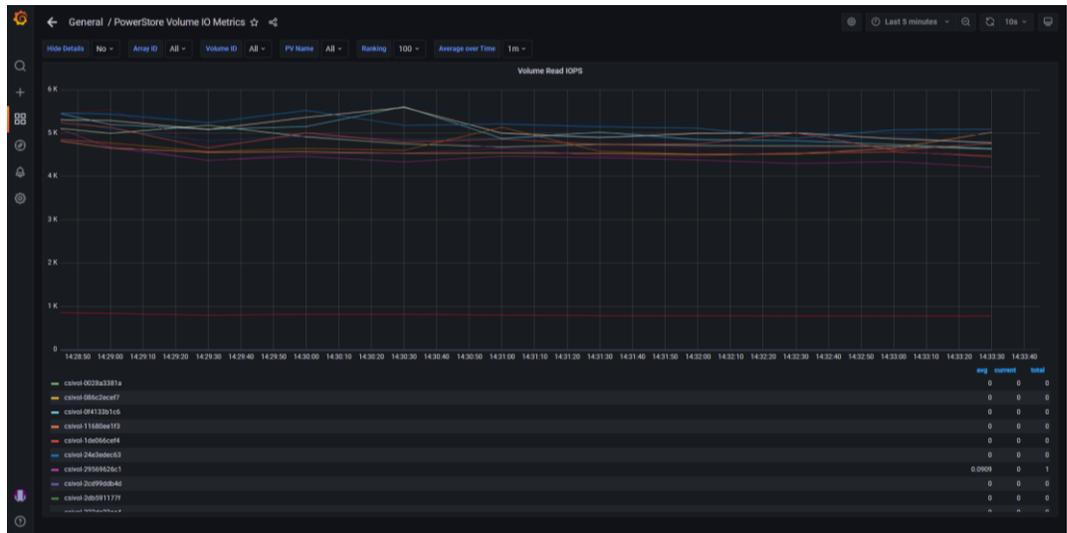


Figure 3. PowerStore read IOPS by volume in Grafana

The PowerStore CSM Replication module allows control of array-based replication features from the K8s control plane. Volume can be replicated and failover and reprotect operations can be performed.

The capabilities of the Kubernetes CSI specification are constantly evolving to support new enterprise storage features. Use the latest Dell Technologies PowerStore CSI driver and keep up to date on new releases. The latest documentation for Dell Technologies CSM modules for all Dell storage products can be found at <https://dell.github.io/csm-docs/docs/>.

SQL Server Kubernetes deployment

PowerStore storage for SQL Server running on Kubernetes is deployed and managed almost entirely through Kubernetes. The PowerStore CSI driver translates Kubernetes storage provisioning commands into PowerStore storage provisioning commands that are orchestrated on the storage appliance. Since common storage provisioning tasks such as creating, deleting, and expanding volumes are done through Kubernetes, the DBA or Kubernetes administrator can provision storage without knowing the details of PowerStore or requiring direct access. This ability enables self-service of common storage provisioning tasks.

Once the PowerStore CSI driver is deployed, one or more Kubernetes storage classes will be created. These storage classes describe the details of the CSI driver, encapsulate the connectivity and protocol options of the storage, and contain defaults for common storage provisioning values. Below is an example of powerstore-xfs storage class definition:

```
allowVolumeExpansion: true
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    meta.helm.sh/release-name: powerstore
    meta.helm.sh/release-namespace: csi-powerstore
  creationTimestamp: "2020-10-07T13:33:14Z"
```

```

labels:
  app.kubernetes.io/managed-by: Helm
name: powerstore-xfs
resourceVersion: "221004"
selfLink: /apis/storage.k8s.io/v1/storageclasses/powerstore-xfs
uid: cd013881-522e-4e16-bbd0-1f3c31315c94
parameters:
  FsType: xfs
provisioner: csi-powerstore.dellemc.com
reclaimPolicy: Delete
volumeBindingMode: Immediate

```

Kubernetes storage is provisioned as Persistent Volumes and Persistent Volume Claims. More information about Persistent Volumes and Persistent Volume Claims can be found at [Kubernetes.io](https://kubernetes.io). For example, to create three volumes for a SQL Server pod running on Kubernetes, the definitions are stored in a YAML file:

```

# cat pvc.yaml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: mssql-data
  annotations:
    volume.beta.kubernetes.io/storage-class: powerstore-xfs
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 8Gi
---
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: mssql-data2
  annotations:
    volume.beta.kubernetes.io/storage-class: powerstore-xfs
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 5Gi
---
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: mssql-log2
  annotations:
    volume.beta.kubernetes.io/storage-class: powerstore-xfs

```

```
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 2Gi
```

To create the persistent volume claims, the YAML file above is used as input. Running the following command will create the volumes on PowerStore and create a volume alias that can be used when creating the SQL Server pod.

```
#kubectl create -f pvc.yaml
```

Next, to deploy SQL Server a definition file is also used. The `persistentVolumeClaim` attributes in the definition file refer to the Persistent Volume Claims created in the previous step.

```
#cat sql.yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: mssql-deployment
spec:
  replicas: 1
  selector:
    matchLabels:
      app: mssql
  template:
    metadata:
      labels:
        app: mssql
    spec:
      terminationGracePeriodSeconds: 30
      hostname: mssqlinst
      securityContext:
        fsGroup: 10001
      containers:
        - name: mssql
          image: mcr.microsoft.com/mssql/server:2019-latest
          ports:
            - containerPort: 1433
          env:
            - name: MSSQL_PID
              value: "Developer"
            - name: ACCEPT_EULA
              value: "Y"
            - name: SA_PASSWORD
              valueFrom:
                secretKeyRef:
                  name: mssql
                  key: SA_PASSWORD
```

```

    volumeMounts:
      - name: mssqldb
        mountPath: /var/opt/mssql
      - name: mssqldata
        mountPath: /var/opt/mssql/data2
      - name: mssqllog
        mountPath: /var/opt/mssql/log2
    volumes:
      - name: mssqldb
        persistentVolumeClaim:
          claimName: mssql-data
      - name: mssqldata
        persistentVolumeClaim:
          claimName: mssql-data2
      - name: mssqllog
        persistentVolumeClaim:
          claimName: mssql-log2
---
apiVersion: v1
kind: Service
metadata:
  name: mssql-deployment
spec:
  selector:
    app: mssql
  ports:
    - protocol: TCP
      port: 1433
      targetPort: 1433
  type: NodePort

```

To deploy SQL Server inside of a Kubernetes pod, first a secret is created to store the password:

```
#kubectl create secret generic mssql --from-literal=SA_PASSWORD="P@ssword1"
```

Next, using the YAML file above as input, kubectl create will deploy the SQL Server instance on Kubernetes:

```
# kubectl create -f sql.yaml
```

Azure Arc-enabled SQL Managed Instance

Arc-enabled SQL Managed Instance (MI) allows SQL Server to be deployed from the Azure portal to an on-premises Kubernetes cluster. Since Arc-enabled SQL MI deploys an availability group, shared (read-write many) storage is required in addition to block volumes. PowerStore with combined block and file is uniquely suited to run these workloads on a single appliance. Dell Technologies and Microsoft performed a joint performance and scalability study on running Arc-enabled SQL MI on PowerStore. Complete details can be found in the paper [Dell PowerStore with Azure Arc-Enabled Data Services](#).

