Dell EMC PowerStore – SAN Storage and NAS System

Dell EMC announced the midrange unified storage system called PowerStore in May of 2020. In 2021, Dell EMC announced an additional PowerStore model targeted at entry level enterprise. The Dell EMC PowerStore system is derived from elements of earlier midrange storage platform and new functionality including the ability to run applications on the storage system. The operating environment is called PowerStoreOS and is deployed on purpose-built hardware in two different ways: one is native on the bare metal hardware and is designated by T as a suffix of the product model and the other installs in a VM with VMware ESXi running on the PowerStore hardware, designated with an X suffix. Multiple models of PowerStore vary by the amount of cache and the type of processor. All modes use solid state devices of flash or SCM Optane and devices in the base controller enclosure use NVMe protocol. All models can scale out to four pairs of nodes of the same T or X model type. T models offer unified block and file with NAS support.

The PowerStore 500 model was added by Dell EMC in 2021 as an entry level addition to the PowerStore family. Unlike the other models, PowerStore 500 is available as a T model only. PowerStore 500 can scale out in a cluster of up to 4 appliances, or mixed with other T models.

Many of the software functions of the Dell EMC PowerStore systems were extracted from Unity software and PowerStoreOS is a container-based operating environment. PowerStore systems use a SuSE Linux kernel and support Docker containers. The filesystem software in the T models support multiple NAS Servers, which were previously called virtual data movers. PowerStore Manager is the element manager, deployed as a web GUI, implemented with HTML5.

The dual controllers are packaged in a 2U enclosure along with up to 21 NVMe protocol 2.5” devices and 2 to 4 NVRAM modules depending on the model. Additional devices that are SAS attached can be added with up to three additional enclosures. Software features are included in the base price and include snapshots, asynchronous remote replication, VVOLs and VASA support. Telemetry data is collected and analyzed with a SaaS product called CloudIQ. Health checks are done and proactive maintenance is scheduled based on the analytics.

The SSDs in PowerStore have a lifetime warranty as long as a maintenance contract is in place. In addition, the maintenance price is guaranteed not to be increased from the initial purchase.

Evaluator Group Coverage
Related coverage on the Evaluator Series Research website includes:
SAN and NAS matrices
SAN and NAS Evaluation Guides
Dell EMC PowerStore Product Brief
Competitive Product Analysis and Product Briefs

Highlights

- Dell EMC Unified Midrange System (T models) and block storage only (X models)
- File – Updated 64 bit OS
- Container based – Linux OS
- All flash and SCM storage
- Support for RAID 5 with either 4+1 or 8+1 protection
- HTML5 embedded PowerStore Manager
- Encryption with use of self-encrypting devices and internal key manager
- Block storage support over iSCSI and Fibre Channel
- NVMeOF available with Fibre Channel and NVMe/TCP
- File support over CIFS/SMB, NFS, FTP/SFTP
- VMware vSphere API's for Array Integration (VAAI) for hardware array support of advanced features
- VASA 2.0 and VVOL support
- Dell EMC Support Assist – monitoring and CloudIQ link
- Software functions included in system
- Snapshots, remote replication, file and data migration
- Native MetroSync replication
- Metro Volumes
- X Models support running applications with VMware
    - vMotion support for applications and data
    - replication to public cloud
- Battery backup for graceful shutdown
- 2U Enclosure with 2 controllers
- 21 devices in controller enclosure

- 25 devices in added enclosures up to 3
- Lifetime flash guarantee while under support maintenance
- Guaranteed maintenance price – no increases
- QoS – Host I/O limits on LUNs, filesystems, snapshots
- Inline compression and deduplication in software for all flash storage pools using Intel QuickAssist
- Policy-based file tiering to clouds
- Anytime upgrade – with additional charge on top of premium support package

Contents

Overview

The Dell EMC PowerStore consists of six models with different cache sizes and processors. The T designated models have the PowerStoreOS installed on the the purpose-built hardware while the X models have VMware ESXi installed on the hardware and PowerStoreOS runs on a VM. Both series can scale out from a single controller pair up to four controller pairs, but must do so with models of the same T or X type. The T series supports unified block and file access. The X series can run applications and interacts with VMware Cloud Foundation and vRealize.

There are two deployment modes for PowerStore. **Unified mode** is the default and reserves memory and storage for file operations. 1.5 GB of space per filesystem is reserved. **Block Optimized mode** can be configured and will allow only block or vVol support (no file support). The deployment mode is set at installation and cannot be changed unless starting for new install state.

