

Dell EMC PowerStore: Introduction to the Platform

Abstract

The white paper provides an overview of the Dell EMC™ PowerStore™ platform. It details the value proposition, architecture, and various deployment considerations of the available PowerStore appliances.

April 2021

Revisions

Date	Description
April 2020	Initial release: PowerStoreOS 1.0
September 2020	Minor updates
December 2020	PowerStoreOS 1.0.3 updates: Discovery through static IP
January 2021	Hardware overview and Metro node updates
April 2021	PowerStoreOS 2.0 updates

Acknowledgments

Author: Ethan Stokes

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.

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

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

Copyright © 2020–2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [4/20/2021] [Technical White Paper] [H18149.4]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	5
Audience	5
1 Introduction.....	6
1.1 Terminology	7
2 PowerStore overview	9
3 Hardware overview.....	11
3.1 High availability.....	12
3.2 Base enclosure.....	12
3.2.1 PowerStore 1000–9000.....	13
3.2.2 PowerStore 500.....	18
3.3 Expansion enclosure	23
3.4 Drive model comparison.....	23
3.4.1 NVMe SSD	24
3.4.2 NVMe SCM.....	24
3.4.3 NVMe NVRAM drive.....	25
3.4.4 SAS SSD	26
3.5 I/O module	26
3.5.1 25 GbE optical (4-port)	26
3.5.2 32 Gb Fibre Channel (4-port)	27
3.5.3 10 GbE BASE-T (4-port).....	27
4 PowerStore T models.....	28
4.1 Deployment.....	28
4.1.1 Networking and cabling	28
4.1.2 Discovery	31
4.1.3 Initial configuration.....	32
4.2 Deployment modes.....	33
4.2.1 Unified.....	33
4.2.2 Block optimized.....	34
5 PowerStore X models.....	35
5.1 Deployment.....	36
5.1.1 Networking and cabling	36

5.1.2	Discovery	37
5.1.3	Initial configuration.....	39
6	Conclusion.....	41
A	Technical support and resources	42

Executive summary

Dell EMC™ PowerStore™ is designed with two purpose-built platforms. The 2U, two-node dual socket Intel® Xeon® platform is used for the PowerStore 1000, 3000, 5000, 7000, and 9000 models. The 2U, two-node single-socket Intel® Xeon® platform is used for the PowerStore 500 model. PowerStore provides a data-centric, intelligent, and adaptable infrastructure that supports both traditional and modern workloads. This white paper details the hardware platform, drives, and various physical components. This document includes cabling and deployment guidance for all PowerStore model appliances.

Audience

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

1 Introduction

In this constantly changing world of increasing complexity and scale, the need for an easy-to-use, intelligent storage system has only grown greater. Organizations that use new applications and solutions require dependable storage and often face the challenge of doing more with less. Dell EMC PowerStore addresses this challenge by packaging a powerful storage system into a cost- and space-efficient profile. Some of the key PowerStore value propositions listed as follows.

Active/active architecture: PowerStore uses both nodes to serve host I/O and run data operations in an active/active manner. This design efficiently uses all available hardware resources and optimizes performance, cost, and density in data centers.

NVMe platform: PowerStore is designed to use the latest storage and interface technologies to maximize application performance and eliminate bottlenecks. PowerStore can maximize performance with NVMe flash storage and supports Intel Optane™ storage class memory (SCM) which approaches the speed of DRAM. PowerStore supports front-end NVMe connectivity with NVMe over Fibre Channel for a complete end-to-end NVMe solution.

AppsON: Integration of the PowerStore container-based architecture with onboard VMware ESXi™ results in a new level of consolidation for enterprise storage. This consolidation provides the benefits of a local, on-array application environment and integrates with the VMware® vSphere® management environment and server resources. This ability allows users to bring applications closer to storage by running applications as virtual machines that run directly on PowerStore. AppsON enables agility for application deployments and allows seamless movement between the PowerStore appliances and VMware ESXi™ servers. It also helps reduce the server and networking footprint for space-efficient edge and remote deployments. Complemented by joint engineering work with VMware and Intel, AppsON uses intellectual property to bypass the hypervisor. This ability enables bare-metal NVMe performance with full support for plug-and-play functionality and PCIe fault containment.

VMware integration: PowerStore is designed to have deep integration with VMware vSphere including VAAI, VASA, event notifications, snapshot management, VMware vSphere Virtual Volumes™ (vVols), and virtual machine discovery and monitoring in PowerStore Manager.

Unified offering: PowerStore has a single architecture for block, file, and vVols. This architecture uses the latest technologies to provide flexible functionality without sacrificing the cost-effective nature of enterprise storage. PowerStore provides storage in multiple formats to applications, ranging from physical and virtual volumes to containers and traditional files. This ability provides the ultimate workload flexibility and enables IT to simplify and consolidate infrastructure.

A modern, simple interface: PowerStore Manager, the PowerStore management interface, is built with the data-center administrator in mind. Using browser-native HTML5, PowerStore Manager can be used across various operating systems and web browsers without requiring an external management server or appliance.

Inline data reduction: Data reduction technologies play a critical role in environments in which storage administrators are attempting to do more with less. PowerStore data reduction supports this effort by optimally reducing the amount of physical storage that is required to save a dataset. PowerStore data reduction provides space savings by using software data deduplication and compression through hardware assist. The storage system always enables and intelligently controls data reduction.

Native data protection: Security and availability of data are critical concerns for many organizations, and PowerStore storage offers multiple solutions to address this need. Snapshots provide point-in-time copies of block, file, and virtual machine data that can be used for backup and restoration purposes. Asynchronous

replication offers an IP-based replication strategy within a system or between two systems. Data at Rest Encryption (D@RE) ensures that user data on the system is protected from physical theft and can substitute drive disposal processes, such as shredding.

1.1 Terminology

The following terms are used with PowerStore.

PowerStore Manager: An HTML5 user interface used to manage PowerStore systems.

PowerStore T model: Container-based storage system that is running on purpose-built hardware. This storage system supports unified (block and file) workloads, or block-optimized workloads.

PowerStore X model: Container-based storage system that runs inside a virtual machine that is deployed on a VMware hypervisor. Besides offering block-optimized workloads, PowerStore also allows users to deploy applications directly on the array.

Appliance: 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.

Node: Storage controller that provides the processing resources for performing storage operations and servicing I/O between storage and hosts. Each PowerStore appliance contains two nodes.

Cluster: One or more appliances in a single grouping and management interface. Clusters are expandable by adding more appliances to the existing cluster, up to the allowed amount for a cluster.

Base enclosure: Enclosure containing both nodes (node A and node B) and 25 x NVMe drive slots

Expansion enclosure: Enclosures that can be attached to a base enclosure to provide additional storage.

Fibre Channel (FC) protocol: Protocol used to perform IP and SCSI commands over a Fibre Channel network.

File system: Storage resource that can be accessed through file-sharing protocols such as SMB or NFS.

iSCSI: Provides a mechanism for accessing block-level data storage over network connections.

Volume: A block-level storage device that can be shared out using a protocol such as iSCSI or Fibre Channel.

Network-attached storage (NAS) server: File-level storage server used to host file systems. A NAS server is required to create file systems that use SMB or NFS shares.

Network File System (NFS): An access protocol that allows data access from Linux® or UNIX® hosts on a network.

NVMe over Fibre Channel (NVMe-FC): Allows hosts to access storage systems across a network fabric using the NVMe protocol.

PowerStore Representational State Transfer (REST) API: Set of resources (objects), operations, and attributes that provide interactive, scripted, and programmatic management control of the PowerStore cluster.

Server Message Block (SMB): An access protocol that allows remote file data access from clients to hosts on a network. SMB is typically used in Microsoft® Windows® environments.

Snapshot: A point-in-time view of data stored on a storage resource. A user can recover files from a snapshot, restore a storage resource from a snapshot, or provide access to a host.

Thin clone: Read-write copy of a volume, volume group, file system, or snapshot that shares blocks with the parent resource.

PowerStore Command Line Interface (PSTCLI): Interface that allows a user to perform tasks on the storage system by typing commands instead of using the user interface.

vSphere API for Array Integration (VAAI): A VMware API that improves ESXi host utilization by offloading storage-related tasks to the storage system.

vSphere API for Storage Awareness (VASA): A VMware vendor-neutral API that enables vSphere to determine the capabilities of a storage system. This feature requires a VASA provider on the storage system for communication.

VMware vSphere Virtual Volumes (vVols): A VMware storage framework which allows VM data to be stored on individual Virtual Volumes. This ability allows for data services to be applied at a VM-level of granularity and according to SPBM. Virtual Volumes can also refer to the individual storage objects that are used to enable this functionality.

2 PowerStore overview

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

PowerStore brings the simplicity of public cloud to on-premises infrastructure, streamlining operations with an integrated machine-learning engine and seamless automation. It 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.