RAID configurations and storage pools

Since the inception of storage devices that used multiple drives for a single volume, the topic of RAID configuration and storage pools has been highly debated. Designing and maintaining the proper design could be time consuming and the wrong decisions could have a negative impact on performance. In addition, reconfiguring the storage layout often required considerable planning.

PowerStoreOS uses the capabilities of flash-based storage for maximum performance and capacity, eliminating the need for administrators to configure RAID or storage pools. Most existing SQL Server publications that discuss RAID levels and configuration were written based on older technology and apply to spinning drives. Therefore, these concepts do not apply to PowerStore. Dell Technologies recommends following the guidance in this paper and associated publications for best performance and availability. For further questions, contact a Dell Technologies representative.

Validating the I/O path

Before deploying workloads on PowerStore, it is a good idea to validate the I/O path between the server and the array. Running a large block sequential read test using small files should saturate the path between the server and the array. This test verifies that all paths are fully functional and can be used for I/O traffic. Run this test on an isolated server and PowerStore. Running this test in a production environment could cause significant performance issues for active workloads. To validate the I/O path, run a large block sequential read test using the following guidelines:

- Create one volume per node.
- Format the volumes using a 64 KB allocation unit.
- Use a block size of 512 KB for the test.
- Configure the test for 32 outstanding I/Os.
- Use multiple threads; eight is the recommended starting point.

If the displayed throughput matches the expected throughput for the number of HBA ports in the server, the paths between the server and PowerStore are set up correctly.

MPIO

A best practice for SQL Server hosts is to provide multiple paths to the storage for resiliency and performance. This practice protects against possible data loss or downtime, should a physical component fail. In scenarios where multiple paths to the storage are used, MPIO must be enabled and configured. See the document *Dell PowerStore Host Configuration Guide* at Dell.com/powerstoredocs for recommendations.

Performance

Overview

Managing storage performance with PowerStore is simple and intuitive. The performance policy can be adjusted at the volume or vVol level by selecting high, medium, or low. Performance policies are relative in PowerStore, so leaving all volumes at the default of medium provides equal performance to all volumes. The volumes all perform at the highest level until a PowerStore volume is set to a different performance policy.

Selecting a performance policy sets a relative priority for the volume using a share-based system such that when resources are constrained, the resource with the higher

PowerStore Node Affinity

performance policy receives priority. If there is no PowerStore resource constraint, the volumes perform equally.

In some cases, SQL Server can be very storage-resource intensive. For I/O intensive SQL Server databases, multiple volumes are recommended. This use allows the flexibility to balance the workload across both nodes in a PowerStore appliance. The “ownership” of a volume by a node is known as Node Affinity.

PowerStore performs best when the system performance is balanced across both nodes in the appliance. The Node Affinity is set when the volumes are mapped to a host. By default, volume Node Affinity alternates between the two nodes in the appliance when mapping volumes to a host to balance the workload. Usually, this process sufficiently balances the load across both nodes in the appliances. Since every SQL Server workload is unique, and storage access patterns depend on database layout across one or more files, Node Affinity adjustments may be necessary to maximize overall system performance.

Overall system and individual appliance performance can be viewed in the Hardware section on the Performance tab. Scrolling down in the view will reveal performance by node.

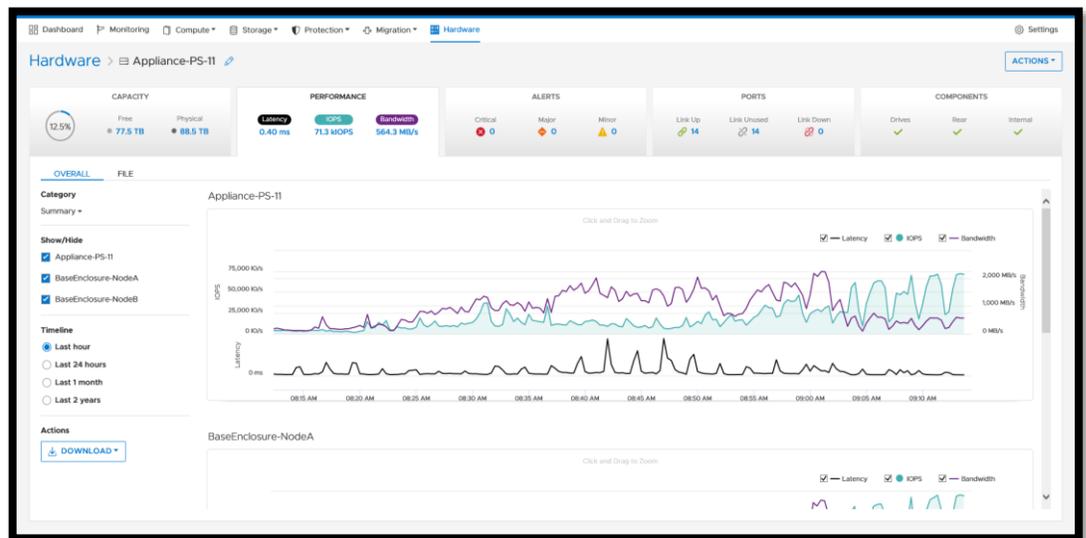


Figure 4. PowerStore hardware performance

Ideally, each node in the appliance will have equal performance. If the workload is unbalanced between the nodes, you can evaluate the Node Affinity of the workload and adjust.

Note: Node Affinity can only be adjusted on block volumes, vVols are system select only.

Volume node affinity is shown in the Volumes list under Storage. By default, this column is not displayed. To add it to the view, click the Columns selection icon next to the filter icon. Scroll down the list, and select **Node Affinity**.

