

Dell EMC PowerStore Release Notes

Release Version: 2.0.0.0 (Build: 1397847)

Release Type: Major (MA)

Contents

Revision History.....	1
Product Description	2
New Features	3
Changed Features.....	13
Resolved Issues.....	13
Known Issues	23
Limitations	38
Environment and System Requirements	40
Installation and Upgrade Considerations.....	40
Where to Get Help.....	42

Revision History

Revision	Date	Comments
A00	June 2021	Original release version 2.0.0.0 (Build: 1376722)
A01	June 2021	Updated build release version 2.0.0.0 (Build 1397847). The following issues were resolved in this release: MDT-294054, MDT-305819, MDT-303583, MDT-305868, and MDT-305874. Refer to the Resolved Issues section for details.
A02	July 2021	Updated resolved issue MDT-237899, and the workaround description for Known Issue MDT-280394.
A03	August 2021	Updated the link to the PowerStore Release Notes – Minor Issues knowledge base article Number: 000155626 .
A04	September 2021	Updated incorrect knowledge base article references and removed links to all knowledge base articles other than the link to the PowerStore Release Notes – Minor Issues knowledge base article Number: 000155626 .

Product Description

The ground-breaking Dell EMC PowerStore achieves new levels of operational simplicity and agility, utilizing a container-based architecture, advanced storage technologies, and intelligent automation to unlock the power of your data. Based on a scale-out architecture and hardware-accelerated advanced data reduction, PowerStore is designed to deliver enhanced resource utilization and performance that keeps pace with application and system growth. Utilizing the proven capabilities of VMware ESXi, PowerStore X models with AppsON provide the unique ability to host data-intensive and storage applications directly on the PowerStore system with a storage-based virtualization environment, with the flexibility of seamless movement of applications between the storage system and external VMware servers. PowerStore T models provide organizations with all the benefits of an enterprise unified storage platform for block, file and vVol data, while enabling flexible growth with the intelligent scale-up AND scale-out capability of appliance clusters.

Highlighted features include:

- Brand new built from scratch all NVMe Platform:
 - All NVMe
 - Active-Active architecture
 - Container based PowerStore OS
 - Block, File, and vVols
 - NVMe Flash and Storage Class Memory (SCM) media support
 - Front End Connectivity:
 - FC: 32 Gb NVMe-FC, 32/16/8Gb FC
 - Ethernet: 25/10/1 GbE iSCSI & File
 - Compact 2U starting form factor
 - Enterprise Data Services & Leading Data Reduction:
 - Inline dedupe and compression
 - Native async replication
 - Snapshots and space efficient thin clones
- Enterprise Data Services & Data Reduction:
 - Inline dedupe and compression
 - Native async replication
 - Snapshots and space efficient thin clones
 - Advanced drive failure protection and sparing technology
- Simple and intelligent Management & Serviceability:
 - Embedded management
 - Built in AI for simple, autonomous storage administration and proactive health analytics
 - CloudIQ
 - VM Visibility
 - New Anytime Upgrade Program
 - Integration with automation framework
- Flexible & Granular Scalability:
 - Scale Up in single drive increments, up to 2.8 PBe per appliance
 - Scale Out to 4 appliances, up to 11.3 PBe per cluster
 - Scale Down
- VMware:
 - VMware vVols 2.0/VASA 3.0 support

- New AppsOn hypervisor-based deployment, allowing to run Virtual Machines on the same appliance as storage without the need for an external server

PowerStore deployment models and configurations

Deployment Model	Model numbers	Support Configurations
PowerStore T	500T 1000T 3000T 5000T 7000T 9000T	Storage-centric deployments. Both block (storage area networks (SAN)) and file (network attached storage (NAS)) services are supported with PowerStore T deployments. Note: Hypervisor deployments are not supported on this model.
PowerStore X	1000X 3000X 5000X 7000X 9000X	Application and storage-centric deployments. Block (SAN-only) storage services with a hypervisor layer installed on the system. The system software is deployed on the hypervisor, which enables deployment of virtual machines (VMs) and customer applications within the hardware. Note: File (network attached storage (NAS)) services are not supported on this model.

Refer to the PowerStore product support pages (<https://www.dell.com/support>) for detailed user documentation.

New Features

The following features have been introduced in this release.

Functional Area	Feature Description	Summary of Benefits
Connectivity-Networks	Support for up to 32 Storage networks can be added to your PowerStore deployment.	Better scalability for network traffic across ports. High Availability Configure more than one storage networks.
Connectivity-Networks	PowerStore T model appliances can be deployed without Top-of-Rack switches when Network Attached Storage, iSCSI Storage networks, clustering, and asynchronous replication are not required.	Lower cost of deployment when Network Attached Storage, iSCSI Storage networks, and asynchronous replication are not required. Additional Top-of-Rack switch can be added later to enable additional storage functionality.

Functional Area	Feature Description	Summary of Benefits
Connectivity-Networks	I/O modules can be added to deployed PowerStore model clusters to expand the amount of usable front-end ports. Adding I/O modules post deployment must be performed by your service provider.	More bandwidth for better network performance. Additional ports for expanding storage networks.
Data Collection Serviceability	This feature provides the ability to relay audit logs to the remote server for monitoring PowerStore. Also, it adds TLS support which provides security and encryption for the logs sent.	PowerStore can send audit log messages to two remote servers. Log message transfers can use a one-way authentication (Server CA Certificates) or an optional two-way authentication (Mutual Authentication Certificate).
Data Protection	Enables the user to initiate a disaster recovery test that simulates a failover operation. The user can test a target PIT to validate the replicated snapshot.	Confirms application recoverability for business continuity on the destination system without impacting production or ongoing replication.
Data Protection	Adds the ability to access the PowerStore Metro Node from PowerStore Manager.	Metro node enables synchronous active-active replication over metro distances with multi-site dual access.
Data Protection	This feature adds Double Drive failure protection support to PowerStore. During the initial configuration, the user can choose Double Drive failure protection or Single Drive failure protection. Once set, this selection cannot be changed. Note: The minimum drive count for Double Drive failure protection is 7.	Double Drive failure protection protects the appliance from having a Data Unavailable (DU) or Data Loss (DL) situation if two drives fail simultaneously.
Data Protection	The user can set/modify a time zone for the snapshot scheduler, so that the snapshots will be taken at the present time, regardless of daylight-saving time changes.	Prevents changes in time of snapshot creation following daylight saving time changes.
Hardware	When you use PowerStore Manager to power off a single appliance in a multi-appliance cluster, the system performs a series of validations prior to powering off the appliance. Based on the results of the validations, the system provides warnings and recommendations to minimize potential disruptions to operations.	Simplifies the process of powering off a single appliance in a multi-appliance cluster.
Hardware	Automated Intelligent Metadata Tiering provides improved system performance by ensuring that the metadata is prioritized over data and is always placed on the highest available tier.	Prioritizing metadata over user data on the highest available tier offers the best \$/IOPS to PowerStore customers. The overall system performance can be improved by adding as low as one high performance drive, such as NVMe SCM, even if the other drives in the appliance have different performance profiles.

Functional Area	Feature Description	Summary of Benefits
Hardware	<p>The PowerStore 500T model is introduced with this release and serves block and file services, and the software stack is deployed directly on the system. The PowerStore 500T hardware consists of a 2U, two node storage solution. Between the front and rear of the enclosure, a mid-plane distributes power and signals to all the enclosure components. On the front of the base enclosure, drives connect to the mid-plane. On the rear of the base enclosure, the nodes and power supply modules connect to the mid-plane. The I/O modules connect directly to the node. Each node contains an internal battery backup module, redundant fan modules, DDR4 memory, and one Intel Cascade Lake processor.</p> <p>Additionally, PowerStore 500T model appliances also offer dedicated low line power (less than 220V) in a native fashion. As a result, low line customers purchasing a 500T will no longer be required to additionally purchase step up transformers (SUTs) in order to operate the PowerStore 500T at low line voltages.</p>	<p>Provides a cost-optimized appliance with a single socket CPU and onboard vaulting.</p> <p>Customers can operate at low line power will be the capital expenditures savings historically required when purchasing step up transformers (SUTs) when native low line support was not supported.</p>
Hardware	Added support for 1.5 TB NVMe SCM drives.	Provides an additional NVMe SCM drive capacity option.
Import	The non-disruptive import feature supports additional combinations of host OS, host multipath software, and host to storage protocols. This enhancement is based on the support for ESXi 6.7 Update 3 and SUSE Linux Enterprise Server 15 for host OS. For the most up-to-date supported versions of software and systems, see the Support Matrix for PowerStore.	This enhancement provides the ability to migrate applications running on ESXi 6.7 Update 3 or SUSE Linux Enterprise Server 15 to PowerStore.
Installation and Upgrade	An appliance that is part of a PowerStore cluster can be removed from the cluster. When an appliance is removed from a cluster a factory reset is performed on the appliance. The factory reset leaves the appliance in an unconfigured state and in a discoverable mode.	When an appliance is removed from a cluster, the appliance can be used in a new cluster or the appliance can be added to another existing cluster.
Internal Migration	Allows you to non-disruptively move storage resources to another appliance in the cluster. This feature is applicable if you want to evacuate space from an appliance, power off an appliance, or remove an appliance from a cluster.	Provides the ability to non-disruptively move storage resources to another appliance in the cluster. Allows you to select multiple volumes for migration and provides Resource Balancer recommendations for target appliances.

