

Dell EMC PowerStore: Introduction to the Platform

Abstract

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

January 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 / Metro node updates

Acknowledgments

Author: Ethan Stokes

Updated: Robert Weilhammer

Acknowledgments

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

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

Copyright © 2020 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. [1/18/2021] [Technical White Paper] [H18149]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	4
Executive summary.....	6
Audience	6
1 Introduction.....	7
1.1 Terminology	8
2 PowerStore overview	10
3 Hardware overview.....	12
3.1 High availability.....	12
3.2 Base enclosure.....	13
3.2.1 Drive model comparison.....	14
3.2.2 Node	16
3.2.3 Internal M.2 boot module.....	16
3.2.4 Fan modules.....	17
3.2.5 Battery backup unit.....	17
3.2.6 DIMMs	18
3.2.7 Embedded module	18
3.2.8 I/O module	19
3.2.9 Power supply	20
3.3 Expansion enclosure	22
3.3.1 SAS SSD	22
4 PowerStore T models.....	23
4.1 Deployment.....	23
4.1.1 Networking and cabling	23
4.1.2 Discovery	25
4.1.3 Initial Configuration.....	26
4.2 Deployment modes.....	27
4.2.1 Unified.....	27
4.2.2 Block optimized.....	27
5 PowerStore X models.....	28
5.1 Deployment.....	29
5.1.1 Networking and cabling	29
5.1.2 Discovery	30

5.1.3 Initial Configuration31

6 Conclusion33

A Technical support and resources34

Executive summary

Dell EMC™ PowerStore™ is designed with a purpose-built, 2U, two-node Intel® Xeon® platform. This platform is used for all PowerStore systems, including PowerStore 1000X/T, 3000X/T, 5000X/T, 7000X/T, and 9000X/T models. 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. Also, this document includes cabling and deployment guidance for both PowerStore T and PowerStore X model appliances.

Audience

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

1 Introduction

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 are the following.

Active/active architecture: PowerStore storage 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.

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 vSphere management environment and server resources. This ability allows users to bring applications closer to storage by running applications as virtual machines running directly on PowerStore. AppsON enables agility for application deployments, and allows seamless movement between the PowerStore appliances and VMware ESXi servers. It also helps shrink the stack by eliminating server and networking footprint for space-efficient edge and remote deployments.

VMware integration: PowerStore is designed to have deep integration with VMware vSphere. Integrations include VAAI and VASA support, event notifications, snapshot management, vVols, and virtual machine discovery and monitoring in PowerStore Manager.

Unified offering: PowerStore has a single architecture for block, file, and VMware® vSphere® Virtual Volumes™ (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 to 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 offload. Data reduction is always enabled and intelligently controlled by the storage system.

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: Multiple appliances in a single grouping. Clusters can consist of one appliance or more. Up to four PowerStore T appliances can be clustered by simply adding additional appliances as required.

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.

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): Network file-sharing protocol, also known as CIFS, used by Microsoft® Windows® environments. SMB is used to provide access to files and folders to Windows hosts on a network.

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): VMware API that allows storage-related tasks to be offloaded to the storage system.

vSphere API for Storage Awareness (VASA): VMware API that provides additional insight about the storage capabilities in vSphere.

Virtual Volumes (vVols): VMware storage framework which allows VM data to be stored on individual Virtual Volumes. This ability allows data services to be applied at a VM-granularity level while using Storage Policy Based Management (SPBM).

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 also offers predictive analytics to easily monitor, analyze, and troubleshoot the environment. PowerStore is highly adaptable, providing the flexibility to host specialized workloads directly on the appliance and modernize infrastructure without disruption. It also offers investment protection through flexible payment solutions and data-in-place upgrades.



Figure 1 PowerStore appliance

The PowerStore platform is available in two different product models: PowerStore T and the PowerStore X models. 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. Beyond the power of a single PowerStore T model appliance, multiple PowerStore T model appliances can be grouped in a cluster. A PowerStore T model cluster can consist of a single-appliance or scale up to four PowerStore T model appliances in a single cluster. A PowerStore cluster enables scaling the compute, storage, and connectivity of the PowerStore solution while managing multiple appliances from a single control plane. A cluster can migrate resources between appliances and intelligently load balance new applications based on storage metrics.