Figure 1: Dell EMC PowerStore Controller (source: Dell EMC)

Dell EMC PowerStore Model Comparison

Table 1: Dell EMC PowerStore

Model	500	1200	3200	5200	9200
Processor	Intel Cascade Lake 12-core / 2.2GHZ	Intel Cascade Lake 10-core / 2.4GHZ	Intel Cascade Lake 16-core / 2.1GHZ	Intel Cascade Lake 24-core / 2.2GHZ	Intel Cascade Lake 28-core / 2.2GHZ
Memory	192 GB	384 GB	768 GB	1,152 GB	2,560 GB
Max Drives (per appliance)	97	93	93	93	93
Max Capacity (RAW) per dual controller	1490 TB	1430 TB	1430 TB	1430 TB	1430 TB

Table 1: Dell EMC PowerStore Model Comparison

Product Architecture – PowerStore

The architecture of the PowerStore system from a hardware standpoint is a traditional dual controller storage system with devices in the controller enclosure attached with PCIe and using NVMe protocol (up to 21 devices. Up to three additional enclosures with 25 SSDs can be attached using SAS. Expansion can only be done if all 21 slots in the base units have devices and if SCM (Optane) devices are not used in the base unity.

These are scale-up systems with the ability to add more capacity. The systems can scale out with the addition of up to three controller pairs. Data on devices is not shared between the node pairs, access is directed from the hosts to the node pairs where data is resident.

- PowerStore Hardware - 2U contoller enclosure with two controllers and 21 devices connected with PCIe. These devices use NVMe protocol and may be flash or SCM Optane SSDs.
- SCM devices can be used as persistent storage or combined with flash to be used for metadata tiering.
- Optional expansion device enclosures with up to 25 SAS 2.5" devices. Only may be used if all 21 slots in base unit and no expansion if SCM used in base unit.
- An M.2 SATA flash device is used for booting. The primary boot device is 240 GB. A secondary 120 GB boot device is used as an alternate and for storing log data.

PowerStore Configuration

Model	500	1200	3200	5200	9200
LUNs	1,000	2,000	2,000	2,000	2,000
LUN Size	256 TB	256 TB	256 TB	256 TB	256 TB
VVOLs	19,000	19,000	19,000	19,000	19,000
Snapshots	10,000	10,000	10,000	10,000	10,000

Table 2: Unit Configuration Limits

NVRAM

The NVRAM is used for persistent storage for write data as a write cache, allowing for write complete to be returned to the host as quickly as possible. Write data is received in DRAM cache and copied to the NVRAM immediately. The NVRAM is a pair (mirrored) with a single pair (two devices) for Models 1200 and 3200 and two

pairs (four devices) for the other models. The NVRAM devices occupy device slots in the base unit. These slots cannot be used for other purposes and with the single pair in models 1200 and 30200, the two unused slots will have device blanks in them.

NVRAM devices have DRAM (for performance reasons), flash chips for persistence, and battery backup units to power the DRAM until data is copied to flash in the event of a power loss. NVRAM is connected via PCIe using NVMe protocol.

Figure 2: Base Unit with NVRAM source: Dell EMC

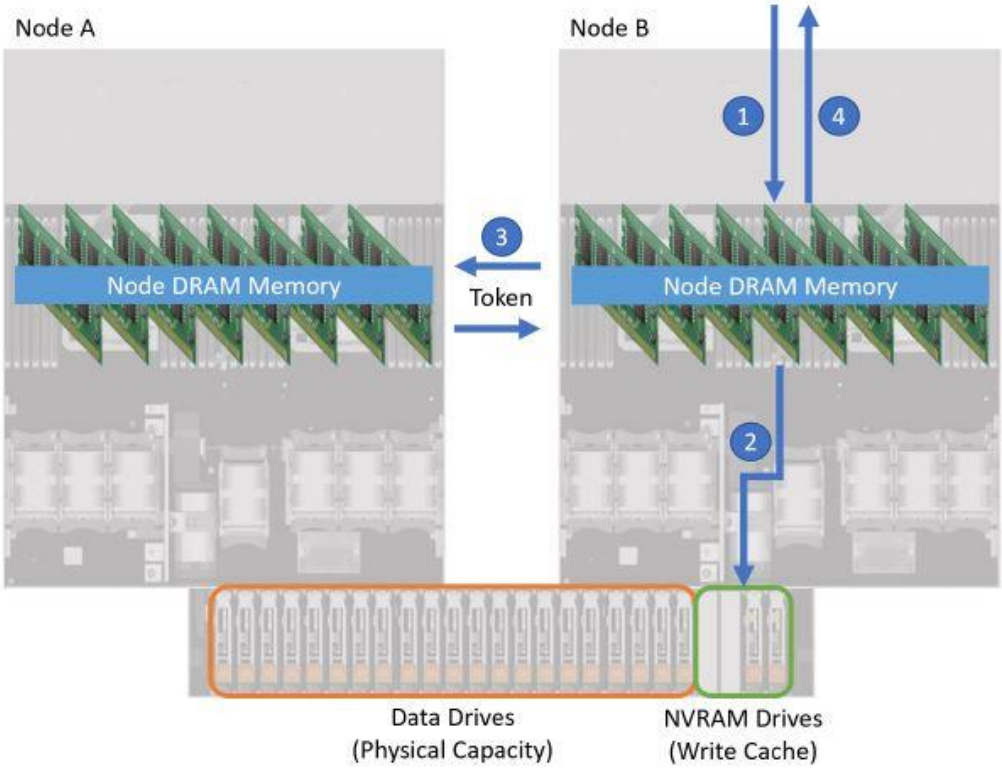

Figure 3: Write Data Flow with NVRAM source: Dell EMC

Software Architecture

PowerStore Operating Environment

The PowerStore operating environment uses an embedded SuSE Linux with containers for PowerStoreOS software and other functionality. Currently, the additional functionality is Dell EMC Support Assist running in a Docker container. The containerized PowerStoreOS can be installed directly on the bare metal for the PowerStore system (T models) or as a VM on ESXi running on the PowerStore hardware (X models).