Functional Area	Feature Description	Summary of Benefits
Monitoring	Provides the ability to measure the performance of iSCSI protocols per IP interface.	Enables the user to measure iSCSI performance per IP interface as part of the ongoing management or troubleshooting.
Monitoring	Provides the ability to identify CPU utilization per node, Provides the ability to identify Volume/VM/vVol affinity.	Enables the user to track CPU utilization per node and balance it manually, managing the volumes affinity if required.
PowerStore CLI	This feature enables you to filter rows of output in the command-line interface using more than two dozen operators along with specific arguments.	This feature allows you to filter the data to see only the relevant data that you want to see.
PowerStore Manager (GUI)	This feature allows for the ability to run an independent health check of the system using PowerStore Manager>Monitoring>System Checks. The health check provides information on hardware, customer objects and system internals.	Customers can run a health check to learn about the cluster health in PowerStore Manager. The health check results may be used to resolve issues prior to performing actions such as system upgrade.
PowerStore Manager (GUI)	<p>A banner display at the top of PowerStore Manager notifies users when something significant has occurred because of a degraded system or an on-going operation that requires immediate attention. Some alerts displayed in the banner are triggers to begin corrective steps when an issue has occurred on the cluster or the cluster appliances, while other messages alert users to system conditions which they may need to be made aware of.</p> <p>The banner colors correspond to the alert severity. Multiple colors may be displayed simultaneously. The different colors allow users to easily assess the situation. The highest severity alert dictates the banner color when there are concurrent alerts being presented.</p>	Immediately alerts users that are currently logged in to PowerStore of an event that may impact their environment.
PowerStore Manager (GUI)	Provides direct mapping of Volume Group collectively for all included volumes to a respective Host/Host Group.	Provides the ability for users to map a Volume Group to a Host/Host Group.

Functional Area	Feature Description	Summary of Benefits
PowerStore Manager (GUI)	PowerStore Manager displays a Cluster Time for the PowerStore cluster on the Cluster Properties page under Settings. The time that is displayed in PowerStore Manager is the system time of the primary node on the primary appliance. The times on the nodes of a cluster are synchronized using a Network Protocol Time (NTP) server by default. However, if the nodes cannot communicate with an NTP server, an alert is generated and the times on the nodes are synchronized with the primary node on the primary appliance.	Customers can view the Cluster Time for a PowerStore cluster in PowerStore Manager.
PowerStore Manager (GUI)	Supports a new role called Storage Operator for operating PowerStore. The new storage user role allows for control of common day-to-day tasks.	Provides a role for simple provisioning and managing within the PowerStore Manager.
PowerStore REST API	This feature allows you to use REST API to initialize PowerStore and add additional systems to the cluster.	You can leverage the API interface to efficiently initialize PowerStore and add additional systems to the cluster.
PowerStore REST API	PowerStore REST API documentation can now be found on the Dell Technologies Developer Portal at https://developer.dell.com	The Dell Technologies Developer Portal https://developer.dell.com , is a single location for sharing and disseminating developer-focused documentation in an easily consumable and consistent format. Using the Developer Portal, developers can discover, explore, and test Dell APIs needed to integrate with other solutions and automate tasks. This enables the 'build once/re-use often' mindset for faster implementation and easier automation of solutions built using Dell's enterprise products.
Security	The Transport Layer Security (TLS) protocol provides the ability to secure communications across networks. PowerStore supports TLS 1.2 by default but includes the ability to reenable TLS 1.1 as necessary.	Provides the highest level of transport security while providing flexibility for legacy environments.

Functional Area	Feature Description	Summary of Benefits
Security	<p>The PowerStore hardware provides the following security features for firmware images and the operating system through the enclosure management software on the platform.</p> <p>Secure Boot and x86 Secure Boot technologies:</p> <p>The enclosure management software authentication and root of trust provides the capability to authenticate boot loader and firmware, and immutable hardware root of trust.</p> <p>Ensure a verified and measured boot.</p> <p>Authenticate firmware images and operating system boot loader at boot time.</p> <p>Digitally signed firmware upgrades ensure that root of trust authenticates all signed upgrade firmware images.</p>	<p>Provides an additional level of security by implementing a complete security check from the bootloader to the OS.</p>
Serviceability	<p>PowerStore supports the transfer of alert information to designated SNMP Managers that act as repositories for generated alert information by the PowerStore cluster. Both SNMP v2 and SNMP v3 trap destinations are supported.</p>	<p>Helps storage administrators to monitor events and alarms from one place rather than monitoring each system.</p>
SupportAssist	<p>To enable SupportAssist for the first time with this release, the associated Dell EMC End User License Agreement (EULA) must be accepted; otherwise, SupportAssist cannot be enabled. Up to two SupportAssist gateway servers can be supported for high availability. One gateway server is designated as the primary while the other is designated as the backup. Upon a failure of the primary gateway server, all alerts, configuration and metrics reporting is directed to the backup. When the primary gateway server is recovered, the reporting is directed back to the primary gateway server.</p> <p>Support for a Policy Manager has been removed in this release.</p>	<p>The PowerStore Manager user interface for SupportAssist has been enhanced to provide easier management of connectivity.</p>

Functional Area	Feature Description	Summary of Benefits
SupportAssist	The Remote Secure Credentials feature within SupportAssist allows authorized Dell EMC service personnel to authenticate to PowerStore without arranging a password with the system's owner beforehand. This feature is enabled at the cluster level and provides Dell EMC service personnel with temporary access to a PowerStore cluster to respond faster to resolve issues.	Dell EMC Support does not need to contact the customer and obtain the service user or admin user password before they can log in remotely, which can cause a significant delay in resolving any issues. Also, the use of 3rd party tools is not needed.
Storage-Block	In the rare event when corruption occurs within an individual volume and its family (associated snapshots and clones), the impact will be contained within the volume and its family and they will be taken offline. The user can then approach Dell Support to resolve the issue.	Improves system availability by enabling fault containment at the volume family level.
Storage-Block	NVMe over Fabrics (NVMe-oF) extends the NVMe protocol from the drive level to the network fabric with NVMe over Fibre Channel (NVMe-FC) support.	NVMe over Fabrics (NVMe-oF) is much more efficient, parallel and scalable than SCSI.
Storage - Block	This feature expands PowerStore's Intelligent Data Reduction capabilities with enhancements to inline deduplication and background deduplication. Under periods of extreme I/O demand, PowerStore's default inline deduplication is temporarily deferred to reduce CPU utilization, dynamically prioritizing the workload I/O. Also, as with many other PowerStore features, Intelligent Data Reduction works proactively in the background on your behalf with background deduplication.	These enhancements to inline deduplication and PowerStore model's data reduction services give the system up to a 20% IOPS boost with no compromise to the 4:1 guarantee and with the system remaining fully automated. Background deduplication provides the best combination of data reduction, write minimization and performance under any condition.
Storage-File	The Grace Period can be set for a File System Tree Quota that is independent of the User Quota, Grace Period. A Tree Quota no longer inherits the Grace Period set for a User Quota on a File system.	Provides more flexibility for setting Tree Quotas for File Systems being used for different functions or organizations.
Storage-File	New metrics for performance of file storage operations at the appliance and cluster level with an option to break down by protocol and node.	Provide more detail into File storage performance.

Functional Area	Feature Description	Summary of Benefits
Virtualization	<p>The internal nodes of a PowerStore X model cluster are visible as ESXi hosts on the Hosts & Host Groups page in PowerStore Manager, allowing mapping of block volumes to these internal ESXi hosts via REST/CLI if needed.</p>	<p>By default, AppsON virtual machines leverage PowerStore's efficient vVol implementation due to its simple nature, design optimizations, and integration within the PowerStore UI. Although it is still recommended to use vVols due to its simple nature, optimizations, and integration within the PowerStore UI, starting with PowerStore version 2.0, PowerStore X model appliances also support VMFS Datastores for the storage of Virtual Machines within AppsON by allowing the mapping of block volumes to PowerStore's internal ESXi hosts via PowerStore's REST API and/or CLI.</p> <p>For details, please see Article Number 000182913: PowerStore: Using VMFs on PowerStore X model appliance internal nodes.</p> <p>Note: Access to this article is restricted to partners and Dell personnel.</p>
Virtualization	<p>PowerStore X model clusters support up to four PowerStore X model appliances that are running the same version of PowerStore software.</p> <p>PowerStore X model appliances can be added to or removed from a PowerStore X model cluster.</p>	<p>Customers can take advantage of the improved performance and scaling abilities of a PowerStore X model cluster with multiple appliances.</p>
Virtualization	<p>For PowerStore X model clusters, empty affinity rules, by means of VM/Host Groups and VM/Host Rules on PowerStore X model clusters are pre-created in vCenter, in order to allow storage and compute co-location if needed. The system creates the following by default:</p> <p>A Host Group that is named Appliance name. The internal ESXi hosts on the appliance are automatically added to this Host Group.</p> <p>A VM Group that is named Appliance name VMs. The VM Group is empty by default.</p> <p>A VM/Host Rule that is named Appliance name. The VM/Host Rule is a Should run on hosts in group, which specifies that VMs in the VM Group should run on the hosts in the Host Group, but can move over to other nodes in case of an HA event or entering Maintenance Mode.</p>	<p>Customers can use the VM/Host Groups and VM/Host Rules on a PowerStore model X cluster to keep compute and storage co-located on the same appliance for performance benefits.</p>

Functional Area	Feature Description	Summary of Benefits
Virtualization	<p>The PowerStore software upgrade for PowerStore X model clusters includes a required ESXi update for the internal nodes of the clusters. The internal nodes of an existing PowerStore X model cluster must be upgraded to VMware ESXi 6.7 P05, Patch Release ESXi670-202103001 (Build: 17700523) before the cluster can be upgraded to the latest version of the PowerStore OS. In addition to the ESXi update, other PowerStore X model components must be upgraded as well.</p> <p>Download the latest PowerStore X model software upgrade .zip file from https://www.dell.com/support and see the Upgrade Considerations section in this document for more information.</p>	The internal nodes of a PowerStore X model cluster use a newer version of ESXi that contains bug fixes and security fixes.
Virtualization	<p>For PowerStore X model clusters, PowerStore Manager displays the ESXi version that is running on the ESXi host nodes of an appliance, the appliances in a cluster, and the cluster itself.</p> <p>An alert is generated if the ESXi versions do not match on Node A and Node B of a PowerStore X model appliance. An alert is also generated if the ESXi versions do not match on the appliances in a PowerStore X model cluster.</p>	Users are able to view the ESXi version of an internal ESXi host node on a PowerStore X model cluster. Users are also alerted if the ESXi versions installed on nodes or appliances of a PowerStore X model cluster do not match.
Virtualization	PowerStore model clusters support online migration of vVols for vVol-based VMs to another appliance in the cluster. However, due to a known issue with VMware's vSphere, only newer versions of vSphere support this.	Customers can migrate the vVols for a vVol-based VM to another appliance in a cluster without powering off the VM. For information on required vSphere versions to support this feature, please see Dell knowledge base article: 000105896: PowerStore: Supported Host operating system for Non-Disruptive migration of Storage resources