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.

Metro node is an external hardware/software add-on feature based on the Dell PowerEdge R640 platform for PowerStore for which it provides active/active synchronous replication as well as standard local use cases.

Additionally, 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 ALUA array based on SCSI response data and therefore 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 additional information about Metro node, go to [Dell Support](#) and reference the Metro node best practices white paper.

3 Hardware overview

The purpose-built PowerStore system is offered in multiple physical hardware models in both PowerStore T model and PowerStore X model appliances. The platform starts with the PowerStore 1000T or 1000X model and scales up to the PowerStore 9000T or 9000X model. The letter **T** or **X** at the end of the model number indicates whether that 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 together when hardware is referenced. See Table 1 for model comparisons.

Also, the system limits change depending on the PowerStore model. More information about system limits can be found on the [PowerStore Info Hub](#).

Table 1 PowerStore model comparison

	PowerStore 1000T/X	PowerStore 3000T/X	PowerStore 5000T/X	PowerStore 7000T/X	PowerStore 9000T/X
Intel CPU (per appliance)	4 x 8 core @ 1.8 GHz	4 x 12 core @ 2.1 GHz	4 x 16 core @ 2.1 GHz	4 x 20 core @ 2.4 GHz	4 x 28 core @ 2.1 GHz
Memory (per appliance)	384 GB	768 GB	1,152 GB	1,536 GB	2,560 GB
NVRAM drives	2		4		
Maximum storage drives (per appliance)	96				
Supported drive types	NVMe SCM, NVMe SSD, SAS SSD ¹				
4-Port Card	25/10/1 GbE optical/SFP+ and Twinax or 10/1 GbE BASE-T				
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	2.5-inch 25-drive SAS SSD				

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 in the environment, such as network or power outages. If an individual component fails, the storage system remains online and continues

¹ SAS SSDs only supported in expansion enclosure

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 25 x 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* on Dell.com/StorageResources.

3.2 Base enclosure

The PowerStore base enclosure supports 25 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 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 once it has been fully inserted in the rack. If the enclosure must be removed from the rails, a bottom latch on each side of the base enclosure must be lifted to pull the base enclosure out. While the base enclosure securely latches onto the rails, there are optional screws underneath each latch which can be tightened for additional security.

The front of the base enclosure contains an LED to display different states of the system. This LED is located in the upper left 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. A minimum of six storage drives are required for PowerStore. 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*.

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 also be used for 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* on the [PowerStore Info Hub](#).



Figure 3 Base enclosure back view

3.2.1 Drive model comparison

Multiple drive types and capacity points are supported on the PowerStore system. 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. Expansion enclosures can only be added to systems which have fully occupied all 21 slots supporting NVMe SSDs in the base enclosure. Expansion enclosure are not supported on systems with NVMe SCM drives or systems which have not been fully populated with 21 NVMe SSDs in the base enclosure. A list of all supported drives can be found on [Dell Support](#).

SAS flash, NVMe flash, and NVMe SCM are considered storage drives in PowerStore, and are formatted with a 512-byte block size. Slots 0 through 20 can be populated with only NVMe SSD or only NVMe SCM drives. NVMe SSD and NVMe SCM drives cannot be mixed 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. Slots 21 through 24 are reserved for NVMe NVRAM drives which serve as additional 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.

3.2.1.1 NVMe SSD

NVMe solid-state drives (SSDs) are dual-ported, high performance, nonvolatile flash drives that are accessed by the PowerStore operating system 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 as follows:

- 1,920 GB
- 3,840 GB
- 7,680 GB
- 15,360 GB

NVMe SSDs are supported in base enclosure slots 0 through 20. NVMe SSDs cannot be mixed with NVMe SCM drives. However, once all free base enclosure slots have been filled with NVMe SSD drives, expansion enclosures with SAS SSDs can be added.