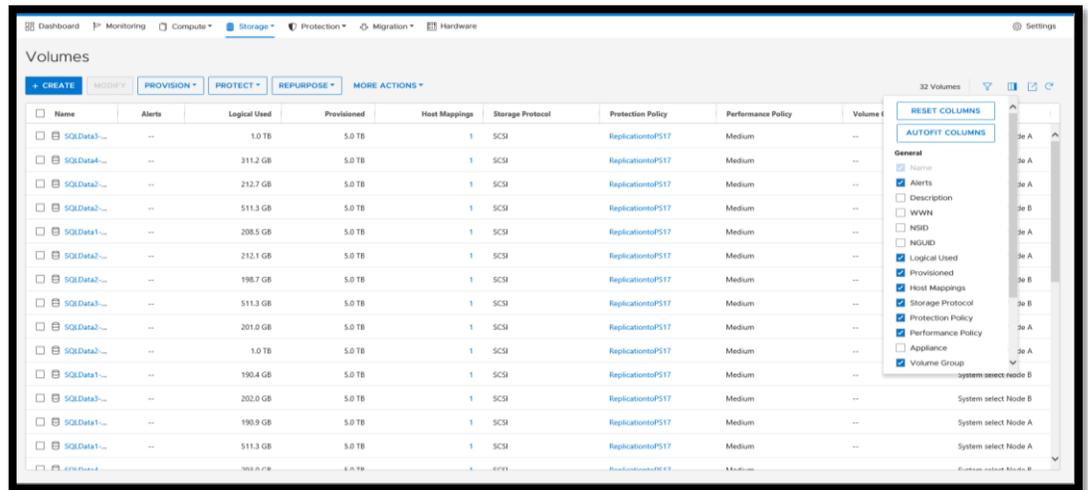


Figure 5. Selecting additional columns in the Volumes list

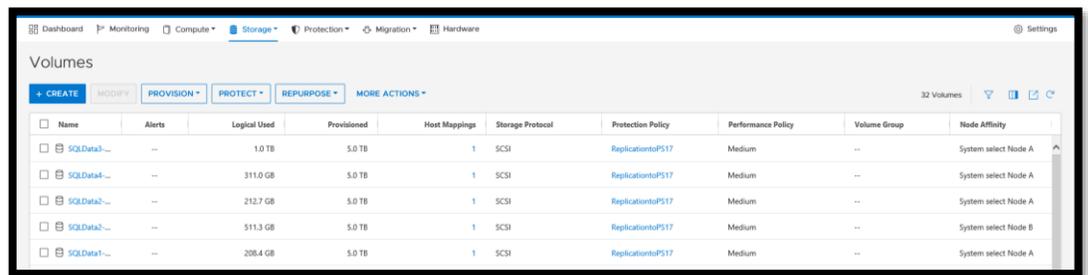


Figure 6. Volumes list with Node Affinity added

“System select node A” and “System select Node B” are the system assigned values. Node Affinity can be set using the PowerStore CLI (CLI) or PowerStore REST API. This example uses the CLI.

The following command will display the details of a volume named “SQLData2-007”:

```
pstcli -destination <PowerStore Manager URL/IP> -u <id> -p <password> /volume -name "SQLData2-007" show
```

To make changes, the volume id is required. You can either find the volume ID from the above command or append “ -select id” to return only the volume id:

```
pstcli -destination <PowerStore Manager URL/IP> -u <id> -p <password> /volume -name "SQLData2-007" show -select id
```

Note: The “id” used here is not displayed in the PowerStore Manager and therefore needs to be discovered.

After the volume ID is discovered, it can be used in the following command to set the Node Affinity. The following command sets the Node Affinity for volume **a5edff43-e9a7-**

48bd-9b19-e46c7d87d7d2 to Preferred_Node_A. Valid values are **Preferred_Node_A**, **Preferred_Node_B**, and **System_Select_At_Attach**.

```
pstcli - destination <PowerStore Manager URL/IP> -u <id> -p
<password> /volume a5edff43-e9a7-48bd-9b19-e46c7d87d7d2 set -
node_affinity "Preferred_Node_A"
```

Note: When Node Affinity is set to **Preferred_Node_A** or **Preferred_Node_B**, it is no longer managed by PowerStore unless it is set to **System_Select_At_Attach**.

Virtualization

When running SQL Server inside virtual machines (VMs), it is recommended to start with the VMware document [SQL Server on VMware Best Practices Guide](#) for performance and scalability. This guide provides guidance for configuring SQL Server, the operating system, and virtual machines for best performance on VMware.

For more tuning recommendations and best practices regarding VMware vSphere, VMware ESX, and PowerStore T or PowerStore X (AppsOn), see the documents [Dell PowerStore: Virtualization Integration](#) and *Dell PowerStore Virtualization Infrastructure Guide* at Dell.com/powerstoredocs.

In addition to these general best practices for virtualization, there are some VMware features and integrations for SQL Server.

VMware vVols vs VMDKs

Snapshots are a great way to protect SQL Server virtual machines and are encouraged whenever possible to provide an additional recovery point to SQL Server backups. However, taking snapshots on VMDK disks can result in a decrease in performance and sometimes lengthy maintenance procedures to maintain snapshot copies, coalesce, and delete.

vVol snapshots are offloaded to PowerStore, using hardware snapshot capabilities. This process removes the traditional issues of VMDK snapshot performance and maintenance. There is no performance penalty for vVols that contain snapshots. In addition, there are no lengthy maintenance procedures.

Also, Dell Technologies has tested the performance of VMware vVols and VMDKs. In our SQL Server workload testing, the performance has been proven to be equal on volumes without snapshots. On volumes that contain snapshots, the performance on vVols is superior. Therefore, vVols are recommended when running SQL Server workloads to provide superior snapshot performance.

A complete discussion of vVols vs VMDK disks is in the VMware article [Understanding Virtual Volumes](#).

Metro Volume

PowerStore Metro Volume allows synchronous replication of VMFS datastores in a vSphere Metro Storage Cluster (vMSC). This enables vSphere stretched cluster capabilities within or across sites. It can also be used to accelerate VM migrations. One of the challenges with migrating SQL Server virtual machines can be the time it takes to

perform a storage migration due to data size. Metro Volume can be used to accelerate migrations between sites, in addition to high availability and fault tolerance. For complete details on PowerStore metro volume, see the document [Dell PowerStore: Metro Volume](#).

AppsON

When running workloads on virtual machines on the PowerStore X appliance or AppsON mode, there are some important considerations when sizing workloads.

General VMware tuning and resource allocation principles apply; consult the VMware article [About vSphere Resource Management](#) for general tuning guidelines. If resources in the vSphere cluster become a bottleneck, guest storage performance will suffer, reducing the effectiveness of running workloads in AppsON mode.

Resource reservation

PowerStore X appliances reserve approximately half the CPU and memory resources for running PowerStore appliance tasks. When sizing workloads to run on the appliance, this needs to be considered. This reservation is visible in vSphere as the CPU and memory will be displayed as consumed by the PowerStore X controller virtual machines. Also the ESXi cluster cannot run at 100% capacity without incurring degraded performance. A minimal amount of resources are required for ESXi overhead. See the VMware articles [Configuring Resource Allocation Settings](#) and [Administering Memory Resources](#) for resource reservation and sizing guidance.

For production applications where guaranteed performance is necessary, make sure to size for possible node failover in the PowerStore cluster.