Functional Area	Feature Description	Summary of Benefits
Virtualization	<p>All PowerStore models now support automatic vCenter Server and VASA Provider registration. PowerStore Manager displays the vCenter Server connection status and VASA Provider status on the vCenter Server Connection page.</p> <p>In addition, new VMware vSphere property called "vSphere Host Name" is now displayed in PowerStore Manager:</p> <p>The internal ESXi hosts of a PowerStore X model cluster are displayed as hosts on the Hosts & Host Group page.</p> <p>Hosts added to PowerStore Manager display their associated ESXi host name in the vSphere Host Name column of the Hosts & Host Groups page.</p> <p>The name of the ESXi host that a Virtual Machine is running on is displayed in the vSphere Host Name column of the Virtual Machines page.</p> <p>The names of the ESXi host, VM, and appliance to which a specific vVol belongs are displayed in the vSphere Host Name, Virtual Machine, and Appliance columns on the Virtual Volumes card of the Storage Container details page.</p> <p>The ESXi hosts to which a vVol Storage Container is accessible are displayed on the ESXi Hosts card of the Storage Container details page.</p> <p>The ESXi version running on each recognized vSphere host is displayed.</p>	<p>This feature allows users to easily use vVols by registering the VASA provider from the PowerStore manager by adding vCenter information and provides vSphere host level visibility into PowerStore objects associated with these hosts.</p>
Other	<p>Improves PowerStore datapath performance for:</p> <p>VM Boot/Shutdown Storm that uses a large amount of UNMAP commands</p> <p>File System formatting</p> <p>Any thin aware application generating large amounts of UNMAP (truncate/trim of files) commands</p>	<p>Performance and efficiency improvement of third-party applications.</p>

Changed Features

The following functionality has changed.

Functional Area	Feature Description	Summary of Benefits
Product versioning	<p>Product releases will be identified as Major and Minor releases.</p> <p>Major releases will be versioned by increasing the first digit of the software version followed by a zero for example the first release is 1.0. Subsequent releases will increase the first digit as follows: 2.0, 3.0.</p> <p>Minor releases will be identified by the second digit of the software version number for example: 1.1, 1.2, 2.1, 2.2.</p>	Easily identify whether a new release is a major or minor software release.

Resolved Issues

The following issues have been resolved in this release.

Issue ID	Functional Area	Description
MDT-262726	Connectivity - Hosts	On a PowerStore X model cluster, the timesynced service was enabled, but not used. The time sync service is provided by the NTP service running on the hosting ESXi. The timesynced service will try to reach the factory default NTP server. It can be disabled by a support representative.
MDT-189246	Connectivity - Hosts	Cannot reuse the Logical Unit Number when unmapped and then mapped to the same host.
MDT-261727	Connectivity - Hosts	A host cannot be removed from a Host Group because another host in that Host Group has an attached or bound vVol.
MDT-240046	Connectivity - Hosts	A PowerStore appliance node may stop processing IOs if it is used with ESX clustered hosts where Atomic Test and Set (ATS) locking and other locking mechanisms are used at the same time.
MDT-234322	Connectivity-Hosts	A volume on a Windows Server cluster that is connected to a PowerStore cluster using Fibre Channel might become inaccessible after performing an upgrade from PowerStore 1.0.2 to a later version of PowerStore.
MDT-244501	Connectivity-Networks	Multiple ARP responses caused virtual network environments such as Cisco ACI to incorrectly detect duplicate IP addresses.

Issue ID	Functional Area	Description
MDT-305819	Connectivity-Networks	The following error occurred while trying to map or unmap to the Default Storage Network or remove the Global Storage IP address after upgrading from PowerStore 1.0.x. "Encountered an irrecoverable error during command execution, Reason not_supported (0xE04020040004)"
MDT-253220	Connectivity-Networks	iSCSI port remains down after retrying an unsuccessful VLAN change.
MDT-197207	Connectivity - Networks	The management port being down alert is reworded from "The first onboard port of node <A B> is down." to "The management port of node <A B> is down."
MDT-258225	Connectivity - Networks	If port 0 and port 1 of 4-port card are configured with speed less than 10 Gbps on PowerStore X, a warning is provided during Initial Network Validation. However, cluster creation is not prevented and will fail during network initialization.
MDT-260658	Connectivity - Networks	Top-of-Rack switch credentials did not include permissions to read basic switch configuration as well as running configuration information.
MDT-260649	Connectivity - Networks	Unexpected cyc_net error occurred, which indicates the validation component unavailability during PUHC. Contact your support representative and provide the (cyc_net_error) message.
MDT-194277	Connectivity - Networks	The create cluster and add appliance operations will be blocked if there are duplicates found in the L2 environment for the provided addresses.
MDT-209202	Connectivity - Networks	If PowerStore is configured with an unstable network environment or the physical switches are experiencing spontaneous packet drops, it may produce alerts about unreachability of some IP addresses and may auto-clear them after a short period of time.
MDT-278594	Connectivity - Networks	Depending on customer's network environment NTP servers used for PowerStore configuration may not respond spontaneously. In such cases PowerStore will produce alerts about NTP availability, that may be auto-cleared after a short period of time. If Users are aware about the health of their NTP servers, they may ignore or acknowledge these alerts. But correctness of functioning of PowerStore's services that require precise time synchronization (replication etc.) is not guaranteed.
MDT-210743	Connectivity - Networks	Add appliance fails due to wrong IP provided after an attempt to add an appliance with incorrect IP addresses. Primary appliance kept records of the IPs from a failed attempt to add an appliance, which prevented the same appliance to be added later to the cluster with correct IPs.

Issue ID	Functional Area	Description
MDT-242362	Connectivity - Networks	When the Intra-Cluster Management (ICM) and Intra-Cluster Discovery (ICD) networks are not available, and IPv6 ICM addresses not pingable, the <code>svc_diag list --info'</code> command will not run on the peer node. The workaround is to run the command on the master node.
MDT-266289	Data Protection	On the destination cluster pre-NDU (non-disruptive upgrade) steps pause all the replication sessions. The pausing of replication sessions for NDU deletes the system snapshot used by replication session. After the NDU completes, when replication session synchronization is attempted (via RPO, REST or PowerStore Manager), it fails because system snapshot was deleted on destination cluster.
MDT-141379	Data Protection	When processing the deletion of many snapshots on a large volume group, PowerStore Manager may freeze for more than one minute. Control should be restored after the deletion is complete.
MDT-85401	Data Protection	The replication session is not displayed if the Replication tab of a resource's protection page is open when assigning a protection policy with a replication rule.
MDT-146780	Data Protection	When deleting a snapshot rule, the snapshot rule may not appear to be deleted, and the following error appears: "The system encountered unexpected backend errors. Please contact support. (0xE0101001000C)". Do not continue using the rule for protection.
MDT-196359	Data Protection	When viewing the table of snapshot rules returned from the CLI command <code>show days_of_week</code> , not all days are displayed correctly.
MDT-219835	Hardware	Rare system bond issues occurred when one of the Ethernet port links stays down after the System Bond was created at startup time.
MDT-250110	Hardware	An HSST lockup occurred on the SAS controller and the workaround was not successful in recovering, resulting in a single node panic/reboot with a failover, reboot, and the node returning to service.
MDT-232113	Hardware	After adding more drives to the array (more than 6 drives), the next appliance shutdown is at risk of not recovering.
MDT-217478	Hardware	After recovering from multi disks malfunction where the system entered "service offline mode", the system was able to serve IOs but some volumes may not have been accessible due to configuration issues.
MDT-270425	Hardware	Replacing an Expansion Enclosure's LCC when the enclosure is mis-cabled may prevent the new LCC from coming online. Because the system was in a bad state due to mis-cabling, the old LCC information was not cleared prior to adding the new LCC information. Code modifications have been made to check that the system/enclosure state is good prior to adding the new LCC information.
MDT-218461	Hardware	A failed service was detected in error, causing an unnecessary node failover.
MDT-233579	Hardware	Access to all expansion enclosures was lost.

Issue ID	Functional Area	Description
MDT-231279	Hardware	Upon issuing managed shutdown to one of the storage controllers, in rare cases the unmanaged shutdown (reboot) might occur instead.
MDT-247447	Hardware	In some cases of node reboot due to kernel panic we are unable to root cause the issue because the kernel dump does not complete in time allowed. This is because we prioritized fast reboot over diagnosable materials. Timing of this process is under review.
MDT-272231	Hardware	Cluster cannot add two new appliances "The system encountered unexpected backend errors."
MDT-253764	Hardware	SYM to fix sym_mom_is_ssd_fru_state_ok so it will check that all healthy PMs sees the drive as OK. SYM will not accept a new drive to DP before it is up on all healthy nodes.
MDT-237899	Hardware	While performing HA testing customer issued command to reboot node_a multiple times, this indeed triggered reboot of node_a and then node_b reboot causing outage. When reboot was initiated for a second time, the command was issued to the second node which has CP as the failover which completed from first reboot request and thus triggered reboot of second node in addition to first node reboot.
MDT-241178	Hardware	The access control account for the certificate created on primary node failed to sync with member nodes in the cluster due to synchronization framework failure because of the underlying OS signals. As a result, when failover happens, the certificate account will not be on the newly elected primary node and the corresponding application will report the missing cert account. The synchronization failure is mitigated and fixed in this release.
MDT-194839	Hardware	During NDU, in rare cases the upgrade is stopped and drive failure tiers are left in a degraded mode while the system continues to service I/Os.
MDT-233808	Hardware	Rare race condition on system boot between iSCSI portal configuration and ESXi iSCSI adapter rescan.
MDT-252811	Hardware	PSU is not marked as disconnected state when its sensors don't get updated in OS.
MDT-232303	Hardware	After adding drives and restarting SYM, sometimes CLST didn't rebuild SSD group. In that case if adding drives to the point that the user added more than half of the total new drives - there might be a possible loss of majority. This causes SYM not being able to restart
MDT-260065	Hardware	BBU is marked as disconnected state when its sensors don't get updated in OS.