3.2.1.2 NVMe SCM

NVMe storage class media (SCM) drives are dual-ported, extreme-high-performance, nonvolatile drives designed with Intel Optane™ technology. NVMe SCM drives have lower latency and improved performance compared to other SSD drives and are also accessed by the PowerStore operating system 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 two capacity points:

- 375 GB
- 750 GB

NVMe SCM drives are supported in base enclosure slots 0 through 20. NVMe SCM drives cannot be mixed with NVMe SSDs. A PowerStore system with NVMe SCM drives does not support any expansion enclosures.

3.2.1.3 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. There are two or four drives per appliance, depending on the PowerStore model (see Table 3).

Table 3 NVMe NVRAM count by PowerStore model

PowerStore model	NVMe NVRAM count
PowerStore 1000T/X	2
PowerStore 3000T/X	
PowerStore 5000T/X	4
PowerStore 7000T/X	
PowerStore 9000T/X	

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 cannot be modified. Extra NVMe NVRAM drives cannot be added later.

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

3.2.2 Node

The purpose-built PowerStore platform is powered by dual-socket Intel® Xeon® processors. The core count varies between 8 and 28 cores for each of the two processors per node depending on the PowerStore model. 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 mid-plane 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

3.2.3 Internal M.2 boot module

A primary and a secondary M.2 SATA device are located inside each node 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. It is used to store the base operating system and log files, and is used for general system operations. The secondary M.2 device is 120 GB. It is used as a recovery device during a primary M.2 failure and is an alternate location for log files.



Figure 4 Internal M.2 boot module

3.2.4 Fan modules

Fan modules (cooling modules) are used 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.

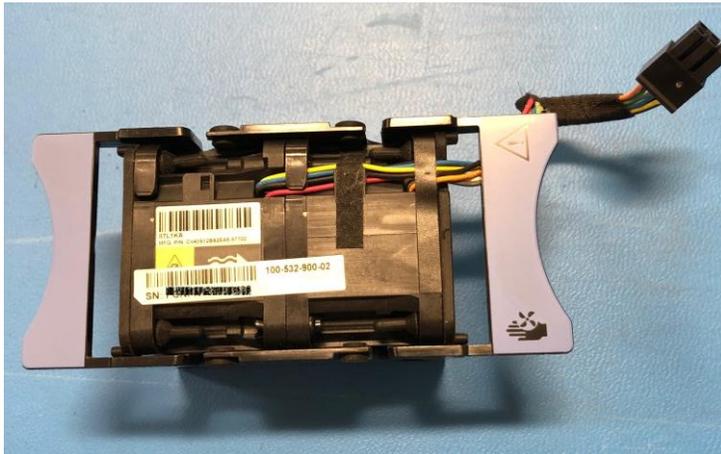


Figure 5 Fan module

3.2.5 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 each NVRAM mirrored pair is powered by both BBUs, 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

3.2.6 DIMMs

Each node contains 24 DDR4 DIMM slots, which are populated in different configurations that are based on the PowerStore model. DIMMs (Figure 7) are between 16 GB and 64 GB in size depending on the configuration. Table 4 outlines the model, memory total, and DIMM configurations per node. 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.

Table 4 Node DIMM configuration

Model	Node memory	DIMM configuration
PowerStore 1000T/X	192 GB	12 x 16 GB
PowerStore 3000T/X	384 GB	12 x 32 GB
PowerStore 5000T/X	576 GB	12 x 16 GB + 12 x 32 GB
PowerStore 7000T/X	768 GB	24 x 32 GB
PowerStore 9000T/X	1,280 GB	20 x 64 GB



Figure 7 DIMM

3.2.7 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, 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. The 4-port card configuration is selected at the time of ordering and cannot be changed later. Ports 0 and 1 must be configured with a link speed of at least 10 GbE on the 4-port card. This ensures that inter-cluster 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

3.2.8 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) (PowerStore T models only)
- 32/16/8/4 Gb Fibre Channel (4-port)
- 10/1 GbE BASE-T (4-port) (PowerStore T models only)

The 25 GbE optical I/O module (Figure 9) supports SFPs running at either 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. Different SFPs or Twinax cables can be mixed on the same I/O module and are hot-swappable. This I/O module supports iSCSI traffic.