Figure 1 PowerStore appliance

The PowerStore platform is available in two different product model types: PowerStore T and PowerStore X. PowerStore T models are bare-metal, unified storage arrays which can service block, file, and vVol resources along with numerous data services and efficiencies. PowerStore T models are perfect for traditional and modern workloads, with examples including relational databases, electronic medical record applications, content repositories, and many more.

PowerStore X model appliances enable running applications directly on the appliance through the AppsON capability. A native VMware ESXi layer runs embedded applications alongside the PowerStore operating system, all in the form of virtual machines. This feature is in addition to the traditional storage functionality of PowerStore X model appliances, which supports serving external block and vVol storage to servers with FC and iSCSI. This innovative design is perfect for storage-heavy applications, providing additional compute and high-performance storage to an existing environment, or any scenario where density, performance, and availability are primary factors.

Beyond the power of a single PowerStore model appliance, multiple PowerStore model appliances can be grouped into a cluster. A PowerStore cluster can consist of a single-appliance or scale up to four PowerStore appliances in a single cluster. The cluster enables scaling the compute, storage, and connectivity of the PowerStore solution while managing multiple appliances from a single control plane. It can also migrate resources between appliances and intelligently load balance new applications based on storage metrics.

Metro node is an external hardware and software add-on feature for PowerStore, and it provides active/active synchronous replication plus standard local use cases. It also provides a solution locally with the local mirror

feature to protect data from a potential array failure. Both use cases provide solutions for true continuous availability with zero downtime.

PowerStore is viewed by metro node as an Asymmetric Logical Unit Access (ALUA) array based on SCSI response data, and is required to follow the four-active, four-passive path connectivity rules. This rule states that both nodes of the metro node must each have four active and four passive paths to all volumes provisioned from the array. For more information about Metro node, see the document [Dell EMC PowerStore Metro Node](#).

3 Hardware overview

The purpose-built PowerStore system is offered in multiple physical hardware models in both PowerStore T and PowerStore X appliances. The PowerStore T model series starts with the PowerStore 500T and scales up to the PowerStore 9000T. The PowerStore X model series starts with the PowerStore 1000X and scales up to the PowerStore 9000X. A letter **T** or **X** may be listed at the end of the model number, indicating whether the specific appliance is a PowerStore T or PowerStore X model, respectively. For each model numeral, such as 1000 (for PowerStore 1000T or 1000X models), the hardware specifications are identical. The two different PowerStore options are often grouped when hardware is referenced, and the T or X is omitted from the model number. See Table 1 for model comparisons.

The PowerStore 500 is a new PowerStore model introduced in the PowerStoreOS 2.0 release. This model runs on a dual-node, single-socket Intel® Xeon® platform and runs the PowerStoreOS 2.0 software. The PowerStore 500 runs the same PowerStoreOS as all other PowerStore T models. PowerStoreOS 1.0 and all PowerStoreOS 1.0 service packs are not supported on the PowerStore 500 appliance.

The system limits vary depending on the PowerStore model. You can learn more about the system limits on the [PowerStore Info Hub](#).

Table 1 PowerStore model comparison

	PowerStore 500	PowerStore 1000	PowerStore 3000	PowerStore 5000	PowerStore 7000	PowerStore 9000
NVRAM drives	0	2		4		
Maximum storage drives (per appliance)	25	96				
Supported drive types	NVMe SCM, NVMe SSD	NVMe SCM, NVMe SSD, SAS SSD ¹				
4-port card	25/10/1 GbE optical/SFP+ and Twinax	25/10/1 GbE optical/SFP+ and Twinax or 10/1 GbE BASE-T				
2-port card	10 GbE optical/SFP+ and Twinax	-				
Supported I/O modules (2 per node)	32/16/8 Gb FC or 16/8/4 Gb FC 25/10/1 GbE optical/SFP+ and Twinax (PowerStore T only) 10/1 GbE BASE-T (PowerStore T only)					
Supported expansion enclosures	None	2.5-inch 25-drive SAS SSD				

¹ SAS SSDs only supported in expansion enclosure.

3.1 High availability

PowerStore features fully redundant hardware and includes several high availability features. These features are designed to withstand component failures within the system itself and environmental failures such as network or power outages. If an individual component fails, the storage system remains online and continues to serve data. The system can also withstand multiple failures if they occur in separate component sets. After the administrator is alerted about the failure, they can order and replace the failed component without impact.

PowerStore is a dual-node architecture which includes two identical nodes for redundancy. It features an active/active controller configuration where both nodes are servicing I/O simultaneously. This increases hardware efficiency, since there are no requirements for idle standby hardware. These nodes, along with up to twenty-five 2.5-inch drives, are enclosed within the base enclosure, all in a 2U form factor.

The following sections cover the different hardware components of the PowerStore platform. Specific sections detail the redundancy and high availability features of that component and how it pertains to PowerStore. For more information about high availability at the software and cluster level, see the document [Dell EMC PowerStore: Clustering and High Availability](#).

3.2 Base enclosure

The PowerStore base enclosure supports 25 all-NVMe 2.5-inch drives in a 2U chassis. The base enclosure is secured into a rack using toolless snap-in rails. The rails ship with every system and allow for easy installation of a PowerStore system. The base enclosure securely latches onto the snap-in rails when fully inserted into the rack. If you must remove the enclosure from the rails, lift a bottom latch on each side of the base enclosure to pull the base enclosure out. While the base enclosure securely latches onto the rails, there are optional screws underneath each latch which you can tighten for additional security.

The front of the base enclosure contains an LED to display different states of the system. This LED is in the upper-left side of the chassis near drive slot 0. The LED states and corresponding system status are shown in Table 2. Each of the 2.5-inch drives contain both a drive power and activity LED, and a drive fault LED. The drive fault LED illuminates amber when a drive becomes faulted. There is also an option in PowerStore Manager to blink a specific drive to identify it using the fault LED. If the drive-power and activity LED is powered on and active, it blinks blue.

Table 2 Base enclosure LED status description

LED state	System status
Blue	Power is on. No fault has occurred.
Amber	Power is on. Fault has occurred within the enclosure.
Blue and amber alternating	Power is on. System is not initialized.
Off	Power is off.

The data storage drives are populated from left to right, starting in slot 0. PowerStore requires a minimum of six storage drives. User data, metadata, and system data are automatically stored and protected across the available storage drives using the PowerStore Dynamic Resiliency Engine (DRE). For more information about PowerStore DRE, see the document *Dell EMC PowerStore: Clustering and High Availability*.

3.2.1 PowerStore 1000–9000

The base enclosure for PowerStore 1000, PowerStore 3000, PowerStore 5000, PowerStore 7000, and PowerStore 9000 models supports twenty-five all-NVMe 2.5-inch drives in a 2U chassis. Figure 2 shows the base enclosure for a PowerStore system with 21 NVMe SSDs and four NVMe NVRAM drives.



Figure 2 Base enclosure front view

The back of the base enclosure reveals the nodes and their connectivity options (Figure 3). Each node has one embedded module and two I/O module slots (optional) for network connectivity. Each node has a dedicated 1 GbE BASE-T service port which can be used for on-site support access and initial configuration of the system. Each node also contains a second 1 GbE BASE-T port. This port is used for management traffic in PowerStore T model appliances and is unused in PowerStore X model appliances. These embedded 1 GbE BASE-T ports are both contained on the embedded module of the node.

For more information about the base enclosure and base enclosure components, see the [Dell EMC PowerStore Hardware Information Guide](#).



Figure 3 Base enclosure back view

3.2.1.1 Drives

The base enclosure is an all-NVMe platform, capable of supporting NVMe SSD, NVMe SCM, and NVMe NVRAM drives. You can populate slots 0 through 20 with NVMe SSD or NVMe SCM drives. Slots 21 through 24 are reserved for NVMe NVRAM drives which serve as additional system write caching. PowerStore 1000 and PowerStore 3000 models contain two NVMe NVRAM drives in slots 23 and 24. PowerStore 5000, PowerStore 7000, and PowerStore 9000 models contain four NVMe NVRAM drives in slots 21 through 24. In models that only use two NVMe NVRAM drives, slots 21 and 22 are not available for storage drives.

Systems running PowerStoreOS 1.0 cannot mix NVMe SSD and NVMe SCM drives in the same base enclosure. PowerStore systems require a minimum of six NVMe SSD or six NVMe SCM drives, which can be scaled up in single-drive increments.

Starting with PowerStoreOS 2.0, systems can support one or more NVMe SCM drives mixed with NVMe SSD drives for metadata tiering. NVMe SCM drives provide lower latency than NVMe SSD drives and can improve system performance by storing metadata on these low-latency drives. On a system which contains both NVMe SCM and NVMe SSD, the NVMe SCM drives are dedicated to metadata and all user data is stored on NVMe SSD drives. PowerStore requires a minimum of six NVMe SSD drives for systems containing entirely NVMe SSD, or systems containing NVMe SSD supplemented with NVMe SCM. PowerStore requires a minimum of six NVMe SCM drives for system containing entirely NVMe SCM.