SQL Server virtual machines

General best practices apply when running SQL Server virtual machines in AppsOn mode, but there are two items to note. First, VMware vVols will typically be used as storage for VMDKs. While other storage options are still available, PowerStore X is optimized for vVols storage, so this option should be the first choice. Also, use the VMware Paravirtualized SCSI (PVSCSI) Controller as the virtual SCSI Controller. The PVSCSI controller improves performance and reduces CPU overhead.

For more best practices guidance for running SQL Server virtual machines, see the VMware guide [Architecting Microsoft SQL Server on VMware vSphere](#).

Multi-appliance virtual machine migrations

Starting with PowerStoreOS 2.0, up to four PowerStore X appliances can be added to a PowerStore X cluster. Then, AppsOn workloads can be moved among appliances in the PowerStore X cluster using VMware vSphere vMotion. When moving virtual machines among appliances in an AppsOn cluster, the storage will not migrate since PowerStore Storage Containers are presented as a vVols datastore across all appliances in the cluster. However, the vVols do reside on a single appliance. Therefore, this must be considered when migrating VMs between clustered appliances. For a VM, running the compute on one appliance when the vVols reside on another appliance will result in degraded performance.

In the Compute-> Virtual Machines view below is a VM running on **PS-9-appliance-1-node-B**.

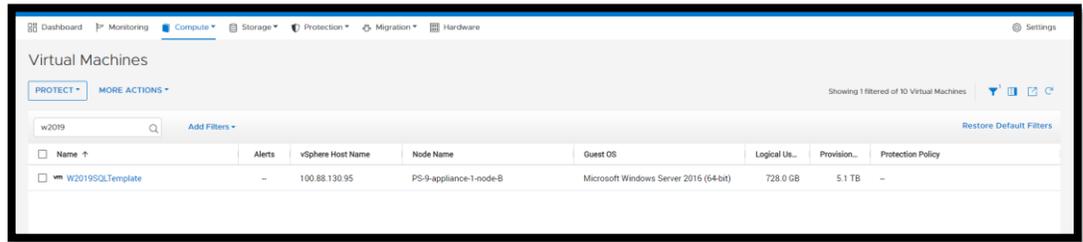


Figure 7. Virtual Machine residence

Examining the vVols that belong to this VM shows that they are on PS-9-appliance-2. These vVols should be migrated to PS-9-appliance-1. This process is done by selecting a volume and then selecting **Migrate**. This action can also be done in the PSTCLI using the `migration_session create` command.

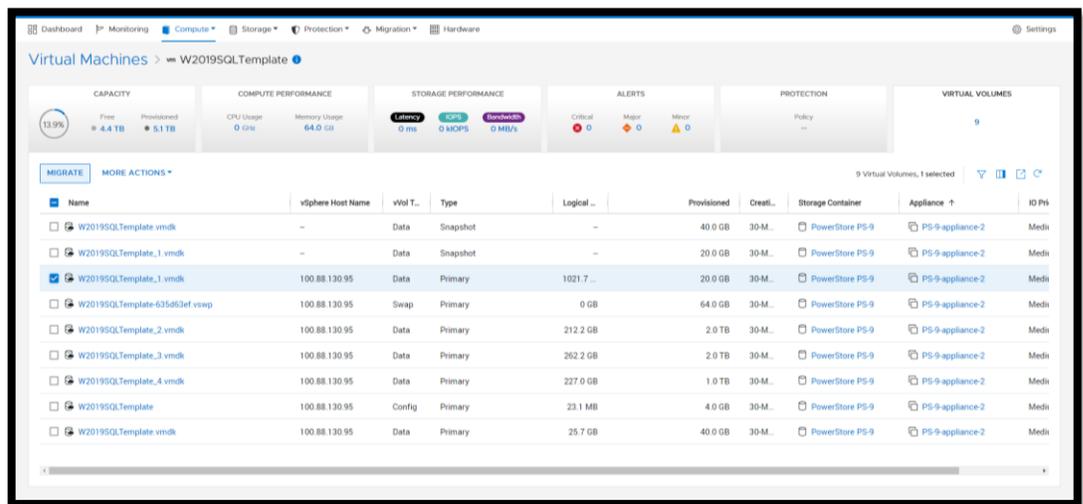


Figure 8. vVol residence

Once the migration has been started, it can be viewed in the Migration section.

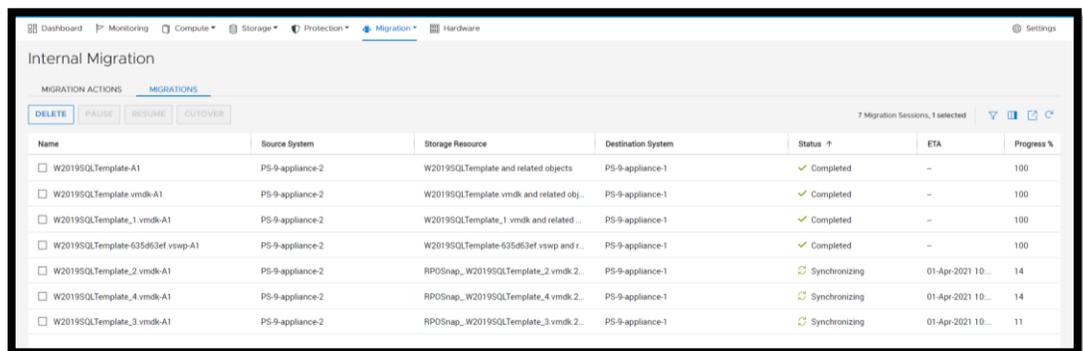


Figure 9. vVol migration

SQL Server Big Data Clusters

The Kubernetes platform performs resource balancing and scheduling within the Big Data Cluster and is not aware of the virtualization layer. When optimizing for performance, it is recommended to balance the workload across all nodes in the AppsON environment. In a simple scenario, there would be an even number of worker nodes in the Big Data Cluster. More complex scenarios may require placing certain nodes in the AppsON environment and using Kubernetes scheduling features to place targeted workloads on AppsON nodes. These concepts are covered in the Kubernetes documentation [Assign Pods to Nodes](#).

Kubernetes nodes in the master role consume relatively few resources. However, they need to be responsive and highly available. Master node VMs could be run on other available nodes in the vSphere cluster when applicable to reserve resources for the AppsON environment.

Regardless of where master role VMs run, based on internal testing, it is recommended to reserve CPU and memory resources for these K8s masters. If resource contention is detected within K8s masters and they cannot immediately get the resources they require, it can panic the K8s resource monitor and workload pods may be restarted.

Reducing SQL Server I/O

Optimizing SQL Server workloads will improve the performance and scale of the storage layer and ultimately the related applications. The following sections include common techniques to reduce SQL Server storage I/O.