Issue ID	Functional Area	Description
MDT-205619	Hardware	<p>During a node replacement procedure if the node is inserted into the chassis and powers on without the embedded module inserted or properly seated, the node will enter service mode and requires service engagement for recovery.</p> <p>This change does not eliminate the error condition or the need to take special recovery steps, however, the firmware was updated to simplify the process.</p>
MDT-261674	Hardware	Firmware issue causes kernel panic leading to node reboot.
MDT-254816	Hardware	Under certain single node failures, NVRAM drives can be marked as offline due to NO_PATH condition. This leads to the system taking the drives offline and initiating a Halt & Vault process, resulting in a DU.
MDT-283877	Hardware	A null pointer was referenced when looking up a host data structure.
MDT-201033	Hardware	PowerStore Manager incorrectly reports all drives are disconnected.
MDT-174416	Import	Hosts groups are not supported for agentless import.
MDT-147910	Import	Importing a volume from a source system to PowerStore fails while switching path on host. This issue only impacts Linux operating systems. If any "blacklist" keyword appears in the multipath.conf file before the blacklist {} block, the host plug-in fails to blacklist the source volume and will cause an import to fail. For imports to be successful, the keyword "blacklist" and curly brace should appear in the same line, for example, "blacklist { " in /etc/multipath.conf file.
MDT-303583	Install or Upgrade	While attempting to upgrade from version 1.0.* to version 2.0.0.0 (Build 1376722), two Pre-Upgrade Health Checks that validate network configuration incorrectly report a problem and block the upgrade attempt. This issue only affects SAN-only configuration using only Fibre Channel hosts without a top of rack switch.
MDT-305868	Install or Upgrade	<p>After upgrading to 2.0.0.0 (Build 1376722), attempts to access PowerStore Manager failed with error, "Unable to load system configuration data. Contact your service provider for assistance. Error: Service unavailable. Please retry in a few minutes (0xE04040010004).</p> <p>This error can happen if volumes were moved from 1 volume group to another creating multiple entries in a volume group membership table. These entries are incorrectly tagged as problematic during upgrade causing upgrade to fail.</p>
MDT-305874	Install or Upgrade	Due to a race condition in PowerStore 1.0.x, inconsistency in the metadata occurs and is not a corruption. After upgrading from PowerStore 1.0.x to PowerStore 2.0.0.0 (Build 1376722) a background job is run. This background job detects this inconsistency and continues raising the following alert "Schedule mapper recovery (later)" after upgrading from PowerStore 1.0.4 to 2.0.0.0 (Build 1376722).
MDT-191912	Install or Upgrade	When running the Initial Configuration Wizard, after the cluster creation step you may experience a timeout or loss of connection. The cluster, however, is still being created.

Issue ID	Functional Area	Description
MDT-238657	Install or Upgrade	The system shows an alert in case of the issues with the connection to the vCenter. The alert presence won't prevent the NDU from happening in PowerStore 1.x versions. Starting with 2.0 versions, the raised alert of this type will block the NDU.
MDT-257003	Install or Upgrade	During the NDU, cyc_rcvy_control service might not be enabled, the cyc_rcvy_control and as a result, sub-sequent tasks related with NDU will fail, including reimaging. s.
MDT-250029	Install or Upgrade	NDU may fail if there are stale locks in the management DB. The locks must be manually removed from the management database via a service technician. There are failure conditions surrounding HA events (node reboots and CP restarts) that can leave stale locks in the management database. In general, the condition is more prevalent in a clustered-PowerStore environment but can occur in single appliance systems with less frequency.
MDT-262439	Install or Upgrade	During NDU a local connection can be reset causing a command failure and results in a failover. If there is another occurrence, the recommendation is to retry NDU. Note: This flow is not necessarily NDU specific.
MDT-225893	Install or Upgrade	After running a factory reset, using CTRL+C made both nodes unusable.
MDT-259371	Install or Upgrade	During a Health Check or Upgrade of a PowerStore X Software Release package, the Pre-Upgrade Health Check can fail when verifying the ESXi PowerStore X VIB versions with the error "Component version is in incorrect format". This issue is specific to PowerStore X models with RecoverPoint installed, that have VIB which erroneously gets checked as part of the VIB PUHC and as a result causes PUHC to fail.
MDT-242768	Install or Upgrade	The PowerstoreX-node_firmware package is expected to be used only on PowerStore X systems as part of its special multi-step upgrade process. The node_firmware package is not intended to be used with PowerStore T systems as the firmware will be upgraded along with the Trident Software release. The newer versions of the PowerStoreX-node_firmware have been updated to block them from being installed accidentally on PowerStore T systems.
MDT-253108	Install or Upgrade	Multiple subsequent NDU failures from PowerStore 1.0 release can leave array in a state where NDU code fails to execute roll forward behavior correctly and continues to fail.
MDT-214506	Install or Upgrade	An attempt was made to update PowerStore 1.0.x to PowerStore 1.0.2 while the management ports on node B were not properly set up. After correcting the problem, the NDU was restarted several times and each time the following error occurred: "Upgrade failed on appliance appliance-1 with the following message: The system encountered unexpected backend errors. Please contact support. (0xE04030010003)".
MDT-215818	Install or Upgrade	While upgrading from PowerStore 1.0.1 to PowerStore 1.0.2 the upgrade failed and was restarted a few times before finally completing. However, the release version shown in PowerStore Manager (GUI) still displayed PowerStore 1.0.1.

Issue ID	Functional Area	Description
MDT-259223	Install or Upgrade	Upgrade from PowerStore 1.0.3 to 1.0.4 failed because of hung commands. This can happen after an unsuccessful attempt to add or remove an appliance and some commands did not complete and remained on the system. When an upgrade is initiated on the system, it will wait for all running commands to complete to start the upgrade. Since these commands are hung, upgrade will fail after waiting for 30 minutes.
MDT-236192	Internal Migration	Internal migration of imported volumes from a remote system using the non-disruptive import variant does not redirect a host application running on either Linux/Windows or VMware to the associated PowerStore volumes.
MDT-143768	Internal Migration	When migration is occurring on a volume, the Mapped Volumes tab of the hosts/host groups that are attached to the migration source volume might show some additional volumes. The migration uses these additional volumes, which are internally created. These additional volumes are deleted when the migration finishes.
MDT-145764	Internal Migration	During migration the temporary volumes that are created during migration are calculated into the vVol capacity metrics. The temporary volumes should not be included in the calculation.
MDT-143749	Internal Migration	On the Internal Migration > Migrations tab in PowerStore Manager, you may see incorrect results when you try filtering results based on matching strings in the Storage Resource column.
MDT-231942	Data Collection	Inaccurate data logging related to expansion enclosure connection status.
MDT-247636	Data Collection	Data collection fails and the data collected daily is not being sent to CloudIQ.
MDT-257551	Data Collection	Appliance metrics data uploaded to CloudIQ are occasionally duplicated.
MDT-257727	Data Collection	Deltas of configuration and metrics can be sent in too close succession due to timing issues, causing confusion on CloudIQ side as packages are processed out of order.
MDT-233577	Data Collection	Inaccurate data logging related to expansion enclosure connection status.
MDT-217411	Data Collection	Data collection cannot complete while the replication database is in read-only mode.
MDT-237832	Data Collection	In certain cases, the system was permitting two data collections to run at a time, causing a failure.
MDT-247214	Data Collection	Unable to collect data because a file changed in the process of being added to the tar file. Tar file returns an error code of 1.
MDT-263836	Data Collection	On systems that generate large amounts of metrics information, the metrics provided by support metrics are too large for the script to process.

Issue ID	Functional Area	Description
MDT-90271	Data Collection	Changing the members of a Volume Group may cause that Volume Group's performance and capacity metrics to be reported incorrectly at a specific timescale for a specific time period. For details see, Dell knowledge base article 000132968: PowerStore: Supported Host operating system for Non-Disruptive migration of Storage resources
MDT-148297	Data Collection	A problem exists in the collection of replication copy metrics which affects the 'data remaining' copy metric for the synchronization of a given replication session. In PowerStore Manager, the graph of the remaining data will show twice the actual value, and in some cases this value shown could be larger than the actual volume size. This problem will self-correct during the next synchronization of the replication session, however the erroneous value will remain on the historical graph of the 'data remaining' until the time scale on the graph expires.
MDT-261351	Data Collection	Configuration capture requests are clashing and preventing some telemetry data from being uploaded to CloudIQ, affecting charting.
MDT-145175	Monitoring	When a volume is migrated, there is a period of time where we have associated source and destination volumes across appliances; however, performance metrics are not being gathered for both. Even so, some performance metrics may be erroneously reported for the cluster for the volume no longer receiving I/O. Metrics will be reported properly for the individual volumes.
MDT-224269	Notifications and Alerts	Alert remains even after HA issue resolved. Alert must be manually removed by a support representative.
MDT-266668	Notifications and Alerts	If PowerStore X is connected to a vCenter with enabled SSO authentication, Users may periodically see alerts about vCenter login problems on PowerStore X side. This alert should be ignored.
MDT-231435	Notifications and Alerts	In some cases, particularly after a node reboot, a false alert may display that a node is offline.
MDT-183731	Notifications and Alerts	An unhelpful error message was returned while configuring SMTP. The error was caused by permission issues on the SMTP server.