The expansion enclosures support SAS drives for expansion capacity beyond the all-NVMe base enclosure. Expansion enclosures are not supported on systems containing entirely NVMe SCM drives. Systems running PowerStoreOS 2.0 and containing a mix of NVMe SSD and NVMe SCM drives do support adding an expansion enclosure. For more details about PowerStore drives, see section 3.4.

3.2.1.2 Node

The purpose-built PowerStore platform is powered by dual-socket Intel® Xeon® processors. Each purpose-built PowerStore system contains two nodes, which are used for high-availability and load-balancing purposes.

Each node is 1U in size and stacks vertically in the base enclosure, with the top node inverted. The bottom PowerStore node is node A, and the top PowerStore node is node B. Each node can access each drive through the midplane connection inside the base enclosure. Each node contains the following components, which are detailed in the following sections.

- Internal M.2 boot module
- Fan modules
- Battery backup unit
- DIMMs
- Embedded module
- I/O module
- Power supply

3.2.1.3 Internal M.2 boot module

Each node has a primary and a secondary M.2 SATA device inside the system on a riser card between DIMM slots 11 and 12 (see Figure 4). The primary M.2 device is 240 GB and the primary boot device for the node. PowerStore uses this device to store the base operating system, log files, and for general system operations. The secondary M.2 device is 120 GB. PowerStore uses this device for recovery during a primary M.2 failure, and it is an alternate location for log files.



Figure 4 Internal M.2 boot module

3.2.1.4 Fan modules

PowerStore uses fan modules (cooling modules) to provide cool airflow to the node interior to ensure that the internal components remain at optimal operating temperatures (see Figure 5). Each node contains seven redundant fan modules that are connected to the motherboard within the node. A node can tolerate a single fan module fault, and the surviving fans increase their speed to compensate for the faulted module. If two fan modules fault within the same node, the node performs a protective thermal shutdown. A protective thermal shutdown gracefully powers off the node, and any resource fails over to the surviving node.




Figure 5 Fan module

3.2.1.5 Battery backup unit

If system power is lost, the battery backup unit (BBU) provides power to the NVRAM drive slots and the baseboard management controller (BMC). This action allows the NVRAM drive drives to vault their volatile data to nonvolatile storage within the same drive and persist the information. Once the NVRAM drives have completed their vault, the BMC powers off the system. The BBU in node A provides power for drive slots 21 and 23. The BBU in node B provides power for drive slots 22 and 24. The NVRAM drives are in mirrored sets consisting of drives in slots 23 and 24. If the PowerStore model supports four NVRAM drives, there is another mirrored set in slots 21 and 22. The node BBUs are configured so that both BBUs power each NVRAM

mirrored pair, ensuring that there is no single point of failure. Each BBU contains sufficient charge to accommodate multiple back-to-back power failures. Once power is resumed, the BBUs gradually recharge.



Figure 6 Battery backup unit

3.2.1.6 DIMMs

Each node contains 24 DDR4 DIMM slots, which are populated in different configurations that are based on the PowerStore model. All host data is written to the NVMe NVRAM drive from DRAM before the host is acknowledged to protect against data loss upon system power failure.



Figure 7 DIMM

3.2.1.7 Embedded module

Each node contains a single embedded module which has different connectivity components. The embedded module supports network connectivity for data storage, management and service access, cluster communication, and SAS connectivity to expansion enclosures. The embedded module contains the following components:

- 4-port card
- Non-maskable interrupt (NMI) button
- Mini-SAS HD back-end ports
- System management port (1 GbE) (used with PowerStore T models only)
- Service port (1 GbE)
- USB port

There are two 4-port card options that are supported in the embedded module:

- 10/1 GbE BASE-T (4-port)
- 25/10/1 GbE optical/SFP+ and Twinax (4-port)

The 4-port card (Figure 8) selected for the embedded module on node A must be the same 4-port card in the embedded module on node B. You can select the 4-port card configuration at the time of ordering but cannot change the configuration later. You must configure ports 0 and 1 with a link speed of at least 10 GbE on the 4-port card. This configuration ensures that intercluster traffic, which uses ports 0 and 1, has sufficient

bandwidth. Ports 2 and 3 support all advertised speeds (10/1 GbE BASE-T or 25/10/1 GbE optical/SFP+ and Twinax).



Figure 8 Embedded module with 10/1 GbE BASE-T 4-port card

3.2.1.8 I/O module

Each node on PowerStore systems can support up to two I/O modules that provide extra connectivity. For the two nodes in a base enclosure, the I/O modules that are configured must match between nodes. During a node failure, matching I/O modules ensure that the peer node can begin servicing I/O using the mirrored I/O module.

PowerStore systems support the following I/O modules:

- 25/10/1 GbE optical/SFP+ and Twinax (4-port)
- 32/16/8/4 Gb Fibre Channel (4-port)
- 10/1 GbE BASE-T (4-port)

For more details about I/O modules, see section 3.5.

3.2.1.9 Power supply

The PowerStore platform contains two power supply units (PSUs) in the base enclosure and offers PSUs with two wattage options that are based on the model. PowerStore 1000T/X, PowerStore 3000T/X, and PowerStore 5000T models support both the 1800 W high-line and 2100 W high-line PSUs. For environments which only provide low-line power, you can use the 2100 W PSU with a step-up transformer. Do not select the 1800 W PSU for environments with low-line power. PowerStore 5000X, PowerStore 7000T/X, and PowerStore 9000T/X models support only the 2100 W PSU. For environments which only offer low-line power, you can use the 2100 W PSU with a step-up transformer. Table 3 shows the PSU specifications per PowerStore model.

Table 3 Power supplies

Model	Power supply wattage	Connector support
PowerStore 1000T/X, 3000T/X, 5000T	1800 W ¹ or 2100 W ²	1800 W: C13/14 or C13/20 2100 W: C19/C20
PowerStore 5000X, 7000T/X, 9000T/X	2100 W ²	2100 W: C19/20

¹ Supports high-line power only

² Supports high-line power and low-line power with step-up transformer

Figure 9 shows the supported connector types.



Figure 9 C19/C20, C13/C20, and C13/C14 connectors

A single power supply (Figure 10) can power the entire base enclosure during a power supply failure. Power supplies can be replaced without having to remove the node. Power supplies are offered for AC power only.



Figure 10 2100 W power supply

3.2.2 PowerStore 500

The PowerStore 500 base enclosure supports 25 all-NVMe 2.5-inch drives in a 2U chassis. The PowerStore 500 is introduced in the PowerStoreOS 2.0 release and requires PowerStoreOS 2.0 as a minimum supported software version.

The back of the base enclosure shows the nodes and their connectivity options. Each node has one embedded module and two I/O module slots (optional) for network connectivity. Each node has a dedicated 1 GbE BASE-T service port which can be used for on-site support access and initial configuration of the system. Each node also contains a second 1 GbE BASE-T port which is used for management traffic. The embedded module of the node contains both of these embedded 1 GbE BASE-T ports.

For more information about the base enclosure and base enclosure components, see the *Dell EMC PowerStore Hardware Information Guide for PowerStore 500 Models* on the [PowerStore Info Hub](#).

3.2.2.1 Drives

The base enclosure is an all-NVMe platform, capable of supporting NVMe SSD and NVMe SCM drives. You can populate slots 0 through 24 with NVMe SSD or NVMe SCM drives. PowerStore 500 systems use internal DRAM for write caching and do not support NVMe NVRAM drives.

PowerStore 500 systems can support one or more NVMe SCM drives mixed with NVMe SSD drives for metadata tiering. NVMe SCM drives provide lower latency than NVMe SSD drives and can improve system performance by storing metadata on these low-latency drives. On a system which contains both NVMe SCM

and NVMe SSDs, the NVMe SCM drives are dedicated to metadata, and all user data is stored on NVMe SSD drives. PowerStore requires a minimum of six NVMe SSD drives for systems containing entirely NVMe SSD, or systems containing NVMe SSD supplemented with NVMe SCM. PowerStore requires a minimum of six NVMe SCM drives for system containing entirely NVMe SCM.

PowerStore 500 systems do not support expansion enclosures. For more details about PowerStore drives, see section 3.4.

3.2.2.2 Node

The purpose-built PowerStore 500 platform is powered by a single-socket Intel® Xeon® Processor. Each purpose-built PowerStore 500 system contains two nodes, which are used for high-availability and load-balancing purposes.

Each node is 1U and stacks vertically in the base enclosure, with the top node inverted. The bottom PowerStore node is node A, and the top PowerStore node is node B. Each node can access each drive through the midplane connection inside the base enclosure. Each node contains the following components, which are detailed in the following sections.