PowerStoreOS optimizes usage of multiple processor cores and multi-threading for performance. The PowerStore systems supports non-disruptive upgrades for both firmware and hardware.

- 64 bit based operating environment for enhanced memory and I/O address space support
- Support for multi-core Intel processors
- Flash SSD and SCM Optane device support
- The ability to intermix iSCSI and FC networks attachments simultaneously
- Virtual provisioning support

Included Software Area	Function
Management	• PowerStore Manager - element manager • Proactive Assist – configure remote support, online service • Quality of Service – block and vVols • CloudIQ storage analytics • Thin provisioning and dynamic storage pools • Compression, deduplication, zero-detect
Protocols	• File – NFS, CIFS/SMB, FTP/SFTP (T models only) • Block – Fibre Channel (NVMe), NVMe/TCP, iSCSI, vVols
Local Protection	• Encryption – using self-encrypting drives • Snapshots and thin clones • AppSync Basic • Anti-virus linkage • Dell EMC Common Event Enabler, Event publishing agent • Quotas
Remote Protection	• Native Asynchronous – block and file • Native MetroSync
Migration	• Block migration from Dell EMC Unity, VNX, SC, PS

Table 3: Included Software for PowerStore

Optional Software Area	Function
Additional charge software	• Dell EMC SRM • AppSync Advanced • Data Protection Suite • Dell EMC RP4VM • PowerPath • VPLEX

Table 4: Optional Software for PowerStore

VMware VAAI Support

The PowerStore system supports the VMware vSphere APIs for Array Integration (VAAI), a set of mechanisms that allow processing for certain data-related services—copying data when creating a new VM, for example—to be offloaded from the ESX host to a storage array. The intent of these APIs is to streamline the functioning of the ESX server and speed-up delivery of storage-supported services.

Full copy — Enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data. A single SCSI Extended Copy command is used to copy a set of contiguous blocks.

Block zeroing — Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines. The SCSI Write Same command is used to copy a zero block for the number of blocks required.

Hardware-assisted locking — Provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX host farms sharing a datastore.

Dead Space Reclamation — Reclaims blocks from VMFS deleted files. The SCSI Unmap command is used to free physical space.

Out of Space Condition Support — A VM is paused if disk space is exhausted as reported by the storage system. This allows administrators to mitigate the situation rather than causing VM to fail.

VMware VASA Support

The PowerStore systems support the VMware vSphere APIs for Storage Awareness (VASA) version 2.0 and 3.0 with an internal provider. VASA can display the features of the physical storage devices, allowing vSphere administrators insight into storage capabilities. The PowerStore system has a plugin called a vendor providers, which is the integration 'glue' that sits between vCenter and the storage array. Vendor providers retrieve the storage capabilities from the array and pass these onto vCenter, which in turn can display these capabilities in the user interface. VASA may also provide information about storage health status, configuration information and capacity.

VMware VVOL Support

PowerStore supports VVOLs either through block mode VMDK or attached through the NFS interface. Volumes for VMware virtual machines can be defined within VMware without the need for storage administrator to configure LUNs. Instead, the storage administrator creates a Protocol Endpoint which is part of the physical storage fabric and treated like a LUN and is discovered during a scan by ESXi. VASA APIs will bind VVOLs to the Protocol Endpoint (PE). Within the Protocol Endpoint, a Storage Container is created which is a logical grouping of the VVOLs that will be created in VMware by the VMware administrators. A VASA provider for the storage system discovers the Storage Containers and reports to vCenter. There may be multiple Storage Containers per storage system.

VMware vRealize Operations (vRO) and Ansible Module

PowerStore has a plugin for vRealize Operation to automate or assist in operations such as provisioning. In addition, Ansible scripts are provided for automation purposes.

Container Support

PowerStore supports containers, up to 50, with CSI (container storage interface) software. CSI is a plugin used to provide persistent storage for containers.

Storage Pools (Dynamic Pools)

PowerStore uses storage pools to abstract the capacity of solid state devices. Capacity is allocated based on chunks, which represent fixed capacity sizes from the drives that comprise the storage pool. An administrator creates storage pools and PowerStore Manager will create RAID groups across the devices used in the pools. LUNs or filesystems are created from the chunks in a pool. Pools can be created across multiple media types but Dell EMC recommends creating pools from the same media type because of different capacity sizes.

Space is assigned to LUNs and filesystems in 256MB for thin provisioned LUNs. I/O is distributed across the LUN, which is called a Dynamic LUN. The same RAID protection must be used across the entire pool. LUNs may be reduced in size (shrunk) through Windows Server Manager and requires a VDS provider. LUNs and filesystems may be expanded if capacity is available.