Issue ID	Functional Area	Description
MDT-249135	Notifications and Alerts	<p>On a full system (90% expanded) an alert is raised indicating a possible data corruption.</p> <p>Alert info: Description :A storage object resource is offline and requires repair in order to bring back online (mapper recovery later) Details: "Major","0x00200f02","A storage object resource is offline and requires repair in order to bring back online.,"Appliance","power01-appliance-2".... [AlertNotifier] [INFO] [com.emc.cyclone.contexts.eventprocessing.eventframework.application.alertnotifier.AlertNotifierStateTasks vert.x-eventloop-thread-0] [778272ab-d75c-4cf7-86de-062507f2dd1a d512ad48-8613-434f-93cc-1717ba414e97 processSyslogNotification_execute] Alert Notification : DATA_PATH_SCHEDULE_MAPPER_RECOVERY_LATER (0x00200F02) Objectid:A2 Type:appliance ApplianceSN:A2 Raised:2021-01-26T01:32:56.734053Z Severity:Major A storage object resource is offline and requires repair in order to bring back online. Jan 26 03:40:34 CKM01204605182-B control-path[84951]: 2021-01-26</p>
MDT-251269	Notifications and Alerts	<p>SYM will remove the mom object from its own database and XMS will also do the same. However, the code that clears up alerts towards control path is not being called at this time which might causing a pending alert.</p>
MDT-245640	PowerStore Manager(GUI)	<p>When rebooting or powering off a node from PowerStore Manager, the system may not display any acknowledgement that the process has begun. The Reboot or Power Down button may still appear clickable. Despite the lack of acknowledgment, the operation will be triggered successfully after a single click of the button.</p>
MDT-227404	PowerStore Manager(GUI)	<p>Users in the Brasilia time zone (BRT), which is referred to as "Americas/Sao_Paulo" in the IANA time zone database, may notice that timestamps in the PowerStore Manager appear an hour later than they should. This only occurs from November 1st to February 15th and is only an issue with the conversion done in the GUI. This confuses the end user as they believe that their snapshots and other scheduled activities are occurring an hour later than they should be. However, all activities are occurring at the correct time and since the actual timestamps are stored internally as UTC they remain accurate.</p>
MDT-245142	PowerStore Manager(GUI)	<p>On the Settings Support Materials page, the page has a PowerStore Manager (GUI) imposed limit of 100 entries which can prevent the full list of captures from being displayed if the number of actual entries exceed 100.</p>
MDT-189872	PowerStore Manager(GUI)	<p>In the PowerStore Manager (GUI) on the "Host & Host Group" page, a user will see "--" instead of OS name for a Host Group if all hosts have the same OS.</p>
MDT-193084	PowerStore Manager(GUI)	<p>Cannot login to PowerStore Manager (GUI) after NDU from PowerStore 1.0 to SP1.1 and the following error is returned "RESTAPI failed with 503 Service Unavailable."</p>
MDT-86305	PowerStore Manager (GUI)	<p>When you power down a node in PowerStore Manager, the screen does not automatically refresh the state to indicate that the node is getting powered off.</p>

Issue ID	Functional Area	Description
MDT-159008	PowerStore Manager(GUI)	In PowerStore Manager, on the Monitoring > Events tab, the Apply button may remain grayed out after you tried filtering based on a specific start and end date/time.
MDT-126148	PowerStore Manager(GUI)	Signposts or large tooltips that display detailed information while hovering over values in the Policies tables can be truncated.
MDT-146302	PowerStore Manager(GUI)	On the Host Mappings tab for a volume, the value in the logical unit number column appears as "--" when the logical unit number for a host is 0.
MDT-145344	PowerStore Manager (GUI)	While performing a Refresh Using Related Volume or Volume Group operation, you might receive an exception error when you select Choose One from the drop-down before you click Refresh. If this occurs, you must restart PowerStore Manager.
MDT-142439	PowerStore REST API	The "*" and "@" characters cannot be used for the Host Access List Name from PowerStore Manager.
MDT-254352	Security	Following 5 special characters and the unicode characters will be supported in the LDAP user distinguished name, member/memberUid, userPrincipalName, sAMAccountName and common name. \\ () # NUL
MDT-135211	Security	Job_id field in the audit_event REST API is currently not supported and contains an invalid value.
MDT-294054	Storage - Block	During an NDU, or when a node is rebooted/restarted, hosts will lose access to any LUNs that are configured as LUN 0. All other LUNs will not be affected. In order to regain access to LUN 0, the host needs to be logged out and logged back in again.
MDT-238415	Storage - Block	The restore and refresh operations will overwrite the contents of the storage resource with the data from the selected snapshot or clone. Data not flushed from the host cache before the operation could cause data integrity issues after the restore or refresh.
MDT-265880	Support	A race condition was causing the service responsible for uploading files to SupportAssist to fail, requiring a restart of the Control Path container.
MDT-234621	Support	When performing a factory restart using the svc_factory_reset script, there is an issue in which the Master License Key for PowerStore X clusters is being wiped.
MDT-238662	SupportAssist	A race condition in the Telemetry system caused file uploads to SupportAssist to stop.
MDT-221099	SupportAssist	The svc_remote_support command can't force-disable SupportAssist in a multi-appliance cluster.
MDT-257286	SupportAssist	NDU may fail due to a SupportAssist process that is still running. If this happens, restart NDU.

Issue ID	Functional Area	Description
MDT-211234	SupportAssist	A failed ADD_APPLIANCE job may be seen in the list of jobs. This is an erroneous message.
MDT-269756	SupportAssist	A regression caused a key to not be properly stored on Node B of the appliance, preventing SupportAssist enablement from working when Node B is the active node.
MDT-233572	SupportAssist	Node reboot can occur due to the processing of telemetry data being sent to SupportAssist. Each series of telemetry information is sent as a unique file. Under some circumstances file processing of the telemetry information can result in too many files being open concurrently resulting in a node reboot.
MDT-252480	Virtualization	A disruption might occur if vCenter starts an enterMaintenanceMode operation, prompting the CP to begin sending shutdown requests using the service script svc_node shutdown.
MDT-260884	Virtualization	The pdaemon component, which polls information about ESX vmnics/vnics, may hang on a vSphere API call (e.g. when getting info about hosts and VMs). If pdaemon hangs, PM (which is monitoring its status) detects it's unresponsive, and issues a request to kill pdaemon, and if succeeded, to restart it right away. The request to kill pdaemon fails due to inaccurate handling of pdaemon's files, and PM doesn't continue with pdaemon's restart either.
MDT-236381	Virtualization	Unavailability of ESX on one of the nodes due to the 'multiple pdaemon' problems.
MDT-186680	Virtualization	High CPU/memory usage alarms are always shown in vCenter for PowerStore X controller VMs.
MDT-276020	Virtualization	SSH command not running on node.
MDT-235836	Virtualization	vCenter v.7.0u1 is not supported.
MDT-148124	Virtualization	Metrics data may be unavailable in PowerStore Manager if a PowerStore X model appliance cannot communicate with vCenter Server due to a vCenter Server failure or connection problem. This issue also affects metrics that do not directly depend on vCenter Server statistics, such as Capacity metrics.
MDT-246284	Other	stitch.py is not properly handling fields with multibyte characters.

Known Issues

The following high severity issues are reported in this release. For a list of the minor issues, see the [PowerStore Release Notes – Minor Issues knowledge base article Number: 000155626](#).