- Internal M.2 boot module
- Fan modules
- Battery backup unit
- DIMMs
- Embedded module
- I/O module
- Power supply

3.2.2.3 Internal M.2 boot module

A single M.2 SATA device is located inside each node across from the CPU and DIMM complex. The M.2 device is 240 GB and is the primary boot device for the node. It stores the base operating system, log files, and vaulted cache data, and is used for general system operations.



Figure 11 Internal M.2 boot module

3.2.2.4 Fan modules

Fan modules (cooling modules) provide cool airflow to the node interior to ensure that the internal components remain at optimal operating temperatures (see Figure 12). Each node contains six redundant fan modules that are connected to the motherboard within the node. A node can tolerate a single fan module fault, and the surviving fans increase their speed to compensate for the faulted module. If two fan modules fault within the same node, the node performs a protective thermal shutdown. A protective thermal shutdown gracefully powers off the node, and any resource fails over to the surviving node.




Figure 12 Fan module

3.2.2.5 Battery backup unit

If the system loses power, the battery backup unit (BBU) provides power to the node to enable cache vaulting. This feature allows the node to encrypt and back up dirty cache data during power loss from the system DRAM to the internal M.2 boot module. This BBU and vaulting process replaces the BBU and NVRAM self-vaulting process of the PowerStore 1000 to PowerStore 9000 models. After each node completes the vault of cached data, the node powers off. Each BBU contains sufficient charge to accommodate multiple back-to-back power failures. After power is resumed, the BBUs gradually recharge. The BBU is located between the third and fourth fan, in the middle of the enclosure.



Figure 13 Battery backup unit

3.2.2.6 DIMMs

Each PowerStore 500 node is configured with DDR4 DIMMs. To protect against data loss, all host data is mirrored to the peer node before the host is acknowledged. Each node is equipped with a BBU to vault cached data to the internal M.2 boot module upon system power failure. For more detail on the PowerStore 500 write caching architecture, see the document the [PowerStore: Data Efficiencies](#).



Figure 14 DIMM

3.2.2.7 Embedded module

Each node contains a single embedded module which has different connectivity components. The embedded module supports network connectivity for data storage, management and service access, and cluster communication. The embedded module contains the following components:

- 25/10/1 GbE optical/SFP+ and Twinax 4-port card (optional)
- 10 GbE optical/SFP+ and TwinAx 2-port card
- Mini-SAS HD back-end ports
- System management port (1 GbE)
- Service port (1 GbE)

The 4-port card is optional, and you can select it at the time of ordering. The 4-port card is required to support unified deployment mode (file services) and clustering with other appliances. If you do not select a 4-port card, you can deploy the system only in block-optimized mode. For more information about the deployment modes, see section 4.2.

If you configure the 4-port card for the embedded module on node A, you must also configure a 4-port card on node B. You cannot add or remove the 4-port card after you complete the initial configuration. If you select a 4-port card, you must configure ports 0 and 1 with a link speed of at least 10 GbE on the 4-port card. This configuration ensures that intercluster traffic which uses ports 0 and 1 has sufficient bandwidth. Ports 2 and 3 are not supported for customer use and are reserved for future PowerStoreOS releases.

You can use the embedded 2-port card for front-end iSCSI connectivity and replication. This 2-port card is a fixed 10 GbE card which comes standard on all PowerStore 500 systems. The 2-port card supports SFPs and passive TwinAx running at 10 GbE speeds.



Figure 15 Embedded module with 4-port card

3.2.2.8 I/O module

Each node on PowerStore 500 systems can support up to two I/O modules that provide extra connectivity. For the two nodes in a base enclosure, the I/O modules that are configured must match between nodes. During a node failure, matching I/O modules ensure that the peer node can begin servicing I/O using the mirrored I/O module.

PowerStore 500 systems support the following I/O modules:

- 25/10/1 GbE optical/SFP+ and Twinax (4-port)
- 32/16/8/4 Gb Fibre Channel (4-port)
- 10/1 GbE BASE-T (4-port)

For more details about I/O modules, see section 3.5.

3.2.2.9 Power supply

The PowerStore 500 platform contains two power supply units (PSUs) in the base enclosure and comes with a PSU which supports 2100 W high-line and 1450 W low-line power. For environments which only provide low-line power, you can use the PSU with a step-up transformer. The PowerStore 500 PSU supports C13/C20 or C13/C14 connectors. Figure 16 shows the supported connector types.



Figure 16 C13/C20 and C13/C14 connectors

A single power supply (Figure 17) can power the entire base enclosure during a power-supply failure. You can replace the power supplies without removing the node. Power supplies are offered for AC power only.



Figure 17 1450 W power supply

3.3 Expansion enclosure

The purpose-built PowerStore 1000 to PowerStore 9000 systems support a 25-drive 2U expansion enclosure (Figure 18) using 2.5-inch SAS drives for extra capacity. Expansion enclosures are not supported with PowerStore systems using only NVMe SCM drives or PowerStore 500 systems. These expansion-enclosure restrictions only apply to the specific base enclosure of the appliance that the expansion enclosure is being added to. Expansion enclosure restrictions on an appliance do not impact other systems within a multi-appliance PowerStore cluster.



Figure 18 25-drive 2.5-inch 2U expansion enclosure (front)

The back of the expansion enclosure (Figure 19) includes LEDs to indicate power and fault status. There are also LEDs to indicate bus and enclosure IDs.



Figure 19 25-drive 2.5-inch 2U expansion enclosure (back)

3.4 Drive model comparison

PowerStore supports multiple drive types and capacity points. The base enclosure is an all-NVMe platform, capable of supporting NVMe SSD, NVMe SCM, and NVMe NVRAM drives. The expansion enclosures support SAS drives for expansion capacity beyond the all-NVMe base enclosure. You can add expansion enclosures only to systems which have fully occupied all 21 slots that support NVMe SSDs in the base enclosure. Expansion enclosures are not supported on systems with only NVMe SCM drives in the base enclosure. A list of all supported drives is available on [Dell Support](#).

SAS SSD, NVMe SSD, and NVMe SCM are considered to be storage drives in PowerStore and are formatted with a 512-byte block size. You can populate slots 0 through 20 on PowerStore 1000 to 9000 models and slots 0 through 24 on the PowerStore 500 model with only NVMe SSD or NVMe SCM drives. Starting in PowerStoreOS 2.0, you can mix NVMe SSD and NVMe SCM drives in the same enclosure. Earlier releases of PowerStoreOS do not support mixing these drives. PowerStore systems require a minimum of six NVMe SSD or six NVMe SCM drives, which can be scaled up in single-drive increments. Slots 21 through 24 in PowerStore 1000 to 9000 models are reserved for NVMe NVRAM drives which serve as extra system write caching. Based on the PowerStore model, there are either two NVMe NVRAM drives in slots 23 and 24, or four NVMe NVRAM drives in slots 21 through 24. In models that only use two NVMe NVRAM drives, slots 21 and 22 are not available for storage drives.

3.4.1 NVMe SSD

NVMe solid-state drives (SSDs) are dual-ported, high-performance, nonvolatile flash drives that the PowerStore operating system accesses with NVMe. NVMe is a protocol that allows access directly with the PCIe bus. NVMe is designed to capitalize on the low latency of high-performance media.

NVMe SSDs serve as the storage tier for PowerStore, are used for user data or metadata, and come in multiple capacity points. See Table 4 for a complete list of capacity points and supported platforms.

Table 4 NVMe SSD drive size and support

Storage type	Usage	GB	PowerStore 1000–9000	PowerStore 500
NVMe SSD	User data/ metadata	1,920	✓	✓
		3,840	✓	✓
		7,680	✓	✓
		15,360	✓	✓

NVMe SSDs are supported in base enclosure slots 0 through 20 on PowerStore 1000 to PowerStore 9000 models and in base enclosure slots 0 through 24 on PowerStore 500 models. Starting with PowerStoreOS 2.0, you can mix NVMe SSDs with NVMe SCM drives. PowerStoreOS 1.0 and associated service packs do not support mixing NVMe SSD and NVMe SCM drives.

3.4.2 NVMe SCM

NVMe storage class media (SCM) drives are dual-ported, extreme-high-performance, nonvolatile drives that are designed with Intel Optane™ technology. NVMe SCM drives have lower latency and improved performance compared to other SSD drives, and the PowerStore operating system accesses them with NVMe. NVMe is a protocol that allows access directly with the PCIe bus. NVMe is designed to capitalize on the low latency of high-performance media.

NVMe SCM drives serve as a storage tier for PowerStore, are used for user data or metadata, and come in three capacity points. See Table 5 for a complete list of capacity points and supported platforms.