Memory

Allocating the proper amount of memory to SQL Server can increase performance and avoid unnecessary I/O. SQL Server performs all I/O through the buffer pool (cache) and uses a large portion of its memory allocation for the buffer pool. Ideally, when SQL Server performs I/O, the data is already in the buffer pool and it does not need to go to disk. This type of I/O is referred to as logical I/O and is the most desirable because it results in the best performance. If the SQL Server data does not need to reside in the buffer pool, it needs to access disk resulting in physical I/O. Proper memory allocation is critical to SQL Server performance and can improve storage performance as well. Often, SQL Server and storage performance can be further improved by adding additional memory to the SQL Server instance. Generally, allocating more memory is better, but there is a point of diminishing returns that is unique to each environment.

Buffer pool extension

Starting with SQL Server 2014, the buffer pool can be extended to a file on the file system to provide additional space to cache data or index pages. Using this feature can provide significant performance benefits without adding memory to the database server in some cases. With more pages cached on the server, the I/O load on the array is reduced. When placing the buffer pool extension on the array, create a separate volume for the buffer pool extension and do not configure snapshots for the buffer pool extension volume. The buffer pool data is repopulated by SQL Server when the instance is restarted. Therefore, data recovery does not apply.

Persistent memory

SQL Server 2016 introduced support for persistent memory (PMEM). The capabilities of PMEM are being expanded with SQL Server 2019 to cover more scenarios and support the Linux operating system. Often, PMEM can be used to accelerate challenging I/O workloads and make I/O patterns more efficient with SQL Server PMEM features such as tail-of-log-cache and Hybrid Buffer Pool. Virtualized environments running Hyper-V or VMware can also take advantage of PMEM making it a wise investment. Dell servers support PMEM starting with Dell PowerEdge 14th-generation servers. For more information, see [Storage Class Memory in SQL Server 2016 SP1](#) and the [Dell NVDIMM-N Persistent Memory User Guide](#).

Instant File Initialization

By default, SQL Server writes zeros to the data file during the allocation process. The process of zeroing out the data files consumes I/O and acquires locks as the SQL Server data pages are written. This activity can occur for minutes or even hours depending on the file size. While this result may seem minor, writing zeros to these files can occur at disruptive periods when time and performance are critical. Example scenarios include database auto growth, expanding a full data file, replication, or restoring a database as part of a disaster-recovery event. When Instant File Initialization is enabled, SQL Server will skip the process of zeroing out its data files when allocating space. Dell Technologies recommends enabling Instant File Initialization.

Resource Governor

The Resource Governor was added in SQL Server 2008 to allow database administrators to limit the CPU and memory resources a SQL query can consume. This feature was enhanced in SQL Server 2014 to allow I/O resources to be limited as well. For example, the Resource Governor can be used to reduce the impact of a user running an I/O intensive report by limiting the maximum number of IOPS that user can perform. While a query throttled by the Resource Governor takes more time to complete, overall database performance is better.

Database design considerations

Reducing SQL Server I/O requires a holistic approach. Many of the items in this section require involvement from the entire team responsible for SQL Server applications. This team could include the business owner, architect, developer, database administrator, and system administrator. Decisions at the design level have a multiplied downstream impact as data is written and read multiple times and duplicated in various types of database copies. These include databases that are copied for other uses such as testing and reporting, replicated databases, replicated storage, and backups. One of the most challenging aspects of SQL Server is that the I/O pattern and the amount of I/O that is generated can vary greatly depending on the application, even if those applications have databases of the same size. This is because the design of both the database and the data access code control SQL Server I/O.

Database tuning can be one of the most cost-effective ways to reduce I/O and improve scalability. At a high level, consider the following areas when tuning a database to reduce I/O.

Database design

The foundation of the entire database and the schema for how data will be stored and ultimately accessed is determined by the database design. The database design should support both usability and efficient data access. This includes efficient table design and data types as well as indexes, partitioning, and other features that can improve efficiency. It is common for database design to only be focused on usability while performance and scale are overlooked.

Query design

How a query is written can greatly affect the amount of I/O SQL Server must perform when running the query. Queries should return only the required amount of data in the most efficient manner possible. Tune the queries responsible for consuming a relatively large number of resources as well as possible for best performance and scale.

Application design

Consider how applications are using the data and the way it is requested. Sometimes code and component reuse can result in the same data being unnecessarily retrieved repeatedly. All data access should be purposeful.

Maintenance

SQL Server uses a cost-based optimizer to generate query plans for data access. These plans are based on the statistics regarding how data is distributed in the tables. If the statistics are inaccurate, bad query plans may result and unnecessary I/O will be performed. Proper database maintenance includes ensuring that statistics are up to date. Frequent data modifications can also lead to fragmentation within SQL Server data files, producing unnecessary I/O. Fragmentation can be addressed through index reorganization or rebuilds as part of regular database maintenance. The database maintenance process itself can also have a large I/O impact. Typically, every table and index does not need to be rebuilt or reorganized every time maintenance is run. In addition, table partitioning strategies can also be leveraged to make the maintenance process more selective. Consider implementing intelligent maintenance scripts that utilize usage statistics to perform maintenance on an as-needed basis. For mission-critical databases, maintenance activities need to be considered as part of the overall design. If maintenance is not considered as part of the overall process, issues can arise, such as unmanageable sizes and feature incompatibilities that limit available options and strategies.

Management and monitoring

Introduction PowerStore offers multiple features for managing and monitoring storage.

Volume mapping and clustering After creating a volume, host mapping presents storage from the PowerStore array to a host or a cluster of hosts that are defined in a host group. Follow general best practices as outlined in the [PowerStore Host Configuration Guide](#) when mapping storage volumes.

Clustered SQL Server instances require special consideration since correct host mapping is critical to cluster failover. SQL Server Failover Cluster Instances (FCI) rely on Windows Server Failover Clustering (WSFC). Dell Technologies recommends using host groups to uniformly present storage to all the initiators for reduced management complexity. This allows a volume or set of volumes to be mapped to multiple hosts at a time and maintain the same Logical Unit Number (LUN) across all hosts. If volumes in the failover cluster do not have the same LUN number across Windows hosts in the cluster, failover does not work properly. Once the cluster configuration is complete, test the cluster failover to ensure that all failover components, including storage, are functioning as intended.

Creating volumes

PowerStore is virtualized to take advantage of all the drives in the appliance, and in addition there are many types of files that are part of a SQL Server instance. These types of data often have different performance and snapshot requirements. For performance-sensitive applications, Dell Technologies recommends creating at least five volumes for an instance of SQL Server as shown in the following table.

Table 3. Volume provisioning recommendations

File type	Number of volumes per instance	Typical performance requirements	Typical snapshot requirements
User DB data	1+	Lower performance may be acceptable	Frequent snapshots, same volume group as log volumes
User DB transaction log	1+	High performance required	Frequent snapshots, same volume group as data volumes
Data root directory (includes system DBs)	1	Lower performance may be acceptable	Infrequent snapshots, independent schedule
Tempdb data and transaction log	1	High performance required	No snapshots
Native SQL Server backup	1	Lower performance may be acceptable	Snapshots optional, independent schedule
Memory-optimized filegroup (if used)	1+	High performance required	Frequent snapshots, same schedule as log volume

Associate databases that span multiple LUNs within a volume group. This helps ensure that other features such as snapshots, replication, and thin clones will be configured properly and maintain data consistency.