Adding additional drives to a storage pool will rebalance data across all drives in the pool automatically. The addition of drives must be in the increment of the minimum number of drives for the RAID type of the pool. The system will create additional RAID group(s) in the pool when drives are added.

Virtual Provisioning (Thin Provisioning)

Thin provisioning allocates capacity as needed by adding chunks from the storage pool. The minimum capacity consumed for thin LUNs is 1.75 GB – .25 GB of usable capacity and 1.5 GB of overhead for each thinly provisioned LUN. Capacity is added in 256 MB increments. A chunk is 8K in the PowerStore thin provisioned volumes. Additionally, monitoring and threshold enforcement to prevent abuse of "runaway" consumption of storage are included.

Space reclamation returns freed capacity back to the storage pool. Reclamation can be done by shrinking a LUN or filesystem or with software support using an API or a SCSI UNMAP command.

Inline Data Reduction

Inline compression and deduplication for PowerStore works for all flash storage pools and is always on using Intel QuickAssist. Software is used to check for duplicate data and found using a hash of the data, it is not stored and the block table is updated. Data that is not duplicate, is compressed using software algorithm. Highly active data is written directly without the additional processing overhead for compression and deduplication. Data stored that is not compressed may be compressed and deduplicated as a background task. The system automatically selects data that is not highly active to be compressed and deduplicated within the storage group. Compression and deduplication can work independently of all other software features. A pretest of data is performed and if data is not compressible, it is left uncompressed. A compression estimator tool is available from Dell EMC. Writes of zero blocks are also detected and does not result in a write action, only a table update.

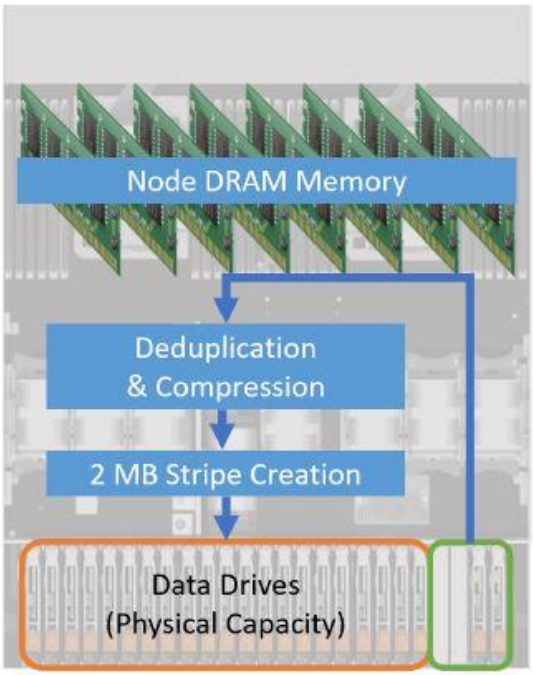

Figure 4: Data Reduction Data Flow source: Dell EMC

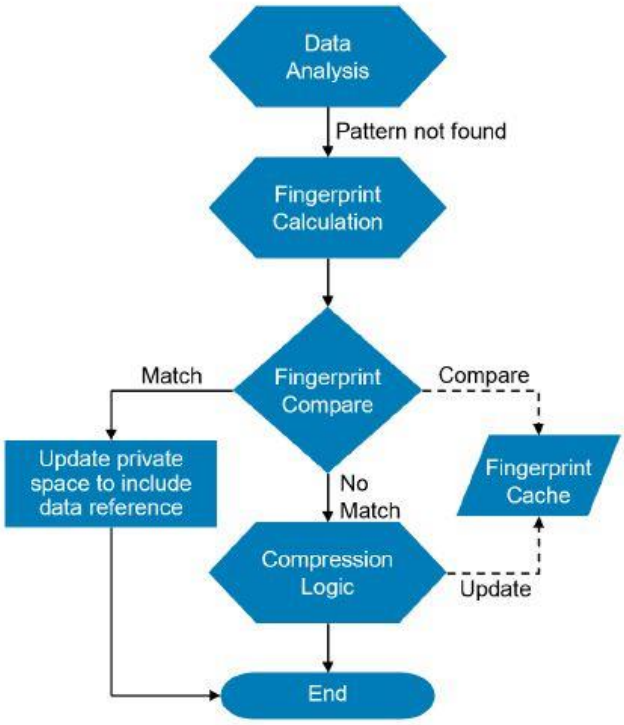

Figure 5: Data Reduction Logic source: Dell EMC

RAID Protection

Initially, RAID 5 protections was the only RAID type supported by PowerStore with RAID 5 groups of 4+1 or 8+1. In 2021, Dell EMC expanded on this by adding support for RAID 6 dual parity.

AppsON – Applications on Storage System

The X models of PowerStore have VMware ESXi installed on the PowerStore hardware with full VMware tools and support. The PowerStoreOS is run as a Virtual Machine within the VMware environment. With this organization, applications may be run on the the storage system. Called AppsON, applications can be hosted with the same capabilities as if running on a standard server with VMware. Features such as vMotion, Site Recovery Manager, vRealize, etc. are available in the AppsON environment. vMotion can be used to move data and to move applications to other VMware-based systems. VMware Cloud Foundation also provides connectivity to cloud environments.