Table 5 NVMe SCM drive size and support

Storage type	Usage	GB	PowerStore 1000–9000	PowerStore 500
NVMe SCM	User data/ metadata	375	✓	✓
		750	✓	✓
		1500	✓	-

NVMe SCM drives are supported in base enclosure slots 0 through 20 on PowerStore 1000 to PowerStore 9000 models and in base enclosure slots 0 through 24 on PowerStore 500 models. Starting with PowerStoreOS 2.0, you can mix NVMe SCMs with NVMe SSD drives. NVMe SCM drives serve as a dedicated metadata tier when mixed with NVMe SSD drives. PowerStoreOS 1.0 and associated service packs do not support mixing NVMe SCM and NVMe SSD drives. A PowerStore system with only NVMe SCM storage drives does not support any expansion enclosures.

3.4.3 NVMe NVRAM drive

NVMe NVRAM drives are dual-ported, extreme-high-performance drives used to enhance the PowerStore caching system. The dual-ported drives are accessible from both nodes and allow the system to easily cache incoming writes. The drives contain dynamic media which can operate at DRAM speeds over PCIe, delivering exceptional performance. Their design allows them to function as nonvolatile media, and PowerStore can quickly store incoming writes and acknowledge the host without mirroring data to the peer node. The drives contain a combination of persistent flash storage within the 2.5-inch NVMe NVRAM paddle card and access to a battery. During a power failure, these features allow the drive to vault the data from the high-performance dynamic media to the persistent flash storage. The NVMe NVRAM drives are 8 GB and are configured in mirrored sets. These drives are supported on PowerStore 1000 to PowerStore 9000 models, and there are two or four drives per appliance, depending on the model (see Table 6).

Table 6 NVMe NVRAM count by PowerStore model

PowerStore model	NVMe NVRAM count
PowerStore 500	0
PowerStore 1000	2
PowerStore 3000	
PowerStore 5000	4
PowerStore 7000	
PowerStore 9000	

The battery backups are wired so that each mirrored set of drives has access to two separate battery backups. This configuration ensures that a faulty battery backup could not result in a failed data vault for an entire mirrored pair. The number of NVMe NVRAM drives are fixed per model, and you cannot modify them. You may not add extra NVMe NVRAM drives later.

For more information about the role of NVMe NVRAM drives in the PowerStore write path, see the document *Dell EMC PowerStore: Data Efficiencies*.

3.4.4 SAS SSD

SAS SSDs are dual-ported, high-performance, nonvolatile flash drives that the PowerStore operating system accesses through the SAS protocol. SAS SSDs are only supported in expansion enclosures and are not supported with PowerStore 500 models.

SAS SSDs supplement the storage tier for PowerStore, are used for user data or metadata, and are available in multiple capacity points. See Table 7 for a complete list of capacity points and supported platforms.

Table 7 SAS SSD drive size and support

Storage type	Usage	GB	PowerStore 1000–9000	PowerStore 500
SAS SSD	User data/ metadata	1,920	✓	-
		3,840	✓	-
		7,680	✓	-

3.5 I/O module

Each node on PowerStore systems can support up to two I/O modules that provide extra connectivity. For the two nodes in a base enclosure, the I/O modules that are configured must match between nodes. During a node failure, matching I/O modules ensure that the peer node can begin servicing I/O using the mirrored I/O module.

PowerStore systems support the following I/O modules:

- 25/10/1 GbE optical/SFP+ and Twinax (4-port) (PowerStore T models only)
- 32/16/8/4 Gb Fibre Channel (4-port)
- 10/1 GbE BASE-T (4-port) (PowerStore T models only)

3.5.1 25 GbE optical (4-port)

The 25 GbE optical I/O module (Figure 20) supports SFPs running at 25 GbE, 10 GbE, or 1 GbE speeds. The optical I/O module ports also support 10 GbE active and passive Twinax, and 25 GbE passive Twinax connections. You can mix different SFPs or Twinax cables on the same I/O module, and they are hot-swappable. This I/O module supports iSCSI traffic.




Figure 20 25 GbE optical I/O module

3.5.2 32 Gb Fibre Channel (4-port)

The 32 Gb Fibre Channel (4-port) I/O module (Figure 21) supports 32 Gb/s and 16 Gb/s SFPs. The 32 Gb/s SFP offers front-end connectivity at 32 Gb/s speeds and can autonegotiate to 16 Gb/s and 8 Gb/s. The 16 Gb/s SFP offers front-end connectivity at 16 Gb/s and can autonegotiate to 8 Gb/s and 4 Gb/s.

Starting with PowerStoreOS 2.0, PowerStore systems with the 32 Gb Fibre Channel I/O module support NVMe over Fibre Channel. NVMe over Fibre Channel support with PowerStore requires 32 Gb speeds, and the Fibre Channel I/O module must be configured with 32 Gb SFPs to support this feature.



Figure 21 32 Gb Fibre Channel I/O module

3.5.3 10 GbE BASE-T (4-port)

The 10 GbE BASE-T (4-port) I/O module (Figure 22) operates at up to 10 Gb/s speeds. It is used for front-end host access and supports both iSCSI and NAS protocols. The I/O module can also autonegotiate to 1 Gb/s speeds as needed.



Figure 22 10 GbE BASE-T I/O module

4 PowerStore T models

PowerStore T model appliances include one of two model types within the PowerStore platform. This distinction is denoted by the letter **T** which follows a given model number, such as PowerStore 9000T. PowerStore T models have the purpose-built PowerStore platform that is detailed in section 3. Built with a microservices and container-based architecture, the PowerStore operating system is designed to take full advantage of the cutting-edge technology of the underlying platform.

PowerStore T models are unified storage arrays. They can serve block storage (with iSCSI, Fibre Channel, or NVMe over Fibre Channel), file storage (with SMB, NFS, FTP, SFTP), and vVol storage (with iSCSI or Fibre Channel). The workflows and operations between block, file, and vVol storage are built on the same engine, and users can control all aspects of the array. The PowerStore single architecture that supports block, file, and vVols allows the platform to accommodate a wide variety of traditional and modern workloads. By providing storage in multiple formats to applications, PowerStore delivers workload flexibility and enables administrators to simplify and consolidate infrastructure.

4.1 Deployment

After you rack the system, a single PowerStore T model appliance requires simple network cabling to the physical Ethernet and management switches. Complete the cabling, and then apply the proper switch configuration to support management and data traffic from the appliance. Finally, connect the power cables and power on the appliance. After the PowerStore system is online, launch the Initial Configuration Wizard and configure the system.

The Initial Configuration Wizard is an HTML5-based configuration wizard that is hosted on the appliance. PowerStore X model appliances also use the same wizard with slight modifications. The Initial Configuration Wizard gathers all necessary information about networking and infrastructure services. The appliance automatically applies this configuration and brings PowerStore into a configured, operational state.

For details about configuring a multi-appliance PowerStore T cluster, adding a PowerStore T model appliance to an existing cluster, or removing an appliance, see the document *Dell EMC PowerStore: Clustering and High Availability*.

For details about PowerStore deployment and configuration in general, see the [Dell EMC PowerStore: Quick Start Guide](#). This document contains an overview of PowerStore deployment and directs readers to all necessary documentation and resources for a successful installation.

4.1.1 Networking and cabling

PowerStore T model appliances require two physical Ethernet switches with Layer 2 connectivity and at least one management switch. The Ethernet switches ensure high availability for iSCSI, NAS, replication, external storage import, data migration, and intercluster traffic. Properly configuring and cabling to the physical Ethernet switches ensures that all capabilities of PowerStore are ready for use when the Initial Configuration is complete.

To ensure that the Ethernet switches can provide high availability to PowerStore, configure them with one of the following Layer 2 interconnect options:

- Multi-chassis Link Aggregation Group (for example, Virtual Link Trunking, Virtual PortChannel, or Multi-Chassis Trunking)
- Reliable L2 uplinks
- Direct Trunk Link

You must meet the minimum cabling requirements listed in Table 8.

Table 8 Minimum cabling requirements for PowerStore T models

PowerStore connection	Switch connection
Node A port 0 of 4-port card embedded module	Physical Ethernet switch 1
Node A port 1 of 4-port card embedded module	Physical Ethernet switch 2
Node B port 0 of 4-port card embedded module	Physical Ethernet switch 2
Node B port 1 of 4-port card embedded module	Physical Ethernet switch 1
Node A 1 GbE management port	Management switch
Node B 1 GbE management port	Management switch

Starting with PowerStoreOS 2.0, single-appliance PowerStore T systems can complete initial configuration with only the management switch connection. This requirement simplifies PowerStore installation and reduces dependencies for the initial configuration. Systems that use NAS, iSCSI, external storage import, or clustering capabilities are not required to be properly configured and cabled to the physical Ethernet switches after initial configuration. The minimum cabling requirements for systems which fit this criteria are listed in Table 9.