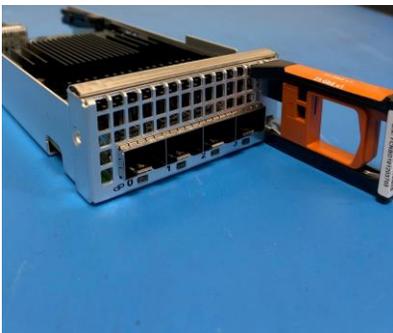


Figure 9 25 GbE optical I/O module

The 32 Gb Fibre Channel (4-port) I/O module (Figure 10) 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.

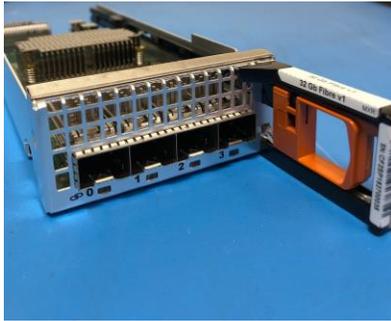


Figure 10 32 Gb Fibre Channel I/O module

The 10 GbE BASE-T (4-port) I/O module (Figure 11) 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 11 10 GbE BASE-T I/O module

3.2.9 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. For PowerStore 1000T or X, PowerStore 3000T or X, and PowerStore 5000T, both the 1800W high line and 2100W high line PSUs are supported. For environments which only provide low line power, the 2100W PSU can be used with a step-up transformer. The 1800 W should not be selected for environments with low line power. For PowerStore 5000X, PowerStore 7000T or X, and PowerStore 9000T or X, only the 2100W PSU is supported. For environments which only offer low line power, the 2100W PSU can be used with a step-up transformer. Table 5 shows the PSU specifications per PowerStore model.

Table 5 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

¹Supports high-line power only

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

Figure 12 shows the supported connector types.



Figure 12 C19/C20, C13/C20, C13/C14 connectors

A single power supply (Figure 13) 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 13 2100 W power supply

3.3 Expansion enclosure

The purpose-built PowerStore system supports a 25-drive 2U expansion enclosure (Figure 14) using 2.5-inch SAS drives for extra capacity. After the free storage drive slots in the base enclosure have been entirely filled with NVMe SSDs, an expansion enclosure can be added. A fully populated system requires 21 NVMe SSDs in slots 0 through 20 of the base enclosure.

Expansion enclosures are not supported with PowerStore systems using NVMe SCM drives. 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 14 25-drive 2.5-inch 2U expansion enclosure (front)

The back of the expansion enclosure (Figure 15) includes LEDs to indicate power and fault status. There are also LEDs to indicate bus and enclosure IDs.



Figure 15 25-drive 2.5-inch 2U expansion enclosure (back)

3.3.1 SAS SSD

SAS SSDs are dual-ported, high-performance, nonvolatile flash drives that are accessed by the PowerStore operating system through the SAS protocol. SAS SSDs are only supported in expansion enclosures and can only be added to a PowerStore system after the base enclosure has been filled with drives.

SAS SSDs supplement the storage tier for PowerStore, are used for user data or metadata, and are available in the following capacity points:

- 1,920 GB
- 3,840 GB
- 7,680 GB

Expansion enclosures and SAS SSDs are only supported on systems with NVMe SSDs, and they cannot be mixed with systems using NVMe SCM drives.

4 PowerStore T models

PowerStore T model appliances include one of two model types within the PowerStore platform. PowerStore T models have the 2U, two-node 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 or 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. Providing storage in multiple formats to applications delivers workload flexibility and enables administrators to simplify and consolidate infrastructure.

4.1 Deployment

Once the system has been racked, a single PowerStore T model appliance requires simple network cabling to the physical Ethernet and management switches. After completing the cabling, ensure that the proper switch configuration has been applied to support management and data traffic from the appliance. Finally, connect the power cables and power on the appliance. Once online, PowerStore can be easily accessed to 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 a similar wizard. 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* on the [PowerStore Info Hub](#). The *Quick Start Guide* contains an overview of PowerStore deployment and directs readers to all other 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. The Ethernet switches ensure high availability for iSCSI, NAS, replication, external storage import, data migration, and inter-cluster traffic. Properly configuring and cabling to the physical Ethernet switches ensures that features 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

PowerStore T model appliances also require at least one management switch.