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-245165	Cluster - Creation	Each PowerStore appliance is pre-configured with an Administrator user (admin), which has a default factory password. The PowerStore Initial Configuration Wizard will require you to change this password. Once you have changed the admin user password, you must complete the cluster creation process using this appliance as the primary one. This appliance can no longer be configured as a non-primary appliance in a cluster, because only the primary appliance is permitted to have the non-factory default password.	Contact your service representative if you mistakenly change the password on an appliance and want it to be a non-primary appliance in a cluster.
MDT-239230	Connectivity- Hosts	When a host becomes unavailable, mounting a Storage Container will fail.	Bring all hosts in the cluster back online and retry mounting the Storage Container. Alternatively, mount the Storage Container manually only to available hosts using the vSphere UI.
MDT-284390	Connectivity- Hosts	Bootimg from SAN is not supported with FC-NVMe on AIX in PowerStore 2.0.0.x.	None
MDT-152854	Connectivity - Hosts	Running an unmap operation might cause a host I/O latency impact of more than 15%.	Disable auto-unmap for applications that have this feature. Schedule unmap operations manually to reduce the impact on hosts.
MDT-139488	Connectivity - Hosts	Xcopy operations can impact host I/O performance.	None
MDT-273055	Connectivity - Hosts	A user can view internal host and volume mappings via GUI but cannot perform mapping/unmapping. Mapping/unmapping actions can be performed via REST interface.	None
MDT-288876	Connectivity - Network	On rare occasions, the create cluster operation may fail with the following error, "Unconfigured appliance with the link local address belonging to the primary appliance was not found, Control Path (CP) container should be restarted." due to failed registration to the internal Zeroconf service.	First try to restart the CP container using the "svc_container_mgmt restart CP" command. If a restart does not work, perform a factory reset on the primary appliance using the "svc_factory_reset" command. Contact your service provider for more information or assistance.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-272385	Connectivity-Networks	Large amounts of external IPv6 traffic exposed to ethernet ports can cause iSCSI performance degradation.	To avoid the problem, use separate VLANs for Storage networks and networks that contain a lot of non-storage related IPv6 traffic. Also, route management VLANs and other VLANs with large amounts of IPv6 traffic to dedicated storage ports.
MDT-284055	Connectivity-Networks	In certain network environment configurations, the system is unable to discover unconfigured appliances through the network and may experience node reboots.	Contact your service provider.
MDT-211358	Connectivity-Networks	PowerStore Discovery IP by Zeroconf technology stopped advertising on PowerStore 3000X model.	Perform either one of the following actions to work around this issue: Either, wait a few minutes for the state of the system to change and try again. or Reset Drive, Mezz or I/O Module if available and wait for few minutes and try again, then reboot the primary node.
MDT-55667	Data Collection	Storage container capacity metrics are not available until I/O operations have begun.	Look at the capacity metrics again after I/O operations on the storage container have been performed.
MDT-240072	Data Collection	Uploaded data collections can be duplicated on massive file transfers. In this case make sure to take the largest of the duplicated data collections.	None
MDT-280001	Data Protection	In certain cases the "replication session resume" command gets stuck because of an internal error. This will cause PowerStore NDU to fail.	Contact your service provider.
MDT-282294	Data Protection	The modify volume command failed after it was issued to soon after a High Availability event.	Retry the modify volume command.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-220994	Data Protection	The replication session may not be deleted once a Protection Policy is unassigned from the storage resource.	<p>To delete the replication session, do the following:</p> <ol style="list-style-type: none"> 1. If Protection Policy is not assigned to the storage resource, reassign the same Protection Policy that was previously assigned to the storage resource. 2. Unassign the Protection Policy from the storage resource.
MDT-133408	Data Protection	When a volume group is created with member volumes and is protected by a protection policy that includes a replication rule, the volume group members may be delayed in being replicated to the remote system until the next RPO cycle. The default replication rule RPO value in PowerStore Manager is 1 hour, but it can range from 5 minutes to 24 hours.	If the remote system connection is healthy, the member volumes should automatically be replicated to the destination during the next RPO cycle. To update the member volumes on the destination before the next RPO cycle, navigate to Protection-> Replication, select the volume group session and perform a "Synchronize" operation.
MDT-242915	Data Protection	When replication destination cluster loses connectivity to source cluster, the system issues specific alerts/events, but the state of the replication session may still display Operating Normally status. Admins should check the alerts issued by the remote system and resolve the connectivity issues to recover replication sessions.	None
MDT-252199	Hardware	The "Do Not Remove" LED light is not working correctly on PowerStore 500T model appliance nodes. After powering down a node on a PowerStore 500T model appliance, the "Do Not Remove" LED light on the other node does not light up.	There is no workaround just be sure to remove the powered down node before continuing.
MDT-289908	Hardware	A newly inserted disk remains locked, and the rebuild of the PowerStore Model appliance times out.	Restart the node.
MDT-286236	Hardware	In rare cases, a node that has been powered off does not remain powered off.	Power off the node again.
MDT-288084	Hardware	After replacing a node, a firmware issue causes the replaced node to try to boot from an external server that is unavailable instead of restarting from the internal M.2 boot module.	Contact your service provider.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-285058	Hardware	In some cases, after the appliance stops servicing IOs and then goes back online, the "stopped servicing IOs" alert is not cleared automatically even though the system is working correctly.	If the hosts are connected and the IOs are being serviced, the alert can be ignored.
MDT-117061	Hardware	When beginning cluster creation, the hardware status is incorrectly indicated as not configured for clustering.	Wait a few minutes and try again.
MDT-190232	Hardware	The svc_node shutdown command may reboot the node instead of powering it off.	Run the svc_node shutdown command again when the node has finished rebooting.
MDT-263147	Hardware	If you add an appliance to a cluster when another appliance in the cluster is powered off, the add appliance operation will fail.	None
MDT-260313	Hardware	One or both nodes may not automatically power back on after a short AC power outage of less than 50 seconds. The appliance should power on properly after the first AC failure, but it is vulnerable to subsequent short AC failures.	<p>If only a single node does not power back up PowerStore Manager will display a major alert for the impacted BaseEnclosure-Node - with the following message "Node has been physically removed or shut down" (Event code 0x00304404). Log into the system service account and use the "svc_node power_on" command to power up the peer node.</p> <p>If neither node of the appliance is accessible after 30 minutes following a short AC power failure the appliance will need to be reset. To reset the appliance, remove the AC power for at least 120 seconds and then restore AC power and both nodes should power up and return to service within 30 minutes.</p>
MDT-137439	Import	If an import from a Unity or VNX2 system to PowerStore is cancelled by the user or fails for some reason, it is possible that subsequent attempts to import the volume will also fail.	Remove the Unity or VNX2 as a remote system and then add it again. This action will clear the issue and allow the import to proceed.
MDT-252560	Import	Import of a volume failed with the error "Import failed because the mirror service could not be enabled on the Destination volume".	Run the verify operation on the remote system to recover data connections. Once the remote system verify operation completes successfully, try the import operation again.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-258429	Import	If multiple storage networks are configured on PowerStore and each of these networks has a dedicated Global Storage Discovery IP Address (GSDIP), non-disruptive import fails with Error Code 0xE0301002002C.	See Dell knowledge base article 000184766: PowerStore: Non-disruptive Import failure when more than one Global Storage Discovery IP Address is configured.
MDT-286363	Internal Migration	After migrating volumes, removing an appliance from the cluster failed with the following error message: "Failed to delete target port group."	Contact your service provider.
MDT-290965	Internal Migration	In a PowerStore multi-appliance cluster, internal migration of volumes or volume groups used by Microsoft Windows Failover Cluster may result in hosts losing access to storage.	Internal migration of volumes or volume groups used by Microsoft Windows Failover cluster must be done with no active I/O during cutover. I/O can be quiesced by powering down the host or unmapping the volumes. Failure to quiesce I/O during cutover could lead to data unavailability.
MDT-285903	Internal Migration	The migration failed due to an High Availability (HA) event and a system snapshot remains on the migration source object.	Contact your service provider to clean up the remaining snapshot.
MDT-167020	Internal Migration	Migration of a storage resource failed with the error "Transit object with Handle <id> : status online on only one node." One of the nodes of the migration source appliance was down when the synchronize phase started. Both nodes must be online for migration.	Bring the offline node back online and then try the migration again.
MDT-140927	Internal Migration	After a volume is migrated, there may be a dip in capacity and performance metrics reporting for that volume at the end of the migration operation.	None
MDT-273051	Internal Migration	If all objects are selected, or more objects are selected than the destination has the capacity for, a failure message may appear when attempting to migrate those objects from one appliance to another using the assisted migration workflow.	Select objects individually and migrate volumes by using the manual migration workflow instead of using the assisted migration workflow.
MDT-263177	Internal Migration	One or more vVols might be inaccessible to the host after an internal migration completes successfully.	The inaccessible vVols can be made accessible to the host by manually updating the compute resource within vMotion. For details, see Dell knowledge base article 000105896: PowerStore: Supported Host operating system for Non-Disruptive migration of Storage resources.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-239624	Internal Migration	<p>When a VM is being deployed from a VM template, vSphere uses the clone operation to provision vVols of the newly created VM. PowerStore internally translates it into space-efficient clones (the same technique used for vVol snapshots). As soon as a Data Path (DP) family size limit exists, a base vVol cannot have more than a certain number (1000) of derivatives (snapshots, space-efficient clones). This causes the following limitations:</p> <ul style="list-style-type: none"> - Cloning from a VM template fails if a base vVol of the VM template has too many derivatives (snapshots or clones, or both) already created (1000). - When the maximum number of family members limit is reached, vVol migration cannot occur as it needs to create an internal vVol snapshot which fails. - All vVols are placed to the same appliance in a multi-appliance cluster, so the load is not balanced between appliances. 	<p>Use several VM templates created from the same OVF template so that each VM template is used for a reasonable number of VM clones to avoid hitting the DP family limit. This action allows vVols to be distributed between appliances in a multi-appliance cluster for load balancing.</p>
MDT-242178	Internal Migration	<p>A known ESXi issue might cause a vVol migration to fail if a vVol is being migrated to an appliance on which the vVol was previously located.</p>	<p>Reboot the ESXi host to send an unbind request for the vVol and to allow subsequent migrations of the vVol to the appliance.</p> <p>Note: If the cluster contains more than two appliances, the vVol can be migrated to another appliance in the cluster.</p>
MDT-283025	Install or Upgrade	<p>A PowerStore software upgrade fails if the PowerStore X model cluster cannot communicate with the vCenter Server during the upgrade.</p>	<p>Contact your service provider to recover from the failed upgrade.</p>
MDT-288380	Install or Upgrade	<p>During upgrade if Physical Large Block (PLB)/Virtual Large Block (VLB) metadata corruption is discovered during VLB conversion, the VLB conversion keeps retrying as a low priority background job until the corrupted metadata is fixed through a file system consistency check (FSCK).</p>	<p>Run FSCK as indicated by the Fault Alert. When the metadata corruption is cleared, the system automatically continues the VLB conversion until it completes.</p>

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-164726	Install or Upgrade	If the Pre-Upgrade Health Check fails during an upgrade from PowerStore OS 1.0 to PowerStore OS 1.0.1 or later, clicking the failed operation link under Recent activities might result in an error because of a localization issue with the event message.	None
MDT-219735	Install or Upgrade	If a PowerStore 1.0.3.0.5.007 software upgrade package is uploaded to a PowerStore T model cluster with a single appliance running PowerStore OS version 1.0.1.0.5.002 or 1.0.1.0.5.003, then an appliance running PowerStore OS version 1.0.2.0.5.003 is added to the cluster and another PowerStore 1.0.3.0.5.007 software upgrade package is uploaded to the cluster, performing an upgrade will fail before the upgrade process completes on the primary appliance and the second appliance will not be upgraded.	Contact your service provider.
MDT-156905	Install or Upgrade	After a software update, the NAS servers may not fail back automatically to the original node.	Manually fail back the NAS servers from PowerStore Manager.
MDT-135505	Install or Upgrade	Drive firmware updates may not get applied to a single drive or a few drives within the appliance. This may occur because of a timing issue between the two nodes in the appliance.	Download and install the drive firmware again to ensure it gets updated on the drives that were missed in the previous attempt.
MDT-170949	Install or Upgrade	Running PowerPath 7.1 and earlier on a Linux host causes an I/O error during the software upgrade of PowerStore OS 1.0 to PowerStore OS 1.0.1.	Apply the next/upcoming PowerPath 7.1 patch, and then run the PowerStore software upgrade to avoid the I/O error.
MDT-265301	Install or Upgrade	When performing a software upgrade on a PowerStore T model appliance running in Unified mode, the inter-node iSCSI connection might fail when one node is rebooted, which can cause the NAS to panic.	Allow NAS-HA to take care of uninterrupted NAS operations during a software upgrade. If data unavailability occurs after the software upgrade is complete, contact your service provider.
MDT-256340	Monitoring	Config capture may fail to calculate deltas on very large systems under very heavy load.	None