There is a balance between manageability and performance that is unique to each environment. Generally, fewer volumes are easier to manage while more volumes enable better performance and optimize for advanced features such as storage replication, filegroup backups, and piecemeal restores.

Performance critical volumes can be added to a Watchlist for enhanced visibility. To do this, select the volumes and under the More Actions menu, select Add to Watchlist.

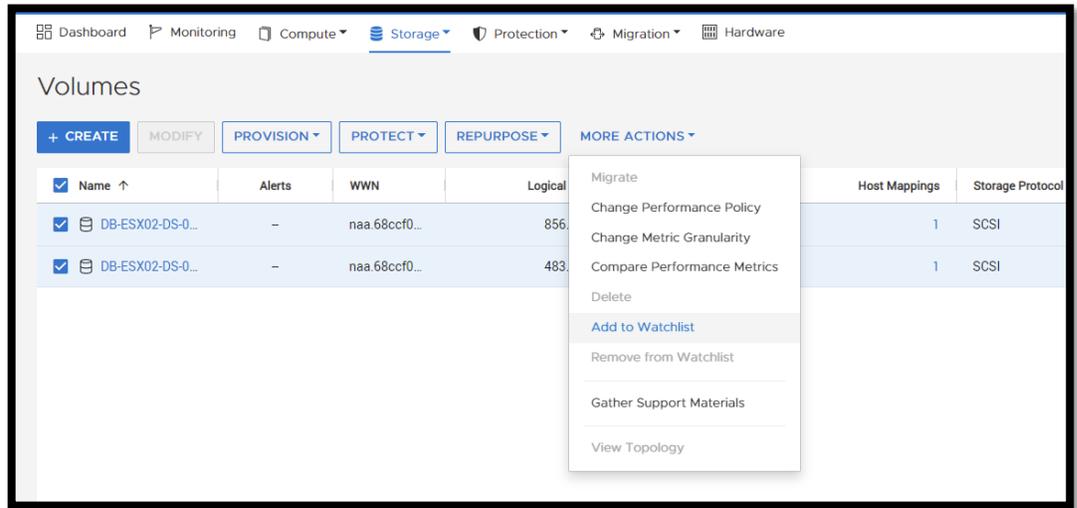


Figure 10. Adding volumes to the Watchlist

Once the volumes are added to the Watchlist, they will appear on the main PowerStore Dashboard. This provides some basic performance data as well as the ability to navigate directly to them by clicking the volume name.

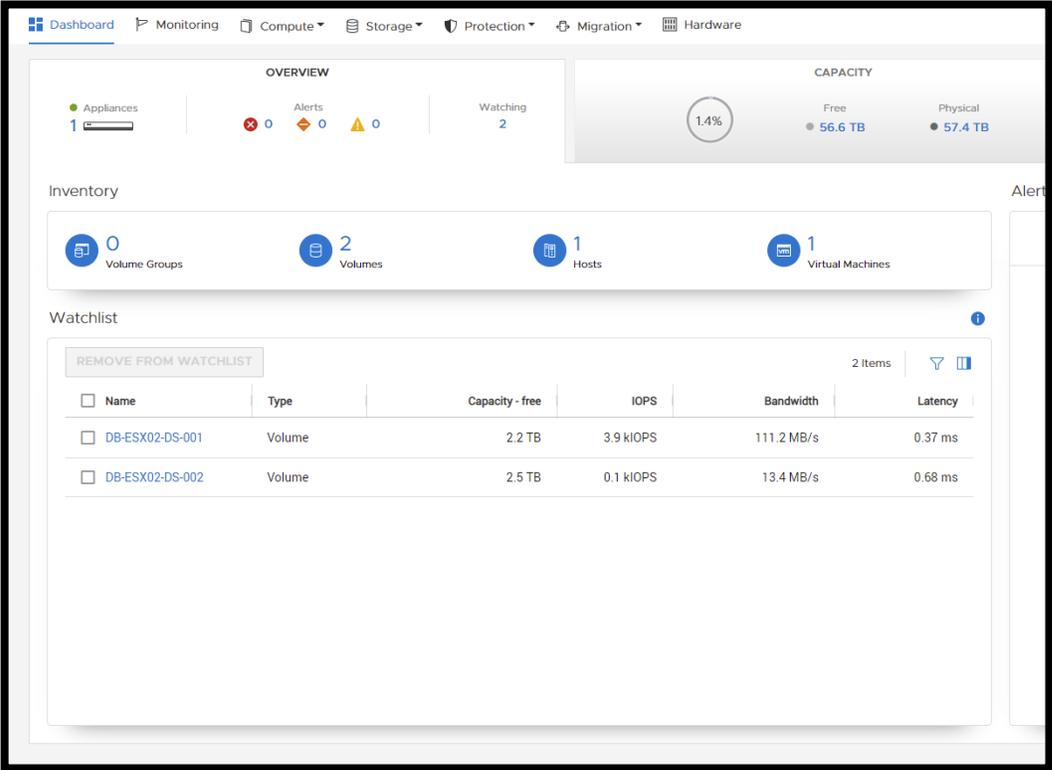


Figure 11. Adding volumes to watchlist

You can also add hosts to the Watchlist in the same manner. If you add a host that already has volumes on the list, those volumes will be combined under the host.