Table 9 Minimum cabling requirements for PowerStore T models running PowerStoreOS 2.0

PowerStore connection	Switch connection
Node A 1 GbE management port	Management switch
Node B 1 GbE management port	Management switch

Figure 23 outlines the minimum cabling that is required for PowerStore T appliances. Systems running PowerStoreOS 2.0 only require cabling to the management switch to complete the initial configuration.

This cabling configuration ensures high availability and optimal traffic shaping for the PowerStore T model appliance. For more details about configuring the necessary switches, see the document [Dell EMC PowerStore: Networking Guide for PowerStore T Models](#).

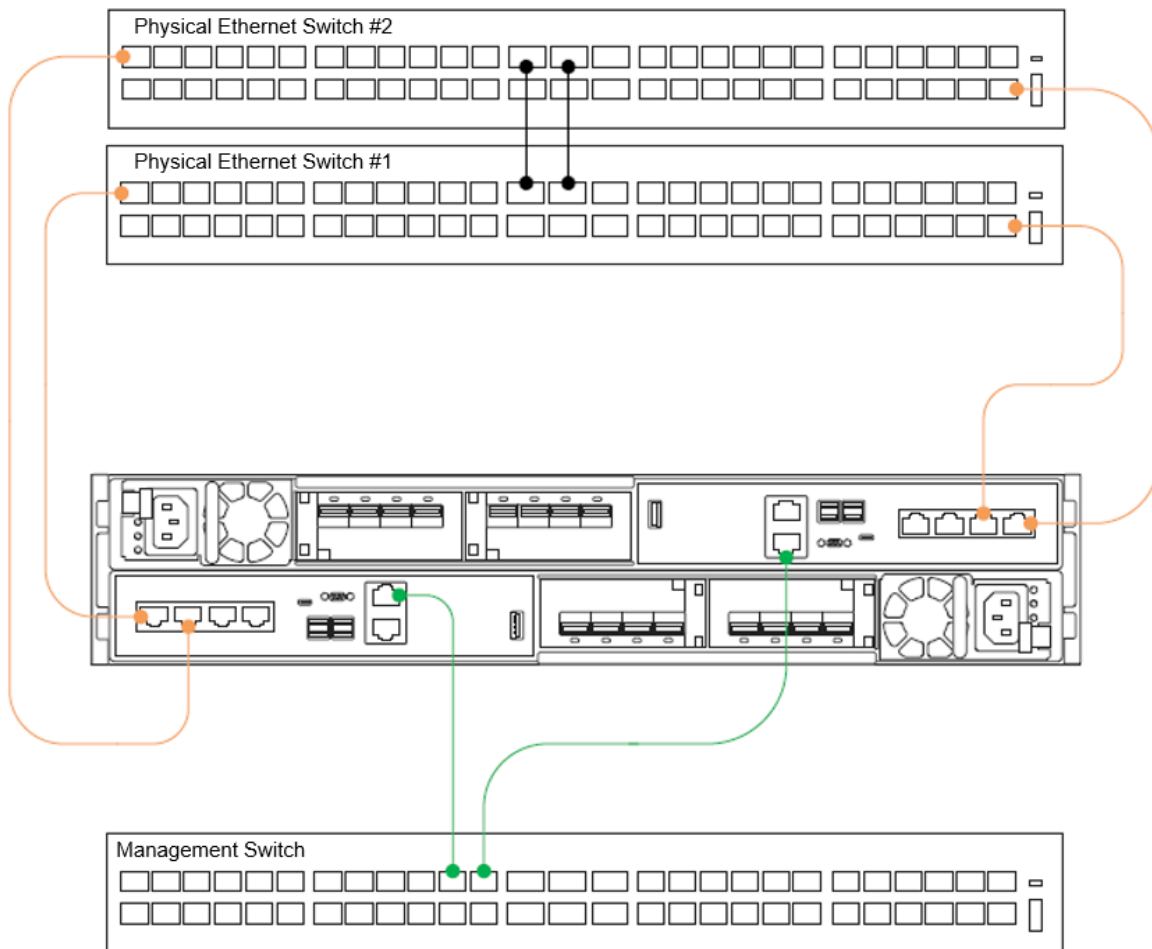


Figure 23 PowerStore T model appliance minimum cabling

PowerStore T model appliances contain a system bond by default on the first two ports of the embedded module 4-port card (Figure 24). PowerStore 500 model appliances that are ordered without a 4-port card do not contain a system bond. This bond provides high availability and potentially increased throughput for iSCSI, NAS, replication, external storage import, data migration, and intercluster traffic. If a link aggregation is not created across the two physical Ethernet switches for this bond, the bond automatically enters an active/passive state. One of the two ports is active for all traffic across the system bond. The passive port remains on standby if the active port experiences a network failure. This configuration is fully supported and delivers high availability during a network issue. If link aggregation is created across the two physical Ethernet switches for this bond, the bond automatically enters an active/active state. This configuration delivers high availability and increased throughput since both physical ports are active.




Figure 24 PowerStore T model system bonds

4.1.2 Discovery

After you cable the PowerStore system and configure the appropriate network settings, perform discovery and initial configuration for PowerStore.

Three discovery options are available:

- Physically connect a workstation to the appliance service port
- Go to the static discovery IP
- Discover the appliance remotely using the PowerStore Discovery Tool application

After the PowerStore system is discovered, step through the HTML5-based Initial Configuration Wizard to complete the configuration.

4.1.2.1 Service port

The preferred method for configuration is to physically connect a workstation to the service port on node A of the PowerStore appliance. A static IP of 128.221.1.249 and subnet mask 255.255.255.0 is set on the workstation. You can log in to PowerStore and begin the Initial Configuration Wizard by directing a browser to 128.221.1.250.




Figure 25 PowerStore Direct Connect Discovery

4.1.2.2 Static IP

Another discovery method available with the PowerStoreOS 1.0.3 release is to use a static IP address. You can access a range of IP addresses to discover the PowerStore and go through the initial configuration.

Perform the following:

1. Discover PowerStore using any of the IP addresses listed below:
 - <https://169.254.0.10>
 - <https://169.254.0.20>
 - <https://169.254.0.30>
 - <https://169.254.0.40>
 - <https://169.254.0.50>
2. Connect the user workstation (also laptop or virtual machine) through an Ethernet cable to the same network (either through the same switch or the same untagged or native VLAN).
3. Ensure that the user workstation has a static IP set in the same 169.254.0.x/16 subnet (for example, 169.254.0.100).
4. The user can open a browser with one of the IP addresses listed above and start the Initial Configuration wizard.

For more information about the static IP address discovery, see the document [Dell EMC PowerStore: Network Planning Guide](#).

4.1.2.3 PowerStore Discovery Tool

If you directly connect a workstation to the PowerStore, or using a static IP is not possible, download the PowerStore Discovery Utility from [Dell Support](#) and install it on a remote machine. The PowerStore Discovery Utility uses zero-configuration technology to detect the PowerStore system Avahi broadcasts. The Avahi broadcasts are sent from the PowerStore T model management port on the native VLAN. The broadcasts also require the remote machine to have an interface on the same network to detect the PowerStore system. The PowerStore Discovery Utility detects all unconfigured and configured PowerStore systems on the network. The utility then automatically launches a web browser to the appliance to begin the Initial Configuration Wizard.

4.1.3 Initial configuration

PowerStore T model appliances use two distinct external networks. However, only the management network is required to configure a PowerStore T model appliance. Optionally, you can configure the storage network, which allows support for iSCSI, data import, and replication traffic when the system is configured. To configure the storage network, a minimum of two IPs are required. If you skip the storage network configuration during the Initial Configuration Wizard, you can configure it at any point through PowerStore Manager. Starting with PowerStoreOS 2.0, the storage network has been removed entirely from the Initial Configuration Wizard to streamline installation. You can configure the storage network in PowerStore Manager after you complete the configuration.

Each network requires several IP addresses, along with gateway, netmask, and VLAN information. If the network uses the access or native VLAN on the switch, **VLAN 0** is used in the Initial Configuration Wizard. While networks can share the same VLAN, the management and storage network must be on different subnets. Table 10 details the specific IP requirements for each network.

Table 10 PowerStore T model Initial Configuration Wizard, IP requirements

Network	Number of IP addresses	Purpose
Management	4	<ul style="list-style-type: none"> • Cluster IP • Appliance IP • 2 x node IP
Storage (optional)	2 required, 1 optional	<ul style="list-style-type: none"> • 2 x iSCSI target IP • Global storage discovery IP (optional)

Besides the networking information, the details for DNS servers and NTP servers are required to complete the Initial Configuration Wizard.

For more information about the Initial Configuration Wizard, see the document [Dell EMC PowerStore: Networking Guide for PowerStore T Models](#) or [Dell EMC PowerStore: PowerStore Manager Overview](#).