For X mode, 50% of the CPU and memory is reserved for PowerStoreOS operation. The other 50% is available for user applications.

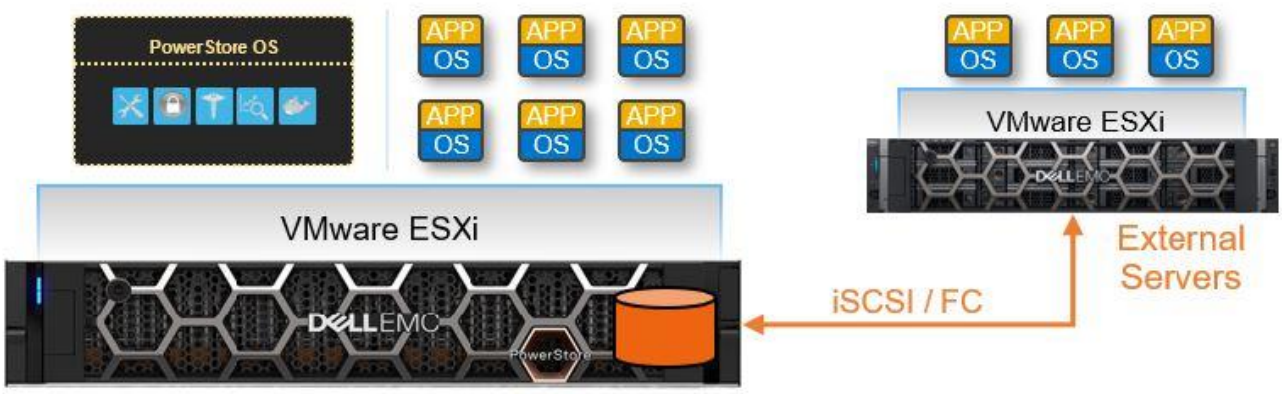

Figure 6: PowerStore X Overview source: Dell EMC

Advanced Features

The advanced features of PowerStore provide the distinguishing elements of the product. Dell EMC has used many of the software features previously available in Unity XT and are now included the base system. Many of the feature names have been changed and updated.

Included Software

PowerStore Manager

The element manager for PowerStore is PowerStore Manager which is an HTML5 implementation, differing from the Unisphere used previously. PowerStore Manager is targeted for the IT generalist with a simple control interface, dashboards for management and status, and wizards to accomplish common tasks. From the highly customizable user interface, simple, intuitive management for the systems can be performed.

Data Migration

PowerStore includes software to migrate volumes from earlier Dell EMC storage systems. Unity, VNX, SC, EqualLogic system data can be migrated to PowerStore and access paths automatically switched.

Encryption

Encryption of data at the array level is supported using Self Encrypting Drives. An external key manager conforming to KMIP is used. FIPS-140-2 level 1 certification is validated for devices except for NVRAM devices.

Snapshot (redirect on write)

The redirect on write solution is a space efficient copy that only requires space for the blocks that are modified. PowerStore snapshots are used for both LUNs and filesystems – a single implementation for the unified environment. Snapshot works with the storage pooling implementation and allocates 8K blocks as minimum increment for snapshots. A Snap Mount Point is created as the LUN or filesystems reference for the snapshot. This allows usage of the LUN or filesystem without requiring another SCSI scan.

All snapshots can be configured with automatic deletion policies to contain the proliferation of snapshots and unneeded consumption of storage. Snapshots can be controlled by PowerStore Manager, CLIs, or through the RESTful API.

Read/write access to the snapshot is allowed. The user is able to do a rollback to restore a snapshot copy to the original volume. Cascading snapshots or a snap of a snap can be performed. Consistency groups are supported for coordinated snapshots.

Native Asynchronous Replication – Block (LUN)

Asynchronous replication for PowerStore is called Native asynchronous Replication. The replication feature is native to the data services and works for LUNs (block mode). Asynchronous replication is done over Ethernet IP connections only. Any Ethernet port may be used for asynchronous replication.

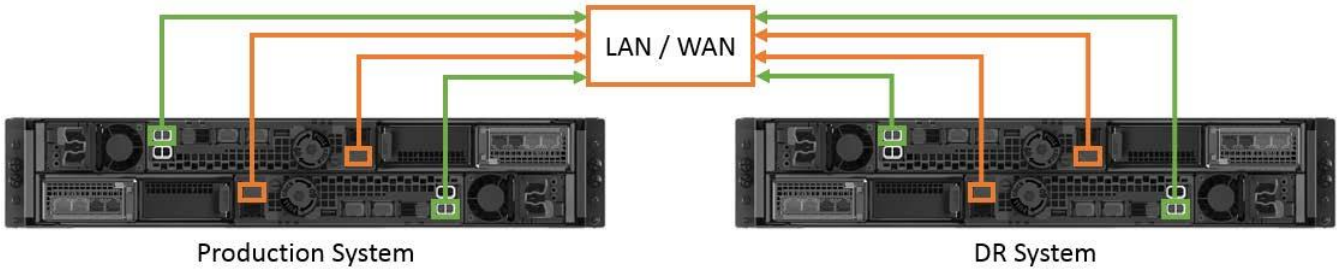

Figure 7: Asynchronous Replication Connectivity

PowerStore snapshots are used for asynchronous replication providing a periodic remote copy implementation. The first replication will be for a full image and subsequent copies will be for changed data only. Dell EMC has provided the following diagram to explain operation.