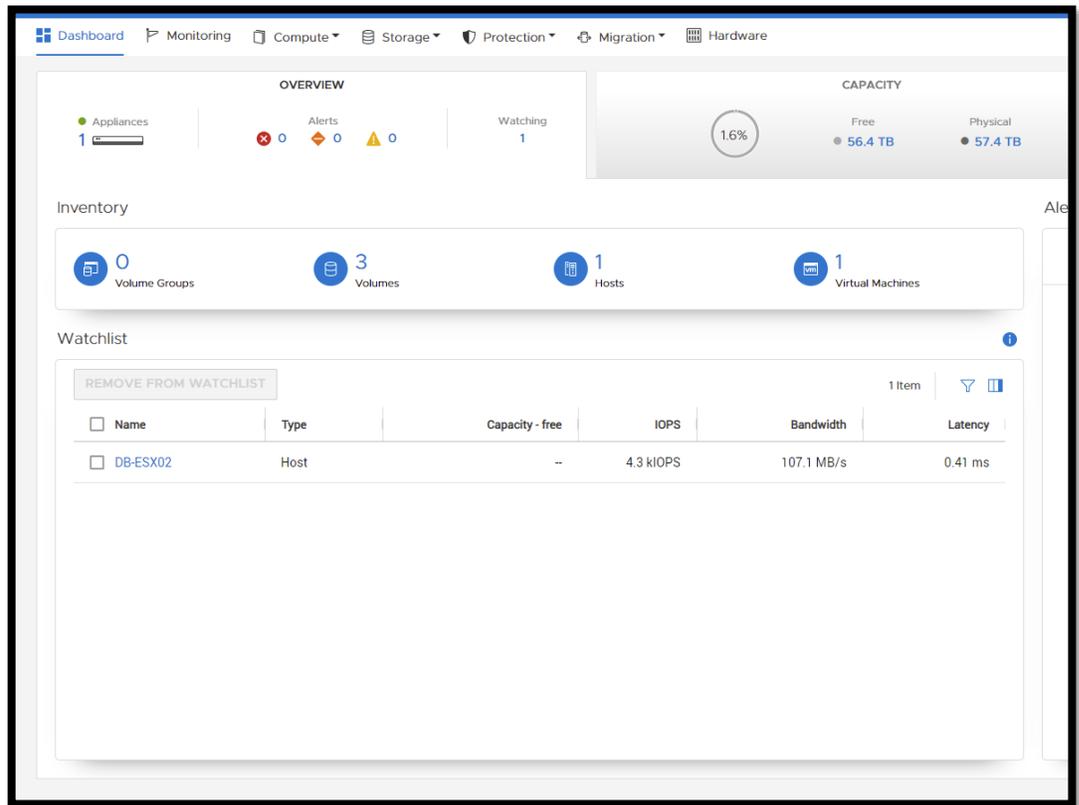


Figure 12. Adding hosts to watchlist

Allocation unit size

When formatting volumes, choose a 64 KB allocation unit size. This aligns with the 64 KB database extent size, which is the allocation unit that is used by SQL Server. The allocation unit size should be the same when formatting volumes on Windows or Linux platforms.

Disk-format wait time

With trim/unmap enabled (default setting in Windows Server), significant wait time occurs when formatting a PowerStore volume that is more than 1 TB. The larger the volume is, the longer the format wait time is, which is a common situation with external storage. This case applies to volumes formatted as NTFS or ReFS.

To avoid a long format wait time when mapping and formatting a large volume, temporarily disable trim/unmap. This setting is disabled using the `fsutil` command from a command prompt with administrator privileges.

Figure 13 shows the commands to query the state and to disable trim/unmap for NTFS and ReFS volumes on a host. A **DisableDeleteNotify** value of **1** means that trim/unmap is disabled, and long format wait times are avoided when performing a quick format.

Changing the state of `DisableDeleteNotify` does not require a host reboot to take effect.

```

Administrator: Cmd Prompt
Microsoft Windows [Version 10.0.17763.914]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Windows\system32>fsutil behavior query DisableDeleteNotify
NTFS DisableDeleteNotify = 0 (Disabled)
ReFS DisableDeleteNotify = 0 (Disabled)

C:\Windows\system32>fsutil behavior set DisableDeleteNotify NTFS 1
NTFS DisableDeleteNotify = 1 (Enabled)

C:\Windows\system32>fsutil behavior set DisableDeleteNotify ReFS 1
ReFS DisableDeleteNotify = 1 (Enabled)

C:\Windows\system32>fsutil behavior query DisableDeleteNotify
NTFS DisableDeleteNotify = 1 (Enabled)
ReFS DisableDeleteNotify = 1 (Enabled)

C:\Windows\system32>

```

Figure 13. Use the fsutil command to query or change the state of trim/unmap

Once the volume is formatted, reenable trim/unmap so the host can take advantage of deleted space reclamation for NTFS volumes.

For more background information about trim/unmap with PowerStore, see the document [Dell PowerStore: Microsoft Hyper-V Best Practices](#).

Thin clones

Thin clone uses a pointer-based technology to create a read/write copy of a volume, volume group, or file system. For data in a snapshot to be accessed, a thin clone must be created. Because the thin clone volume shares data blocks with the parent, the physical used space of the child volume is just the delta changes from after it was created. Thin clones are advantageous in a SQL Server environment because a SQL Server database can be duplicated for testing purposes, all while consuming less storage. For example, if a SQL Server admin must clone a multi-terabyte database server for a developer to run tests, the database can be isolated, tested, and only consume blocks that changed.

Within the PowerStore architecture, thin clones have several advantages for storage administrators.

- The thin clone can have a different data protection policy from the parent volume.
- The parent volume can be deleted, and the thin clones become their own resource.
- Clone databases for testing monthly patches or development.
- It is treated as a regular storage resource and therefore supports many data services such as protection policies, snapshot rules, or replication.
- Can be refreshed quickly to bring latest production changes to lower environment.
- Can be restored quickly to a baseline after testing.

When cloning databases for SQL Server, ensure all data and log volumes for the database are on the same volume or are part of the same volume group. See best practices for creating snapshots of SQL Server databases [below](#). For more information, see the document [Dell PowerStore: Snapshots and Thin Clones](#).

Data encryption

Many SQL Server applications have data encryption requirements, specifically on data at rest. Data at Rest Encryption (D@RE) can be used as an encryption solution for SQL Server without requiring any database or application changes. This also avoids any potential performance impact to the database server or 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.

Data reduction

PowerStore includes zero detection, compression, and deduplication integrated in the core product as part of the data reduction feature. While the amount of reduction varies depending on the dataset, the overall savings are similar to or better than SQL Server database compression. Since data reduction is always enabled on the array, SQL Server database compression and backup compression can typically be disabled, saving CPU resources on SQL Server.

When running SQL Server compute workloads on PowerStore X in AppsON mode, it is recommended to disable SQL Server compression to allow CPU resources to be allocated to other activities on the appliance.

For more information about the PowerStore data reduction benefits, see the document *PowerStore Data Efficiencies* on [Dell.com/StorageResources](#).

Scripting and automation

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

REST API

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

PowerStore CLI

The PowerStore CLI or 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. For more information, see the document *Dell PowerStore Command Line Reference* at [Dell.com/powerstoredocs](#).

Ansible

Ansible is a popular tool for automating IT infrastructure. Since Ansible is widely supported by various hardware and software platforms, it makes it ideal for deployment

automation and IaC scenarios. The Ansible plug in for PowerStore as well as complete documentation can be found at [GitHub.com/Dell/Ansible-PowerStore](https://github.com/Dell/Ansible-PowerStore).

Data protection and disaster recovery

Introduction

PowerStore has integrated snapshot and replication capabilities to protect data, and is all policy driven for ease of administration.

Snapshots and recoveries

To automate and simplify protecting data, PowerStore uses protection policies. These protection policies help to protect data, set retention policies, and help guarantee recovery point objectives for an organization. Protection policies are a set of snapshot and replication rules that are applied to a volume or group of volumes.

In addition, protection policies can be applied to volume groups. By applying a protection policy to a volume group, it allows snapshots, replication, and recovery of the entire group. This allows the protection of a database that spans multiple volumes. To ensure successful recovery of SQL Server databases, ensure all database files, including the transaction log file, are in the same volume group and the write-order consistency option is enabled on the volume group.