4.2 Deployment modes

PowerStore T models support two deployment modes to ensure that the platform delivers maximum performance for each use case. During the Initial Configuration Wizard and after discovery of the PowerStore appliance, you can confirm or modify the deployment mode for the PowerStore T model. The two deployment modes that are offered are **unified** and **block optimized**, which are detailed in the following sections.

You can select the deployment mode during the Initial Configuration Wizard and bring the system online after the configuration is successfully applied. You cannot change deployment modes nondisruptively after the system is configured. These modes require a complete factory reinitialization of the PowerStore appliance to reconfigure the appliance and select a different deployment mode. For this reason, ensure that the proper deployment mode is selected when the system is first configured. The default deployment mode is **unified**, which is the recommended deployment mode since it provides all capabilities of the PowerStore T model appliance.

4.2.1 Unified

The unified deployment mode for PowerStore T models is the default deployment mode. In this configuration, unified PowerStore T model appliances can support block, file, and vVol storage resources. Parts of the underlying hardware resources are reserved for file components and are not available for block resources. For most use cases, the unified deployment mode is preferred because it supports all resources that the PowerStore T model appliance offers. PowerStore 500 model systems must contain a 4-port card to support the unified deployment mode.

4.2.2 Block optimized

The block optimized deployment mode for PowerStore T models is the alternative deployment mode. When you select **Block Optimized** during the Initial Configuration Wizard, a prompt appears that confirms the selection. The appliance reboots and comes online in the block optimized deployment mode. The Initial Configuration Wizard completes with the new deployment mode. Block optimized systems support block and vVol resources. All underlying hardware resources are dedicated to the block performance. Block optimized systems have a higher performance ceiling for block workloads than the same model with a unified deployment mode.

5 PowerStore X models

PowerStore X models make up one of the two model types of PowerStore. This distinction is denoted by the letter **X** which follows a given model number, such as PowerStore 9000X. PowerStore X model appliances are built on the 2U, two-node PowerStore platform that is detailed in section 3.2.1. Built with a microservices and container-based architecture, the PowerStore operating system is designed to take full advantage of the cutting-edge technology of the underlying platform. PowerStore X models have a native hypervisor layer, developed through joint engineering work with VMware and Intel, which enables selected direct device passthrough of VMware ESXi for on-array application processing. This capability allows the system to provide optimal storage performance and low latency, while supporting customer applications running directly on the hypervisor.

PowerStore X models can serve block storage (with iSCSI, Fibre Channel, or NVMe over Fibre Channel) and vVol storage (with iSCSI or Fibre Channel). PowerStore X models can also host native applications directly on the array in the form of vVols. The applications are embedded by running an instance of VMware ESXi directly on each node. Since the PowerStore platform contains two physical nodes, a PowerStore X model appliance consists of two ESXi hosts. Each host is installed directly on the PowerStore node.

Two PowerStore controller VMs run on the ESXi hosts of the PowerStore X model. These VMs provide the active/active redundant capabilities that are standard across all models. The controller VMs run the PowerStore operating system, and each physical PowerStore node contains one ESXi host and one controller VM by default. To ensure proper resource availability for the PowerStore operating system, the ESXi host reserves 50% of the node CPU and memory for the controller VM. The user VMs use the remaining 50% of the node resources. These VMs are protected with high availability (HA) by default, and the PowerStore controller VMs are also redundant.

Figure 26 depicts the architecture for a PowerStore X model appliance. VMware ESXi is installed directly onto the PowerStore platform, and the PowerStore operating system runs as a VM on the VMware ESXi host. Applications can be deployed as VMs directly onto the VMware ESXi host, which use vVols. The platform can serve external-block and vVol storage with Fibre Channel and iSCSI.

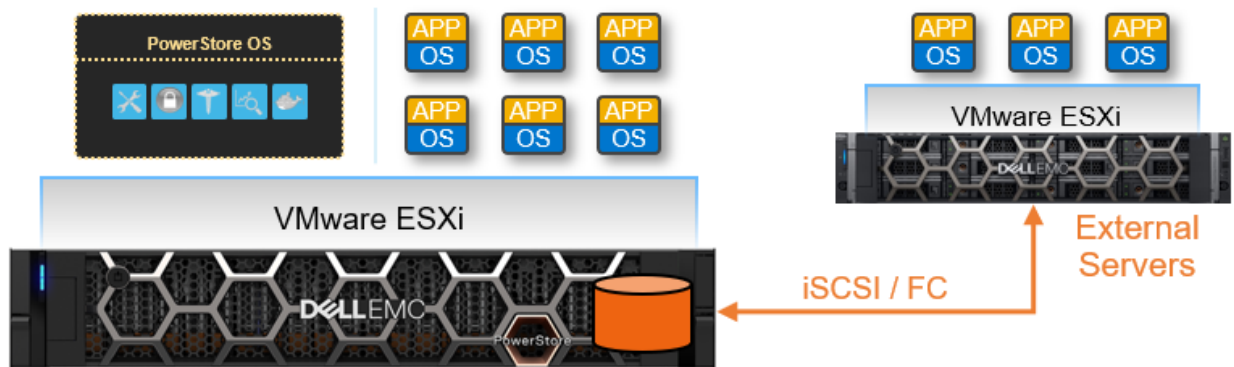


Figure 26 PowerStore X model architecture

For more details about the PowerStore X model components and integration with VMware, see the document [Dell EMC PowerStore: Virtualization Integration](#).

5.1 Deployment

After you have racked the system, a single PowerStore X model appliance requires simple network cabling to the physical Ethernet switches. Completing the cabling, and apply the proper switch configuration to support management and data traffic from the appliance. Finally, connect the power cables and power on the appliance. After the PowerStore system is online, launch the Initial Configuration Wizard and configure the system.

The Initial Configuration Wizard is an HTML5-based configuration wizard that is hosted on the appliance. PowerStore T model appliances also use the same wizard with slight modifications. The Initial Configuration Wizard gathers all necessary information about networking, infrastructure services, and hypervisor details. The appliance automatically applies this configuration and brings PowerStore into a configured, operational state.

Starting with PowerStoreOS 2.0, PowerStore X model appliances support clustering. For details about configuring a multi-appliance PowerStore X cluster, adding a PowerStore X model appliance to an existing cluster, or removing an appliance, see the document *Dell EMC PowerStore: Clustering and High Availability*.

For general information about PowerStore deployment and configuration, see the [Dell EMC PowerStore: Quick Start Guide](#). This guide contains an overview of PowerStore deployment and directs readers to all other necessary documentation and resources for a successful installation.

5.1.1 Networking and cabling

PowerStore X model appliances require two physical Ethernet switches with Layer 2 connectivity. The Ethernet switches ensure high availability for iSCSI, management, replication, external storage import, and VMware vMotion® traffic. Properly configuring and cabling to the physical Ethernet switches ensures that PowerStore features are ready for use when the Initial Configuration is complete.

To ensure that the Ethernet switches can provide high availability to PowerStore, configure them with one of the following Layer 2 interconnect options:

- Multi-chassis Link Aggregation Group (for example, Virtual Link Trunking, Virtual PortChannel, or Multi-Chassis Trunking)
- Reliable L2 uplinks
- Direct Trunk Link

Ensure that you meet the minimum cabling requirements listed in Table 11.

Table 11 Minimum cabling requirements for PowerStore X models

PowerStore connection	Switch connection
Node A port 0 of 4-port card embedded module	Physical Ethernet switch 1
Node A port 1 of 4-port card embedded module	Physical Ethernet switch 2
Node B port 0 of 4-port card embedded module	Physical Ethernet switch 2
Node B port 1 of 4-port card embedded module	Physical Ethernet switch 1

Figure 27 outlines the minimum cabling that is required for PowerStore X model appliances.