The minimum cabling requirements in Table 6 must be met.

Table 6 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

Figure 16 outlines the minimum cabling that is required for PowerStore T appliances.

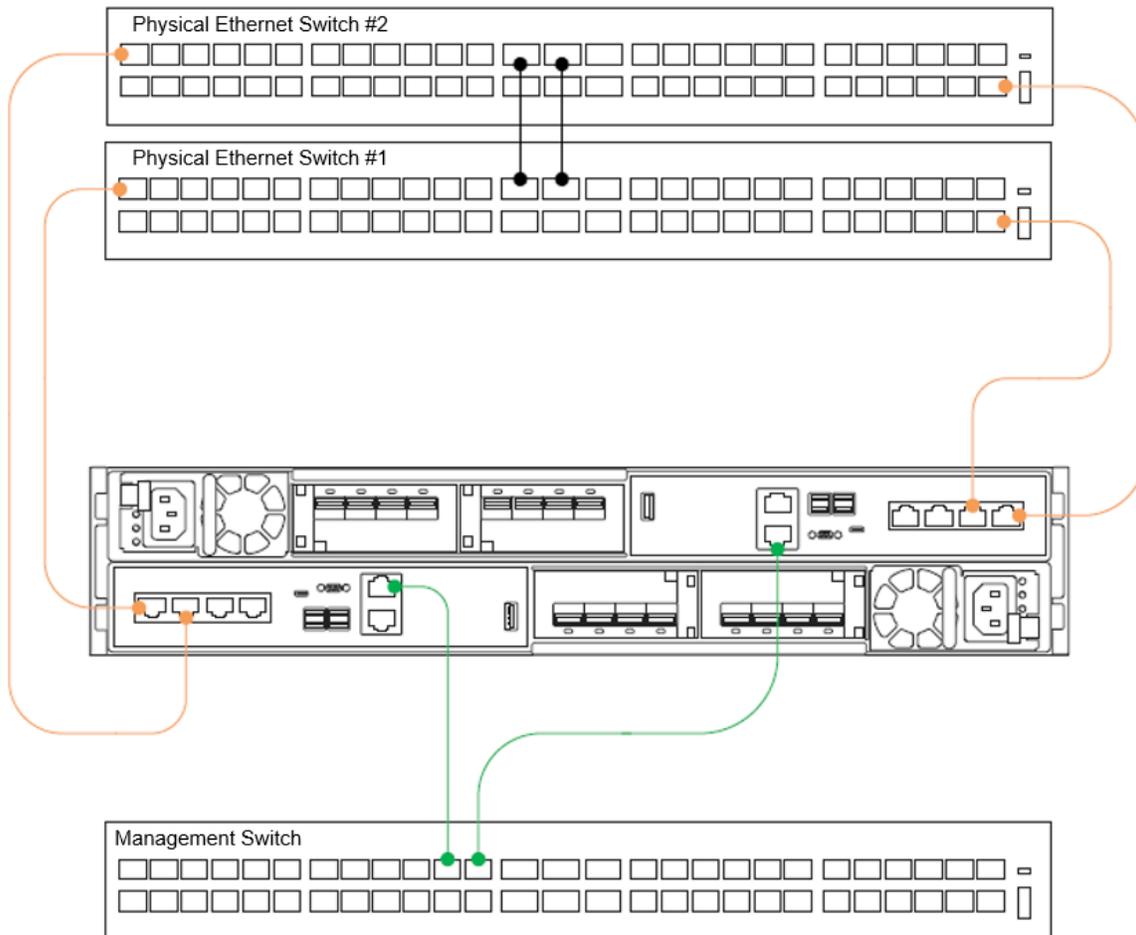


Figure 16 PowerStore T model appliance minimum cabling

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 documents *Dell EMC PowerStore: Network Configuration for Dell PowerSwitch Series Guide* and *Dell EMC PowerStore: Network Planning Guide* on the [PowerStore Info Hub](#).

PowerStore T model appliances contain a system bond by default on the first two ports of the embedded module 4-port card (Figure 17). This bond provides high availability and potentially increased throughput for iSCSI, NAS, replication, external storage import, data migration, and inter-cluster 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.