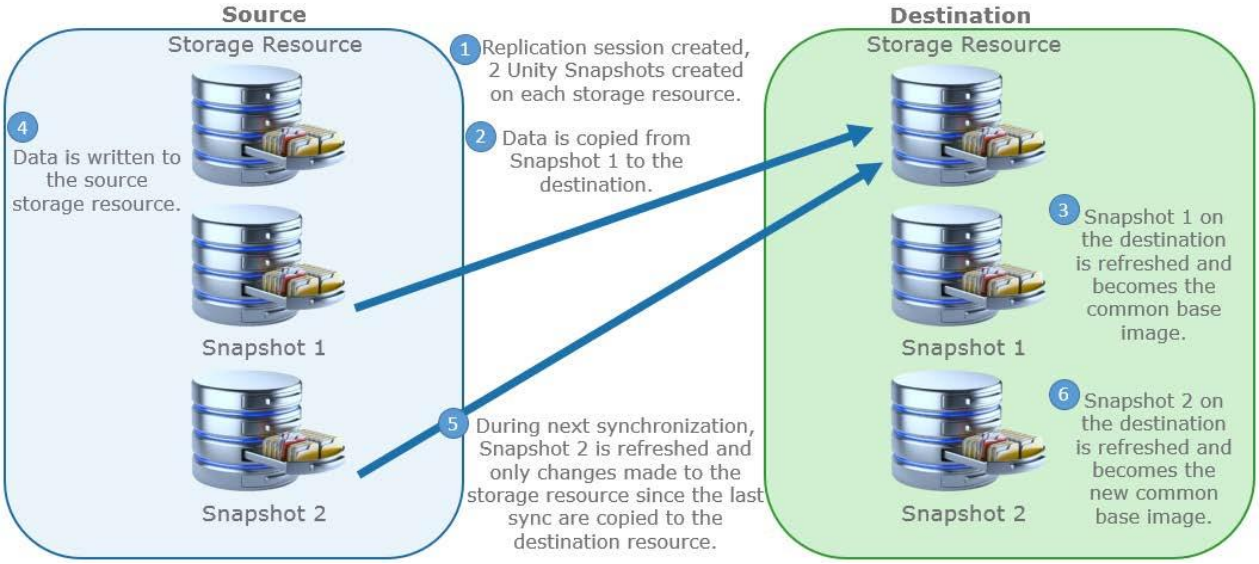

Figure 8: Asynchronous Replication Operation (Source: Dell EMC)

Consistency Groups that contain one or more LUNs are supported. The source and destination system replication resources must be of the same type. For asynchronous replication, these resources are LUNs, Consistency Groups, VMware VMFS Datastores, Filesystems, NAS servers, and VMware NFS Datastores.

Both uni-directional and bi-directional replication are supported as sell as one-to-many and many-to-one. With asynchronous periodic replication, updates at the target site occur at a set interval based on the defined Recovery Point Objective.

MetroSync

PowerStore systems support the MetroSync functionality found in Unity systems. MetroSync provides synchronous file replication over metro areas for disaster recovery, high availability, and data mobility. When using MetroSync, asynchronous replication can be configured to additional sites.

Metro Volumes

Metro Volumes functionality provides block based synchronous replication for PowerStore in a vSphere Metro Storage Cluster.

CloudIQ

Dell EMC uses telemetry data from storage systems including PowerMax for monitoring and storage analytics. Functions include proactive support with health checks and reporting such as capacity consumption and performance. Telemetry data is collected and analyzed with a SaaS product called CloudIQ.

PowerStore NAS

The NAS implementation for PowerStore T models is a 64-bit filesystem that operates in the PowerStoreOS container environment. Virtual file servers are known as NAS Servers and may operate on either controller of the node pair. A NAS Server must be configured before creating a filesystem and is logically isolated from other NAS Servers. Each NAS Server has its own unique information such as name services, file sharing protocol, active directory domain settings, Unix directory services, user mappings, data protection settings and up to 10 interfaces.

NAS Servers use LUNs configured in the system as filesystem storage, leveraging storage pooling.

Model	500	1000T	3000T	5000T	7000T	9000T
Filesystems	500	500	500	500	500	500
Filesystem Size	256 TB	256 TB	256 TB	256 TB	256 TB	256 TB
NAS Servers	50	50	50	50	50	50
Subdirectories per directory	100M	100M	100M	100M	100M	100M
Files per filesystem	32B	32B	32B	32B	32B	32B

File names per directory	10M	10M	10M	10M	10M	10M
ACL IDs	4M	4M	4M	4M	4M	4M

Table 5: PowerStore NAS Specifications

NAS Protocols

PowerStore NAS supports a number of protocols including the following: NFS v3, & v4/V4.1 CIFS/SMB1, 2, 3, FTP, SFTP.