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-249810	Monitoring	If the total system capacity is below 10TB, free page tables are not replenished as this would consume space that might not be required and waste free pages. The lack of free pages in the table causes a performance impact for writes. For better performance, more capacity should be added.	Make sure that the system capacity is higher than 10TB.
MDT-146346	Notifications and Alerts	Some hardware alert states may persist for a short time after the health of the object has been restored. The alert states eventually resolve with the correct state.	None
MDT-245640	PowerStore Manager (GUI)	When rebooting or powering off a node from PowerStore Manager, the system may not display any acknowledgement that the process has begun. The Reboot or Power Down button may still appear clickable. Despite the lack of acknowledgment, the operation will be triggered successfully after a single click of the button.	None
MDT-293294	PowerStore Manager (GUI)	<p>In PowerStore Manager during the selection of storage objects step of a Hardware appliance Migrate action workflow, an unexpected exception can appear if one of the following occurs:</p> <p>While All Available Objects and Eligible For Migration objects are selected, a user selects Not Eligible For Migration objects and then selects Select Objects Manually for migration.</p> <p>While Select Objects Manually and Eligible For Migration objects are selected, a user selects Not Eligible For Migration objects and then selects All Available Objects for migration.</p>	If an unexpected exception appears, reload the page to return to the Hardware page and try to migrate from the appliance again.
MDT-118394	PowerStore Manager (GUI)	Uploading an upgrade package that is larger than 4 GB with Microsoft Edge or Internet Explorer fails.	Try again with a different browser.
MDT-163489	PowerStore Manager (GUI)	Once a node is replaced, the new node hardware information does not appear in PowerStore Manager immediately.	After replacing a node, wait up to 30 minutes for the PowerStore Manager screen to reflect the update.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-261523	PowerStore Manager(GUI)	If Mozilla Firefox is used to upgrade a PowerStore cluster to PowerStore 2.0, a browser cache issue might occur that causes the names of some buttons and labels to be displayed as object names instead of the correct button or label names.	Use a different browser such as Google Chrome or Microsoft Edge. Or, clear the browser cache in Firefox. For instructions, see the KB article "How to clear the Firefox cache" on the Mozilla Firefox support site.
MDT-130894	PowerStore REST API	Two invalid resource_types, FEPort and Initiator, may be returned by the event and alert APIs.	Events or alerts with resource_type FEPort or Initiator should ignore the resource_type values. External OpenAPI clients must disable validation for the event and alert APIs with resource_types of FEPort or Initiator.
MDT-266683	PowerStore REST API	You cannot issue Swagger UI commands with Operator, Storage Administrator, and VM Administrator roles. You must log in with the higher-level Administrator role if you want to complete tasks such as fetching or reading data from PowerStore.	Customers can login in as administrator to use the swagger panel in PowerStore Manager.
MDT-153026	Security	When a drive is inserted into an appliance, it will take time for D@RE to unlock it. Most drives take a few minutes. However, it takes more time for a ColdStream drive to unlock. The time is proportional to the size of the Intel Optane SCM drive, for example, 7 minutes for 375 GB, 15 minutes for 750 GB, and 27 minutes for 1.5TB. If you pull a drive out before it is unlocked, the drive will appear as disconnected. However, the drive will not be automatically removed from the appliance.	If you intend to use the drive, re-insert the drive and let the DARE unlock complete. If you intend to stop using the drive, re-insert the drive, wait for enough time (time to unlock plus about 1 minute for other overhead), and then remove the drive. The drive will then become disconnected and automatically removed from the system.
MDT-213516	Security	DNS is not supported for LDAP.	None

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-282118	Security	<p>Timeout error occurs when trying to log into PowerStore using the LDAP directory service.</p> <p>PowerStore Manager login is set to timeout after 3 minutes. The LDAP timeout period can be set to a value longer than 3 minutes in PowerStore Manager. PowerStore Manager timeout setting supersedes the LDAP setting if LDAP is set to timeout longer than 3 minutes.</p>	<p>Be sure that the Advanced Settings in the PowerStore Manager, Directory Services is configured as follows:</p> <ol style="list-style-type: none"> 1. Go to Settings>Directory Services. 2. Click Edit LDAP Configuration. 3. In the Modify Directory Services Panel, change the Timeout value to less than 180 seconds, or a value less than 3 minutes. 4. Scroll down to the Modify Directory Services panel, and expand the Advanced Settings. <p>Set User Search Settings>Search Path as close as possible to the LDAP user.</p> <p>Set Group Search Settings>Search Path as close as possible to the LDAP groups that have LDAP user as a member.</p> <p>If the above steps still do not help, try to set Group Search Settings> Nested Group Search Level to 1</p>
MDT-284475	Storage - Block	After mapping a LUN or creating a VM, the LUN or VM appears correctly, but it is inaccessible to hosts.	Unmap the non-working LUN and map it again. For VMs, delete the non-working VM and recreate it.
MDT-242368	Storage - Block	Unable to delete a volume group after deleting a replication session, due to undeleted replication session system snapshots.	Contact your service provider for the procedure to delete the snapshots.
MDT-121966	Storage - File	The File system deleted alert is displayed as an active alert in PowerStore T even after the file system has been successfully deleted.	Review the alert details, to confirm the alert is for a file system that has been deleted. Once confirmed, ignore or acknowledge the pending alert.
MDT-121784	Storage - File	After a NAS server is successfully deleted, the alerts against the deleted child objects of the NAS server such as DNS servers, file systems, NFS exports, or SMB shares, are not cleared.	Review the alert details, to confirm the alert is for a child object of a successfully deleted NAS server. Once confirmed, ignore or acknowledge the pending alert.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-139095	Storage - File	While trying to create or delete a snapshot the following error message was returned: "Addition [or deletion] of NFS Export failed due to [The path [path_name] was not found on the system and cannot be exported.]"	Wait for some time and try the action again.
MDT-147688	Storage - File	SMB clients cannot connect to a NAS server when the NAS server name is more than 15 characters. This is because by default the SMB server computer name takes the NAS server name.	Set the NAS Server name with 15 characters or less or set the SMB server's computer name with 15 characters or less. To modify the computer name, select all the characters in PowerStore Manager, and type the new name.
MDT-110930	Storage - File	Access to a snapshot is denied when there is only one snapshot created on the file system, and access to the snapshot is attempted at the time the snapshot is refreshing.	Mount and unmount the export on the client.
MDT-142317	Storage - File	After configuring and running I/O operations to SMB shares and NFS exports events such as an appliance or NAS server reboot, or a file system unmount or remount may cause some negative values to appear in the saved file system performance statistics exported from PowerStore Manager (GUI).	None
MDT-118544	Storage - File	After a request is sent from the REST API with an incorrect GUID, the error message returned displays the invalid UID: 00000000-0000-0000-0000-000000000000.	Ignore the UID in the error message and retry the request with the correct GUID.
MDT-154600	Storage - File	The following alert is sent during upgrade "NAS server [server name] fault tolerance is degraded due to NAS upgrade procedure on a peer NAS cluster node." However, once upgrade is complete the alert is not always automatically cleared.	After upgrade completes, acknowledge the alert to clear it.
MDT-116676	Storage - File	The "NAS node <node name> is down," major alert may be seen after changing the Cluster MTU value. Changing the MTU value may cause degradation or disruption in NAS (file) services.	Wait for the alert to clear before performing any further NAS (File) operations.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-137232	Storage - File	Under rare circumstances, NAS services could be disrupted, if the appliance is close to reaching the maximum used data capacity.	Take one of the following actions on the appliance on which the NAS service is installed: add capacity to the appliance (for example, add drives or expansion enclosures), increase free capacity on the system (for example, deleting unused snapshots, volumes) , or migrate data off the system.
MDT-210085	Storage - File	Unable to add a Tree Quota to a File system because there the description is too large (more than 512 bytes).	Add a shorter description to the Tree Quota.
MDT-195652	Storage - File	Unable to remove the description from an NFS Export.	To change the description, overwrite the current description with a blank character space. Save the NFS Export with no description, then go back and update the description again.
MDT-192427	Storage - File	Occasionally, an automated script used to delete a NAS server fails because it cannot connect to the PowerStore REST API server.	None
MDT-252443	Storage - File	An error occurred when setting a user quota on an SMB share when the user name was created with the SID (security identifier).	Use either one of the following methods to work around this issue: Using the SMB name instead of the Windows SID. Performing a Client connection to the SMB server.
MDT-222985	Storage - File	When creating, or removing an NFS server, from a NAS Server that has been named with over 15 characters, the request to add or remove the NFS Server from the NAS server times out.	Ensure that the NAS server name is less than 15 characters.
MDT-289213	SupportAssist	While attempting to upload data to SupportAssist, a node failover or node reboot occurred. Because of the node failover, or reboot it was unclear if the data successfully uploaded to SupportAssist.	Upload the data to SupportAssist again once the failover, or reboot has completed.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-280394	SupportAssist	If your environment is set up to use the Gateway Connect configuration but is blocking outbound traffic from the Direct Connect host (PowerStore array) connection to Dell support through your gateway server firewall rules, the Remote Secure Credentials feature cannot be enabled properly.	Either change your firewall rules to allow Direct Connect host (PowerStore array) to send out traffic to Dell support or Dell support personnel can use normal remote servicing functions if the customer provides Dell support with credentials for access to the PowerStore cluster. See the PowerStore Security Configuration Guide for information specific to port usage related to SupportAssist.
MDT-148608	SupportAssist	Attempts to upload Support Materials fail. SupportAssist is unable to upload files when there is insufficient free space in the system partition.	Delete old Data Collections to increase the available space in the file system so that SupportAssist can make a backup copy of the file for upload.
MDT-153771	SupportAssist	Enabling either Direct connect with remote access or Gateway connect with remote access may fail.	If enabling either higher tier type, Direct connect with remote access or Gateway connect with remote access fails, wait for the backend issue to clear, then do the following: <ol style="list-style-type: none"> 1. Enable the related lower tier type, Direct connect without remote access or Gateway connect without remote access. 2. Disable the related lower tier type. 3. Re-enable the higher tier type, Direct connect with remote access or Gateway connect with remote access type. If the issue persists, contact your service provider.
MDT-256673	SupportAssist	Cannot enable the SRS gateway on an appliance. Version 3.48 of the SRS gateway has an issue with connectivity from some PowerStore systems.	Use a different version of the SRS gateway.
MDT-109970	Virtualization	Unable to create a vVol storage container in PowerStore, after having just deleted a vVol datastore with the same name from vSphere.	Either create the vVol storage container in PowerStore using a different name or allow some more time before creating the new vVol storage container in PowerStore with the same name that was used in vSphere.
MDT-111956	Virtualization	In rare case, the automated mounting of a user-created PowerStore vVol storage container on an ESXi host could fail with the following error message displayed in vCenter "Timed-out waiting to get datastore information from host".	Manually mount the vVol storage container from the VMware vSphere GUI or CLI.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-156507	Virtualization	In high scale VMware vSphere environments, where the ESXi hosts and vCenter server are highly loaded (for example, simultaneous powering on a large number of VMs), heartbeat messages sent by the ESXi hosts to the vCenter server may not be delivered in time. When the heartbeat is not received in the 60-second timeout interval, the host is considered as not responding, and the host is marked as disconnected. This may result in the vSphere High Availability (HA) to trigger VM migrations to other hosts in the cluster.	In vCenter server, increase the heartbeat timeout interval to 120 seconds. See VMware knowledge base article 1005757: "ESXi host disconnects intermittently from vCenter Server." In case some VMs appear as invalid or orphaned as a result of the ESXi host disconnect, see the VMware knowledge base article 1003742: "Virtual machines appear as invalid or orphaned in vCenter Server." for instructions on recovering the VMs.
MDT-146763	Virtualization	During a vMotion process, some vVols are deleted. If the VASA provider becomes unavailable during a storage vMotion operation, the deletion of the vVols fails. These vVols are left orphaned and may appear as duplicates.	Delete the orphaned vVols using PowerStore CLI or REST API.
MDT-151804	Virtualization	In a PowerStore cluster with multiple VMFS datastores that are mapped to multiple hosts, a performance impact might occur if several VMs are powered off or powered on simultaneously.	Refer to the Dell knowledge base article 000320477: Best practices for VMFS datastores – for Bootstorm or Failover with VMware SRM" for additional details.
MDT-152993	Virtualization	A periodic, automatic cleanup process on vSphere that is performed randomly (every few days) deletes empty config vVols (directories). However, this process also deletes the ISOs folder that is pre-created by PowerStore in case you have not placed any ISO images into the folder.	Recreate the ISOs folder using the following PowerStore CLI command: pstcli -service storage_container -id <id where to create the folder> create_directory -size 100G -name ISOs By default PowerStore creates a 500 GB folder, but you can specify another size if necessary.
MDT-269239	Virtualization	VMware vSphere 7.0 supports SCSI-3 Persistent Reservations (SCSI3-PRs) on a virtual disk (VMDK) level. Support for SCSI-3 PRs enables configurations of a Windows Server Failover Cluster (WSFC) across ESXi hosts with clustered VMDKs as shared disk resources. However, this feature is not currently supported by PowerStore.	None