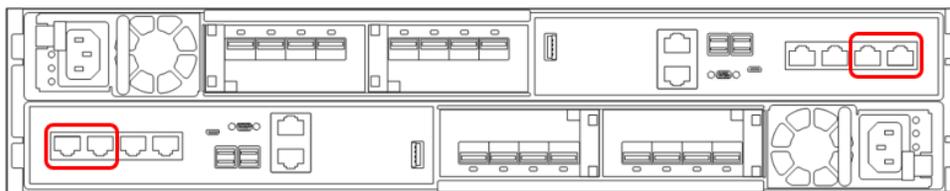


Figure 17 PowerStore T model system bonds

4.1.2 Discovery

After the PowerStore system is cabled and the appropriate network settings are configured, perform discovery and initial configuration for PowerStore. Three discovery options are available, either physically connecting a workstation to the appliance, navigating to a static discovery IP, or discovering the appliance remotely using the **PowerStore Discovery Tool** application. Once the PowerStore system is discovered, step through the HTML5-based Initial Configuration Wizard to complete 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. Users can log in to PowerStore and begin the Initial Configuration Wizard by directing a browser to 128.221.1.250.

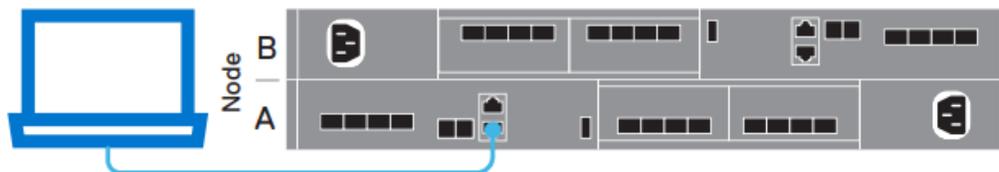


Figure 18 PowerStore Direct Connect Discovery

4.1.2.2 Static IP

Another discovery method is available with the PowerStoreOS 1.0.3 release, through a static IP address. The user can access a range of IP addresses to discover their PowerStore and go through the initial configuration. The users can use any of the IP addresses that are 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>

In addition, a user's workstation (also laptop or virtual machine) needs to be connected through an Ethernet cable to the same network (either same switch or same untagged/native VLAN). Then, the user's workstation needs to have a static IP set in the same 169.254.0.x/16 subnet (for example 169.254.0.100).

Finally, the user can open a browser with one of the IP addresses mentioned 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 directly connecting a workstation to the PowerStore appliance or Static IP are 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 Avahi broadcasts made by the PowerStore system. 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. The storage network can be configured (optionally), which allows support for iSCSI, data import, and replication traffic when the system is configured. If the storage network configuration is skipped during the Initial Configuration Wizard, it can be configured at any point through the PowerStore Manager. Each network requires several IP addresses, along with gateway, netmask, and VLAN information. If the network is using 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 7 details the specific IP requirements for each network.

Table 7 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	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 *Dell EMC PowerStore: Initial Configuration Worksheet* on the [PowerStore Info Hub](#).

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, users 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.

The deployment mode is selected during the Initial Configuration Wizard and the system is brought online after the configuration is successfully applied. Deployment modes cannot be changed non-disruptively 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.

4.2.2 Block optimized

The block optimized deployment mode for PowerStore T models is the alternative deployment mode. When **Block Optimized** is selected during the Initial Configuration Wizard, a prompt is shown that confirms the selection. Then, the appliance reboots and comes online in the block optimized deployment mode. The Initial Configuration Wizard is then completed 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 for PowerStore. PowerStore X model appliances are built on the 2U, two-node 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 X models have a native hypervisor layer which passes through underlying hardware components directly to a virtualized instance of the PowerStore operating system. 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 or Fibre Channel), vVol storage (with iSCSI or Fibre Channel), and 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 controller VM has 50% of the node CPU and memory that is reserved by the ESXi host. The remaining 50% of the node resources is used by user VMs. These VMs are protected with high availability (HA) by default, and the PowerStore controller VMs are also redundant.

Figure 19 depicts a PowerStore X model appliance architecture. 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.