Data Protection

PowerStore NAS offers methods to protect user data including the following:

- NDMP Backup – two-way NDMP is supported

NDMP

The Network Data Management Protocol (NDMP) was developed primarily for backing up NAS systems from remote management applications. NDMP not only can retain permissions for both NFS and CIFS clients, but it also provides the ability to share backup infrastructure, including the backup applications along with SAN tape drives and libraries.

PowerStore NAS requires a read-only filesystem for an NDMP backup, which may be accomplished most easily by utilizing snapshot to create a point in time consistent copy that is used to backup. A snapshot checkpoint is automatically created when an NDMP backup is initiated.

NAS backups, including NDMP backups, operate as a traditional type, known as PAX, or a new volume based backup method know as Volume Backup (VBB). The PAX based method creates backups at the file level, and recursively descends through each directory, backing up each file in that directory.

NDMP Volume Backup reads a set of disk data blocks in an efficient manner compared to the method used for traditional, file-based backups. NDMP Volume Backup works only with Dell EMC-qualified vendor backup software, as listed in the Dell EMC NAS Interoperability Matrix. NAS NDMP Volume Backup can be used when performing a full backup or an incremental backup.

Types of NDMP backup

- Three way backup – the backup is created on tape drives that are SAN attached to PowerStore
- Local Backup – the backup is created on tape drives attached to the PowerStore system
- Remote Backup – the backup is created on remote PowerStore systems tape drives that are attached to the remote system

The following type of restore operations are supported with NDMP volume based backups (VBB):

- Full destructive restore
- File Level restore

Types of PAX Restore:

- Normal
- DAR - direct access restore
- DDAR – directory DAR, and improved form of DAR that enhances performance

The Directory Direct Access Restore (DDAR) optimizes recovery by allowing the NDMP client to directly access backed up data anywhere on the tape set when used with a backup vendor that supports DDAR. DDAR is not compatible with volume backups.

Each Data Mover supports four concurrent NDMP sessions at one time. For example, a system can run three backup sessions and one restore session, simultaneously.

NDMP restrictions:

- NDMP based backups and restores of iSCSI LUNs is not supported
- NDMP does not follow symbolic links within filesystems, therefore NDMP backups do not include the target of the link

VAAI Support for NAS

PowerStore supports VMware vSphere APIs for Array Integration (VAAI) for NAS introduced starting with vSphere 5.0. The features are exploited by VMware to accelerate performance and provisioning time.

- Full File Clone – enables the system to make copies of virtual disks (VMDK files) in the storage system
- Extended Statistics – enables vSphere to query space utilization. Information returned includes the size of the file and the space consumed by the file.
- Space Reservation – enables thick virtual disk files to be provisioned with the lazy-zeroed option.

Quotas

Both soft (advisory) and hard (restrictive) quotas may be set for the amount of storage space consumed, or the number of files utilized. These quota restrictions may be applied on a per user, a per group basis or in combinations.

PowerStore NAS can track quotas using either a block or a file quota policy. The default policy is set to track the number of blocks utilized. Any file smaller than a filesystem block size of 8 KB is counted as 8KB towards the quota since the file consumes one 8 KB block on disk.

If quota policy is set to file-size, disk usage is calculated in logical file size, which means a 1 KB file counts as 1 KB in the quota. When using the file-size quota policy, NAS quota usage is unaffected by file migration and recall operations because these operations do not impact the logical size of files.

Anti-Virus Scanning

PowerStore NAS supports the use of external virus (AV) scanning engines through a mechanism known as AVA (or an Anti-Virus Agent). AVA does not actually perform anti-virus scanning; instead, it acts as an interface and method for allowing an AV agent to scan a file. AVA is an Dell EMC proprietary interface and method of virus scanning. Another developing industry standard is the ICAP interface (Internet Content Adaptation Protocol).

The AV software on an AVA server in conjunction with virus scanning software such as Symantec Norton, McAfee or others is configured to scan network drives on access.

There are two primary methods for utilizing AV scanning, either real-time (as files are accessed) or scheduled scanning. These methods may be used individually or together.

The order of operations is:

1. File open via CIFS to NAS
2. NAS locks the file and ships it to the AVA server(s) via the agent service (AVA) on a remote server
3. Real-time AV software scans the file
4. Once determined clean, AVA agent tells NAS to release the lock

AVA is not a feature that is installed on the operating environment, it is an agent installed on a separate server, together with anti-virus software on a standalone server. AVA by itself does not perform virus scanning.

Policy-Based File Migration

Policy-Based File Migration (formerly File Mover) enables administrators to place file based data according to polices. Administrators may create rules, known as policies that determine whether files should be migrated from primary storage to slower, less-expensive storage devices and back again. Commonly used migration policies are those that migrate older, less frequently used files or files of a particular type.