Snapshots and application backup and restore

Using array-based snapshots is an effective way to protect virtual machine data and establish a recovery point objective (RPO). In the PowerStore architecture, the snapshot schedule is created using protection policies. Each protection policy can define snapshot rules to establish a schedule and retention, and replication rules to specify a destination array and RPO.

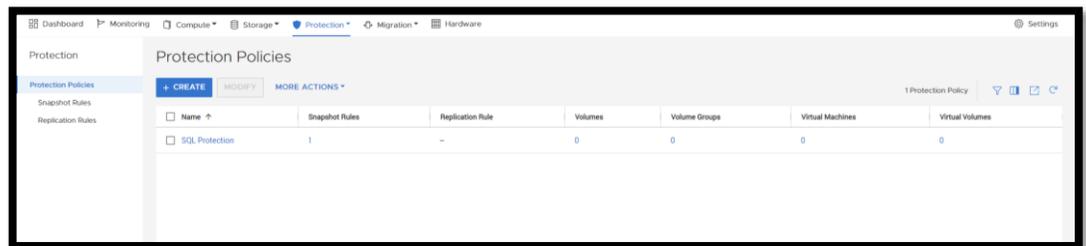


Figure 14. Protection policies screen in PowerStore Manager

PowerStore has three data recovery mechanisms that all behave differently depending on the usage scenario.

- **Thin Clone:** This takes an existing snapshot from a parent volume or volume group and creates a child volume or volume group from that point in time.
- **Refresh:** Using the refresh operation, snapshot data can replace existing data in the volume group. The existing data is removed, and snapshot data from the new source is copied to it in-place. A parent volume or volume group can refresh a child, and a child can refresh a parent.

- **Restore:** The restore operation replaces the contents of a parent storage resource with data from an associated snapshot. Restoring resets the data in the parent storage resource to the point in time the snapshot was taken.

Caution: Use of the refresh and restore operations on active database volumes may cause unexpected results and behaviors. All host access to the volume must be ceased before attempting these operations by either shutting down the SQL Server instance or taking the SQL Server databases offline.

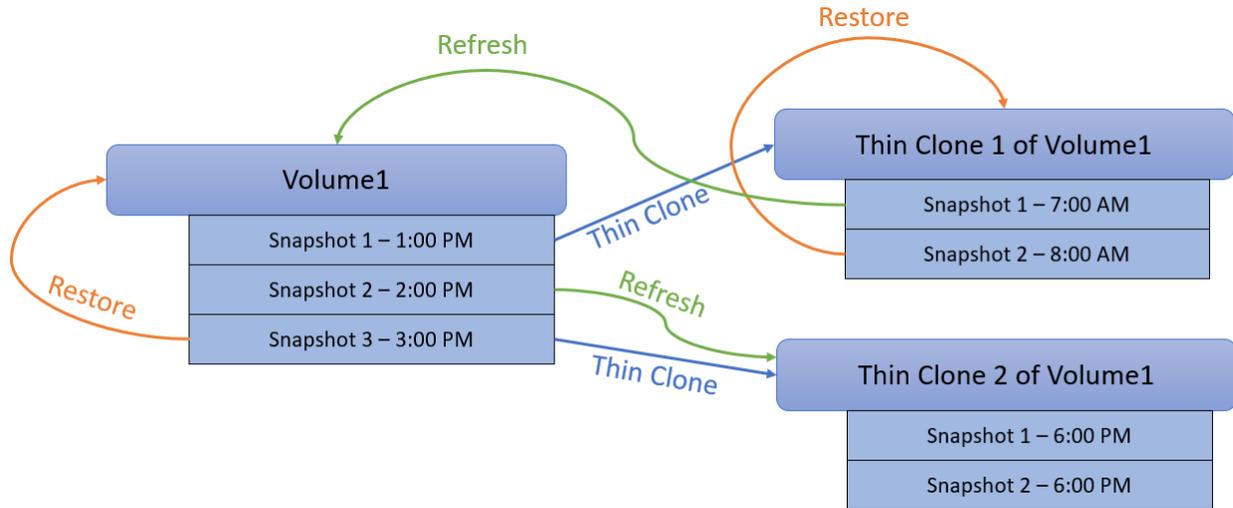


Figure 15. PowerStore snapshot and recovery anatomy

In addition to using snapshots to protect databases, if SQL Server hosts are leveraging boot-from-SAN, the boot volume can be added to the same volume group as the database volumes and the entire host can be protected.

Crash consistency and application consistency

When taking array-based snapshots of virtual machines, it is important to remember that snapshots that are taken without application coordination are considered crash consistent. Crash consistency is the storage term for data that has a snapshot taken in-flight without application awareness. While most modern applications can recover from crash-consistent data, their recovery can yield varying levels of success. For example, when recovering a Microsoft Windows virtual machine, as the operating system boots, it behaves like it experienced unexpected power loss, and potentially invoke check disk (chkdsk) on startup.

Dell AppSync enables taking application-consistent snapshots. When using products such as AppSync, coordination between the array and the application can help to assure that the data is quiesced, the caches are flushed, and the data is preserved in a known-good state. Application consistent snapshots offer a higher probability of recovery success.

Replication and remote recovery

PowerStore offers asynchronous replication to transfer data to another PowerStore array. When the secondary array is at a different location, this can help to protect data from localized geographical disasters. Replication RPOs and options are set within protection policies.

**Integrated copy
data
management**

AppSync is software that enables integrated Copy Data Management (iCDM) with the Dell primary storage systems, including PowerStore. AppSync, integrated with PowerStore snapshots, simplifies and automates the process of generating, consuming, and managing copies of production data. This solution addresses copy management use cases in a consolidated SQL database environment for database recovery and database repurposing. AppSync automatically discovers application databases, learns the database structure, and maps it through the virtualization layer to the underlying PowerStore storage. Orchestration of all the activities required from copy creation and validation through mounting the snapshots at the target host and launching or recovering the application. This solution supports and simplifies Microsoft SQL Server workflows that include refreshing, expiring, and restoring the production database. Additional information about AppSync can be found in the documents [AppSync Integration with Microsoft SQL Server](#) and [AppSync Performance and Scalability Guidelines](#).

References

Related documentation

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

The following links provide additional related resources:

- [PowerStore Info Hub](#)
- Storage Class Memory in SQL Server 2016 SP1
- Dell NVDIMM-N Persistent Memory User Guide
- Dell SQL Server Solutions
- Microsoft SQL Server
- [GitHub.com/Dell/Ansible-PowerStore](https://github.com/Dell/Ansible-PowerStore)
- AppSync Integration with Microsoft SQL Server
- AppSync Performance and Scalability Guidelines
- VMware Documentation
- SQL Server on VMware Best Practices Guide