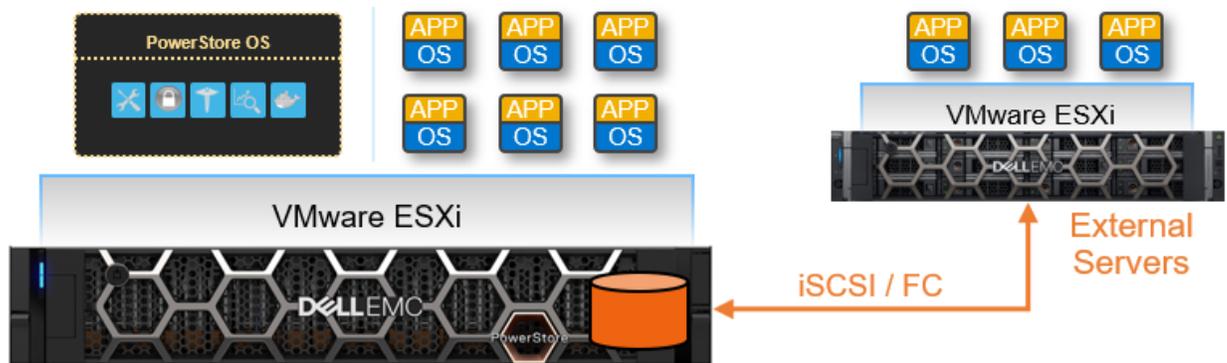


Figure 19 PowerStore X model architecture

For more details about the PowerStore X model components and integration with VMware, see the document *Dell EMC PowerStore: Virtualization Integration* on [Dell.com/StorageResources](https://www.dell.com/storage/resources).

5.1 Deployment

Once the system has been racked, a single PowerStore X model appliance requires simple network cabling to the physical Ethernet switches. After completing the cabling, ensure that the proper switch configuration has been applied to support management and data traffic from the appliance. Finally, connect the power cables and power on the appliance. Once online, PowerStore can be easily accessed to 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 a similar wizard. 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.

For detail on PowerStore deployment and configuration in general, see the *Dell EMC PowerStore: Quick Start Guide* on the [PowerStore Info Hub](#). The *Quick Start 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

The minimum cabling requirements in Table 10 must be met.

Table 8 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 20 outlines the minimum cabling that is required for PowerStore X model appliances.