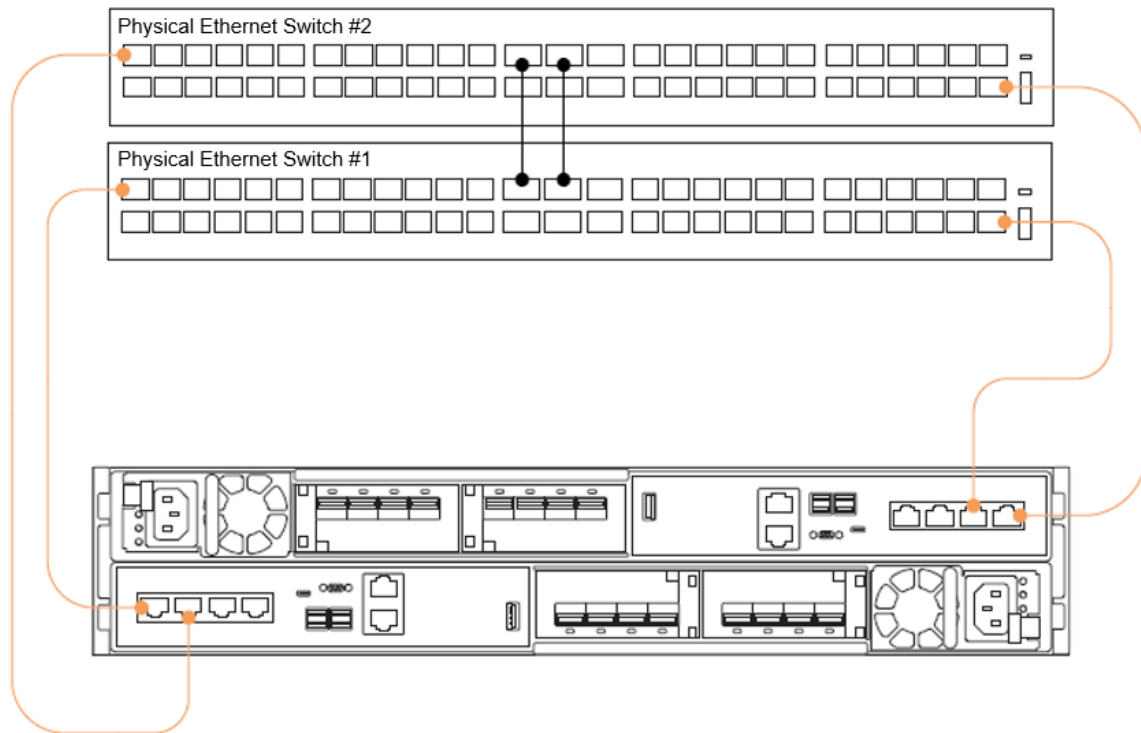


Figure 27 PowerStore X model minimum cabling

This cabling configuration ensures high availability and optimal traffic shaping for the PowerStore X model appliance. For more details about configuring the necessary switches, see the document [Dell EMC PowerStore: Networking Guide for PowerStore X Models](#).

5.1.2 Discovery

After you cable the PowerStore system and configure the appropriate network settings, perform discovery and initial configuration for PowerStore.

Three discovery options are available:

- Physically connect a workstation to the appliance service port
- Go to the static discovery IP
- Discover the appliance remotely using the PowerStore Discovery Tool application

After the PowerStore system is discovered, step through the HTML5-based Initial Configuration Wizard to complete configuration.

5.1.2.1 Service port

The preferred method for configuration is to physically connect a workstation to the service port on node A of the PowerStore appliance. A static IP of 128.221.1.249 and subnet mask 255.255.255.0 are set on the workstation. You can log in to PowerStore and begin the Initial Configuration Wizard by directing a browser to 128.221.1.250. See Figure 28.




Figure 28 PowerStore Direct Connect Discovery

5.1.2.2 Static IP

Another discovery method is available with the PowerStoreOS 1.0.3 release: through a static IP address.

You can access a range of IP addresses to discover the PowerStore and go through the initial configuration.

Perform the following:

1. Discover PowerStore using any of the IP addresses listed below:
 - <https://169.254.0.10>
 - <https://169.254.0.20>
 - <https://169.254.0.30>
 - <https://169.254.0.40>
 - <https://169.254.0.50>
2. Connect the user workstation (laptop or virtual machine) through an Ethernet cable to the same network (either same switch or same untagged/native VLAN).
3. Ensure that the user workstation has a static IP set in the same 169.254.0.x/16 subnet (for example, 169.254.0.100).
4. Open a browser with one of the IP addresses listed above and start the Initial Configuration wizard.

For more information about static IP address discovery, see the document [Dell EMC PowerStore: Network Planning Guide](#).

5.1.2.3 PowerStore Discovery Tool

If you cannot directly connect a workstation to the PowerStore appliance or through a static IP, download the PowerStore Discovery Utility from [Dell Support](#) and install it on a remote machine. The PowerStore Discovery Utility uses zero-configuration technology to detect the PowerStore system's Avahi broadcasts. The PowerStore X system sends Avahi broadcasts from the 4-port card ports 0 and 1 on the native VLAN. The broadcasts require the remote machine to have an interface on the same network to detect the PowerStore system. The PowerStore Discovery Utility detects all unconfigured and configured PowerStore systems on the network. The utility automatically launches a web browser to the appliance to begin the Initial Configuration Wizard.

For more information about discovering and configuring PowerStore, see the document *Dell EMC PowerStore: Quick Start Guide*.

5.1.3 Initial configuration

Three networks are required to configure a PowerStore X model appliance. Each network requires several IP addresses, along with gateway, netmask, and VLAN information. If the network uses the access or native VLAN on the switch, **VLAN 0** is used in the Initial Configuration Wizard. While networks can share the same VLAN, the management and storage network must be on different subnets. Table 12 details the specific IP requirements for each network.

Table 12 PowerStore X Initial Configuration Wizard, IP requirements

Network	Number of IP addresses	Purpose
Management	6	<ul style="list-style-type: none"> • Cluster IP • Appliance IP • 2 x controller VM IP • 2 x ESXi IP
Storage	6 required, 9 for best performance	<ul style="list-style-type: none"> • 2 x iSCSI target IP, 4 x for best performance • 4 x iSCSI initiator IP • Global Storage Discovery IP (optional)
vMotion	2	<ul style="list-style-type: none"> • 2 x vMotion IP

Besides the networking information, you must have the following details to complete the Initial Configuration Wizard:

- DNS servers
- NTP servers
- vCenter IP
- vCenter administrator credentials
- PowerStore credentials

When you configure the system, the storage network provides iSCSI targets to support external iSCSI, data import, and replication traffic. The storage network IP addresses also provide internal iSCSI connectivity between the PowerStore controller VM and then PowerStore ESXi host. The four iSCSI initiators allow the ESXi hosts to establish an iSCSI data path to the controller VM iSCSI target.

Starting with PowerStoreOS 1.0.3, an optimization step is added to the Initial Configuration Wizard. This step will request two additional storage network IP addresses which are automatically configured as iSCSI targets. These optional two iSCSI targets allow the ESXi hosts to have extra paths to the controller VM and improve performance. If you run a PowerStoreOS version that is before version 1.0.3, follow the [PowerStore: PowerStore X Performance Best Practice Tuning](#) knowledge base article to add the additional storage network IPs and complete the best practice configuration.

For more information about the Initial Configuration Wizard, see the document *Dell EMC PowerStore: Networking Guide for PowerStore X Models* or *Dell EMC PowerStore: PowerStore Manager Overview*.

When the Initial Configuration begins, several different tasks are automated from a vSphere perspective. The resources can be automatically deployed into an existing data center. Otherwise, a new resource is created that is based on configuration settings that are specified in the Initial Configuration Wizard. The two ESXi hosts of the PowerStore X model are joined to the target vCenter as a new ESXi cluster. A Distributed Virtual Switch is created for that ESXi cluster to join the two PowerStore nodes from a networking perspective. A

default vVol datastore is mounted and reflects the usable capacity of the appliance. Finally, vSphere HA is configured on the ESXi cluster of the PowerStore X model.

6 Conclusion

PowerStore was designed for the data era. This game-changing platform unlocks the power of data, regardless of its structure or location, helping you adapt and transform your IT environment using current skill sets and without disrupting current operations.

The PowerStore platform delivers data-centric, intelligent, and adaptable infrastructure that supports both traditional and modern workloads. PowerStore provides intelligent automation through programmable, autonomous infrastructure that simplifies management and optimizes system resources. It also enables proactive health analytics to monitor, analyze, and troubleshoot the environment. PowerStore has an adaptable architecture that enables speed and application mobility, and offers flexible deployment models. It also provides choice, predictability, and investment protection through flexible payment solutions and data-in-place upgrades.

PowerStore sets a new standard for storage by delivering compelling simplicity and enterprise features at an affordable price and compact footprint. It meets the needs of resource-constrained IT professionals in both large and small organizations. The purpose-built PowerStore system is offered in PowerStore T and PowerStore X models. With its unified storage and clustering capabilities, the PowerStore solution can scale up and out to satisfy ever-changing requirements. With AppsON functionality, plus high-performance block storage, PowerStore provides flexibility for differing use cases and budgets.

PowerStore is designed with ease-of-use at the forefront. The modern design of the management interfaces is built with best practices in mind. This design makes it easy to provision storage intelligently without having to micromanage every detail. A software feature set that is built with the same mindset allows for automation and simplified upkeep. A strong support ecosystem offers various media for learning and troubleshooting, backed by the quality support model of Dell Technologies. Lastly, users looking to refresh their existing Dell Technologies infrastructure can use the intuitive native-migration capabilities of the PowerStore platform.

The ability to accomplish all these goals with a single 2U appliance provides organizations with the flexibility and agility required to address the challenges of today and tomorrow.

A Technical support and resources

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

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

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