File mover reads file attributes and will move files by matching the rules or polices that have been established. File mover is able to utilize NFS, CIFS or HTTP to scan files. File Mover will not operate on a read-only filesystem, nor does it support the use of the NFS v4 protocol.

File Mover has added support to migrate files to clouds. Cloud targets include VirtuStream cloud Amazon S3 and Microsoft Azure.

When the policy and migration software finds a match, it migrates or moves the file from PowerStore. It does so by reading the file and then writing it to the secondary file server using NFS or CIFS. Next, the policy and migration software issues a command via the File Mover API to convert the file to a stub file. This releases the storage associated with the file and completes the migration operation. When a file is migrated, its data is moved to a file on secondary storage, and the original file is replaced with a stub file. The stub file contains all the metadata associated with the original file including permissions, timestamps, size and other attributes. The stub file also contains additional data associated with the File Mover functionality, including the location of the file content on secondary storage, the identity of the policy and migration software that migrated the file's content, and an information field in which the policy and migration software can store any additional information. A stub file consumes one inode

(one 8KB data block), plus the size of any alternate data streams associated with the file, in the primary file system on the system.

When a write operation occurs, PowerStore uses the information in the stub file to locate the file on secondary storage. It then compares the modification time (mtime) and size it has stored to the current mtime and size of the file on secondary storage. This comparison ensures the stub file and migrated file remain synchronous. When these attributes match, file data is returned to the client to satisfy the read request. If the attributes differ, an error message appears.

Reliability, Availability, Serviceability Features

PowerStore storage systems maintain data integrity and provide data availability via the following mechanisms:

- Hot swappable controllers
- Mirrored write cache
- RAID protection level 5 and 6
- Battery backup and cache vault to flash in NVRAM in case of power failure
- Redundant power supplies
- Redundant cooling modules
- Redundant data paths
- Online upgrades
- Support for data-in-place upgrades

Redundant write cache

The PowerStore system has dual write cache NVRAM that enhances data integrity and provides for higher performance for mirroring and data replication. The cache also has battery backup and is destaged to the internal NVRAM flash to provide persistence.

PowerStore Metro Node

PowerStore systems can achieve business continuity through active/active stretched clusters with the use of an additional hardware appliance called Dell EMC PowerStore Metro Node. PowerStore metro node is a 1U appliance with support for 32Gb/s Fibre Channel. Two 1U metro nodes are required for each PowerStore cluster to achieve business continuity.

Performance

Dell EMC does not support running standard benchmarks of their systems and does not report standard benchmark data. Dell EMC does provide marketing information regarding performance however. Information about PowerStore performance is limited at this point. Dell EMC has disclosed that the bandwidth from a node pair is 25 GB/s.

Evaluator Group Comment: The performance is really an unknown because industry standard benchmarks such as SPC (Storage Performance Council) and IOMARK have not been run and reported. This leaves customers with taking a chance on the vendor numbers or running the system in their own environment to determine if the system will meet requirements. Industry standard benchmark information would be very useful for customers.

Evaluator Group Comments

Dell EMC has consolidated to a midrange storage system with many of the capabilities of the earlier offerings including Unity XT, XtremIO, SC, and PS. The new platform does not have the risks that come with a completely new system as many of the elements come from Unity as a base. The transition to being "containerized" began with Unity and has continued with further advances. Currently NVMeoF is available over Fibre Channel NVMe/TCP. NVMeoF adds significant performance benefits, and is a competitive feature amongst all flash arrays.

The ability to scale out with up to four node pairs (8 controllers) answers a major limitation that Dell EMC has had for their midrange storage systems. The scaling is not with shared backend data between nodes but is a federation of the node pairs with the ability to move data volumes between them for balancing. The scaled out systems are managed as a single system.

The offering of the X models where VMware ESXi is installed on the hardware and PowerStoreOS runs within a VM is interesting. Dell EMC focuses the message that the storage system can be used at the edge where applications can run on the same platform as the storage (as VMs) and the feature of VMware such as vMotion, vRealize, and VMware Cloud Foundation are all available to add value to management, provisioning, along with application and movement. This capability will be applicable beyond just edge instances. However, there is a difference between the X models with VMware and the T models where PowerStoreOS is installed on the bare metal of the purpose built hardware. The X models have added scale out functionality, but cannot be combined with T models and NAS support is not available. This will be a bit confusing for some.

Strengths:

- ***A mature set of functional storage software providing low-risk for customers***
- ***Dell EMC support and service***
- ***High performance system with use of SCM, flash, DRAM caching, and large number of processor cores.***
- ***Integration with VMware and X models allowing applications to run in a VM on the storage system.***
- ***Scale out for both T and X models to four node pairs***
- ***VAAI support for improved virtual machine provisioning***
- ***NVMeoF support with Fibre Channel and NVMe/TCP***
- ***Entry level option with PowerStore 500 model.***

Perceived Challenges:

- ***NVMe protocol for devices is limited to 21 devices in the base controller. Other attached enclosures use SAS devices. This can be somewhat confusing and may require some special administrative effort.***
- ***The capacity scale (maximum) is much less than some competing products.***

More detailed information is available at **http://evaluatorgroup.com**