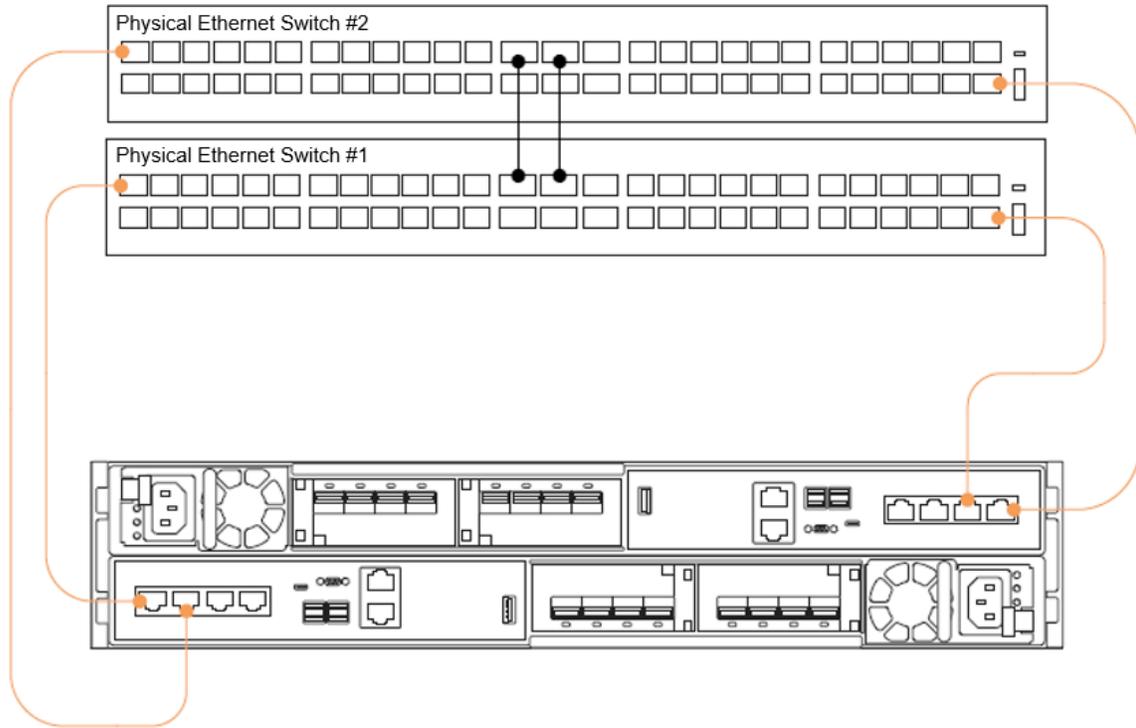


Figure 20 PowerStore X model minimum cabling

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 documents *Dell EMC PowerStore: Network Configuration for Dell PowerSwitch Series Guide* and *Dell EMC PowerStore: Network Planning Guide* on the [PowerStore Info Hub](#).

5.1.2 Discovery

After the PowerStore system is cabled and the appropriate network settings are configured, perform discovery and initial configuration for PowerStore. Three discovery options are available, either physically connecting a workstation to the appliance, navigating to a static discovery IP, or discovering the appliance remotely using the **PowerStore Discovery Tool** application. Once 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 is set on the workstation, and users can log in to PowerStore and begin the Initial Configuration Wizard by directing a browser to 128.221.1.250.

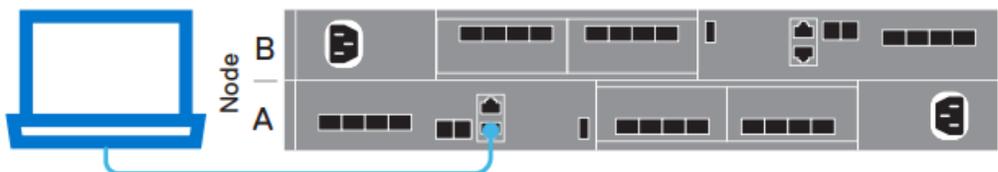


Figure 21 PowerStore Direct Connect Discovery

5.1.2.2 Static IP

Another discovery method is available with the PowerStoreOS 1.0.3 release, through a static IP address. The user can access a range of IP addresses to discover their PowerStore and go through the initial configuration. The users can use any of the IP addresses that are 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

In addition, a user's workstation (i.e. laptop or virtual machine) needs to be connected through an Ethernet cable to the same network (either same switch or same untagged/native VLAN). Then, the user's workstation needs to have a static IP set in the same 169.254.0.x/16 subnet (e.g. 169.254.0.100).

Finally, the user can open a browser with one of the IP addresses mentioned 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 a directly connecting a workstation to the PowerStore appliance or static IP are 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 Avahi broadcasts made by the PowerStore system. The Avahi broadcasts are sent from the PowerStore X 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 then 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* on the [PowerStore Info Hub](#).

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 9 details the specific IP requirements for each network.

Table 9 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 IPs, 4 x for best performance • 4 x iSCSI initiator IPs • Global Storage Discovery IP (optional)
vMotion	2	<ul style="list-style-type: none"> • 2 x vMotion IPs

Besides the networking information, the following details are required to complete the Initial Configuration Wizard:

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

The storage network provides iSCSI targets to support external iSCSI, data import, and replication traffic when the system is configured. The storage network IP addresses are also used for 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. Configuring an additional two iSCSI targets allows the ESXi hosts to have additional paths to the controller VM and improve performance.

For more information about the Initial Configuration Wizard, see the document *Dell EMC PowerStore: Initial Configuration Worksheet* on the [PowerStore Info Hub](#).

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 one 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

The PowerStore platform delivers data-centric, intelligent, and adaptable infrastructure that supports both traditional and modern workloads. A data-centric design optimizes system performance, scalability, and storage efficiency to support any workload without compromise. 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 the unified storage and clustering capabilities of PowerStore T models, the PowerStore solution can scale up and out to satisfy ever-changing requirements. Clustering enables adding an appliance to an existing cluster and migrating resources between them. 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.

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.