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-119414	Virtualization	<p>Deploying a VM fails, and checking the vvol.log vCenter log file reveals the following messages:</p> <pre> 2019-11-15T05:24:35.161Z info vvol[2112740] [Originator@6876 sub=Default] VVolAbandonedObjectScan::ProcessName space examining namespace /vmfs/volumes/<storage-container- identifier>/<vvol-identifier> and 2019-11-15T05:24:35.318Z warning vvol[2112740] [Originator@6876 sub=Libs] 2112740:VVOLLIB : VVolLib_OpenObjectTrackingFile:11348: VVolLib_OpenObjectTrackingFile: Failed to open object tracking file /vmfs/volumes/<storage-container- identifier>/<vvol-identifier> /.vvolObjList with error 6 </pre> <p>where <storage-container-identifier> is the UUID of the storage container and <vvol-identifier> is the alphanumeric identifier of the config-vVol of the VM.</p>	<p>Try to deploy the VM again. For more information, see the Dell knowledge base article 000132277: PowerStore: VM deployment can fail with "A general system error occurred."</p>

Limitations

The following limitations exist in this release.

PowerStore T and PowerStore X model limitations

The following limitations exist in PowerStore T and PowerStore X deployment models.

Functional Area	Description
Importing external storage to PowerStore	<p>A maximum of 16 import sessions is supported in the Ready-For-Cutover state during an import. In some cases, when several dozen import operations are run back-to-back, intermittent failures of alternate import sessions may occur. If this occurs, do the following:</p> <p>Remove the remote system and then add it again.</p> <p>Run fewer set of imports (16 or less) at a time. It is recommended to start all these import sessions with automatic cutover turned off.</p> <p>Once all imports have reached the Ready-For-Cutover state, do a manual cutover.</p> <p>After one set of imports is completed, run the next set of imports after a delay of 10 minutes. This delay allows enough time for the system to cleanup any connections to the source system.</p>
PowerStore Clusters	PowerStore X and PowerStore T appliances are not supported in the same cluster.
PowerStore Management access	Use the PowerStore cluster management IP address to access the PowerStore Manager and the PowerStore REST API. The appliance management IPs are intended primarily for service use.
PowerStore scalability limits	The <i>PowerStore Simple Support Matrix</i> lists the scalability limitations in PowerStore models. The PowerStore Simple Support Matrix is available from PowerStore documentation page at: https://www.dell.com/powerstoredocs
Internal migration	Prior to performing an internal migration, please refer to Dell knowledge base article 000105896: PowerStore: Supported Host OS for Non-Disruptive migration of Storage resources for a list of supported configurations and limitations.
Security	DNS is not supported for LDAP.
SupportAssist	SupportAssist cannot be enabled on PowerStore models configured with IPv6 for the management network. SupportAssist is not supported over IPv6.

PowerStore X model limitations

The following limitations exist in PowerStore X deployment models.

Functional Area	Description
Replication of Virtual Volumes (vVols)	Replication of vVols is not supported
Storage – File	NAS services are not supported in PowerStore X models.

Refer to the PowerStore product support pages (<https://www.dell.com/support>) for detailed user documentation.

Environment and System Requirements

All PowerStore deployment models must meet the following environment and system requirements.

Support Matrix

For compatibility and interoperability information, see the *PowerStore Support Matrix* which can be downloaded from the PowerStore Documentation page at: <https://www.dell.com/powerstoredocs>.

Host Configuration requirements

For host connectivity to PowerStore appliances, please refer to the *Dell EMC PowerStore Host Configuration Guide*.

Installation and Upgrade Considerations

Installation considerations

Before you begin to install PowerStore, review the *PowerStore Deployment Checklist*. Once you have completed all the planning steps in the *PowerStore Deployment Checklist*, you are ready to install your system.

The *PowerStore Deployment Checklist* can be downloaded from the PowerStore Documentation page at: <https://www.dell.com/powerstoredocs>.

Upgrade considerations

Before upgrading PowerStore refer to the *Dell EMC PowerStore Software Upgrade Guide*, which can be downloaded from the PowerStore Documentation page at : <https://www.dell.com/powerstoredocs>.

[Clusters already running PowerStore version 2.0.0.0 \(Build: 1376722\) do not need to upgrade PowerStore version 2.0.0.0 \(Build: 1397847\). The changes in the new build are limited to fixes to address problems encountered with upgrading PowerStore software from version 1.0.* to version 2.0.0.0.](#)

PowerStore T model cluster

To upgrade your PowerStore T model cluster to PowerStore version 2.0.0.0 (Build: 1397847):

1. Download the latest PowerStore T software upgrade package named PowerStoreT-2.0.0.0-1397847-retail.tgz.bin from <https://www.dell.com/support>.

Upgrade package name	Type	Description	Required
PowerStoreT-2.0.0.0-1397847-retail.tgz.bin	Software release	PowerStore OS and node firmware updates	Yes

2. Perform the steps in the PowerStore T model clusters section of the *PowerStore Software Upgrade Guide* to upgrade your PowerStore T model cluster.

PowerStore X model cluster

To view a list of compatible PowerStore OS and ESXi/vCenter versions for PowerStore X model appliances, see the *VMware Licensing and Support for PowerStore X* table in the *PowerStore Simple Support Matrix* which can be downloaded from the PowerStore Documentation page at: <https://www.dell.com/powerstoredocs>.

To upgrade your PowerStore X model cluster to PowerStore version 2.0.0.0 (Build: 1397847):

1. Download the latest PowerStore X model software upgrade .zip file named PowerStoreX-2.0.0.0-1397847-retail.zip from <https://www.dell.com/support>.
2. Extract the software upgrade packages from the .zip file to a folder on your computer. The contents of the PowerStore X software upgrade .zip file are:

Upgrade package name	Type	Description	Required
ESXi670-202103001.zip	ESXi update	Dell EMC qualified ESXi software depot, which is used to update the internal ESXi hosts to ESXi 6.7 P05. Note: ESXi 6.7 P05 is the minimum ESXi version required for PowerStore 2.0.0.0.	Required if the internal ESXi hosts are running an ESXi version earlier than ESXi 6.7 P05. Not required if the internal ESXi hosts are running ESXi 6.7 P05 or later.
PowerStoreX-node_firmware-2.0.0.0-1345890-retail.tgz.bin	Hotfix	PowerStore node firmware updates	Required when upgrading from PowerStore OS 1.0.4 or earlier.
PowerStoreX-2.0.0.0-1397847-retail.tgz.bin	Software release	PowerStore OS updates	Required
PowerStoreX-2.0.0.0_README.txt	README file	PowerStore software upgrade README file	Not applicable
PowerStoreX-PreUpgrade_Package-2.0.0.0-1345067-retail.tgz.bin	Hotfix	Custom Dell EMC qualified VIBs for ESXi hosts	Required when upgrading from PowerStore OS 1.0.4 or earlier.

3. Perform the steps in the *Upgrading PowerStore X model clusters* section of the *PowerStore Software Upgrade Guide* to upgrade your PowerStore X model cluster:

The following table lists the order that the software upgrades must be performed and the locations to which the software upgrades packages must be uploaded:

Upgrade order	Upgrade package name	Upload package to
1. Preupgrade	PowerStoreX-PreUpgrade_Package-2.0.0.0-1345067-retail.tgz.bin	PowerStore X model cluster
2. Node firmware	PowerStoreX-node_firmware-2.0.0.0-1345890-retail.tgz.bin	PowerStore X model cluster
3. ESXi update	ESXi670-202103001.zip	vCenter Server
4. PowerStore OS	PowerStoreX-2.0.0.0-1397847-retail.tgz.bin	PowerStore X model cluster

Note: If the ESXi update package ESXi670-202103001.zip has been previously applied to the ESXi hosts on the PowerStore X model cluster, the ESXi update does not need to be performed again during the upgrade to PowerStore version 2.0.0.0. However, the ESXi hosts on the nodes of the cluster must be rebooted to complete the installation of the preupgrade package and node firmware updates before the PowerStore OS is upgraded. For more information, see the Upgrade the PowerStore OS section of the *PowerStore Software Upgrade Guide*.

Where to Get Help

Dell Technologies maintains support pages for all products at <https://www.dell.com/support>.

The product support pages provide important product information such as product and user documentation, knowledge base articles, drivers and other software installation packages downloads, advisories, knowledge base articles, and more.