DELLEMC

Dell EMC PowerStore Release Notes

Release Version: 2.0.1.1 (Build: 1471924)

Release Type: Service (SR)

Contents

Revision History

Document revision	Date	Comments
A00	October 2021	Original release version 2.0.1.1 (Build: 1471924)

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:

- All NVMe™ Platform:
    - 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 & 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.

Resolved Issues

The following issues have been resolved in this release.

Issue ID	Functional Area	Description
MDT-342090	Data Collection	When an error occurs while reading data from the backend drives, a background process is scheduled to fix the data so that later reads don't have to fix the error. There is a rare possibility that the background process won't handle a situation correctly which will cause some data to be lost.
MDT-338925	Hardware	After the system is rebooted, followed by a double NVRAM failure, the system does not come back up, leading to data unavailability. In certain cases some data will be lost.
MDT-342095	Hardware	After the system data is corrupted a node becomes unavailable.
MDT-335133	Install or Upgrade	While upgrading from PowerStore 1.x to PowerStore 2.0.x, if the upgrade rolls back to PowerStore 1.x, the node which upgraded to PowerStore 2.0.x and then rolled back will be put into service mode.
MDT-336500	Monitoring	Customer received an alert, that the "Node's IO service level has changed," which was caused because lock leakage in the System Manager (SYM) caused a restart after System Manager ran for 75.86 days. The alert auto cleared without any manual intervention.

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.

Issue ID	Functional Area	Description	Workaround/Resolution
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	Booting from SAN is not supported with FC-NVMe on AIX in PowerStore 2.0.0.x.	None
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-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

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-316424	Data Protection	Customer network environment may drop jumbo frames or block them from being transmitted over the network. As a result, the replication session status can change to System Paused during the replication session synchronization.	Decrease cluster MTU size to 1500. In PowerStore 2.0 and higher, there is an option to configure specific ports for replication and set a dedicated VLAN with a different MTU. It is recommended to separate the replication ports and configure the MTU to 1500 to allow flawless replication and keep MTU 9000 for the local area network only
MDT-312619	Data Protection	When replication IPs change, they are not automatically detected. As a result, a replication session's status changes to 'System Paused'.	Run 'Remote System Verify' operation to update the connectivity matrix. The next synchronization cycle will update the replication IP addresses.
MDT-309462	Data Protection	When the user pauses a replication session on the destination cluster during NDU of the source cluster, a post-NDU state may occur, where the source replication session is in System Paused state and the destination replication session is in Paused state.	Pause the replication session on the source cluster (where the session is in System Paused state) and then resume the replication session.
MDT-299182	Data Protection	In a large configuration, pausing or deleting a replication session may take longer when the session is queued for synchronization.	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-220994	Data Protection	The replication session may not be deleted once a Protection Policy is unassigned from the storage resource.	To delete the replication session, do the following: 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.

Issue ID	Functional Area	Description	Workaround/Resolution
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-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.
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.	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. 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.

Issue ID	Functional Area	Description	Workaround/Resolution
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.
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-303168	Internal Migration	The status of a vVol migration session might change to System Paused when a large number of migrations are started simultaneously.	Power down the VM to which the vVol belongs and resume the migration session.
MDT-298613	Internal Migration	A vVol migration session might become stuck with a status of Cutting Over when a large number of vVols are being migrated simultaneously.	Contact your service provider if a vVol migration session becomes stuck with a status of Cutting Over.
MDT-316865	Internal Migration	Migration of a volume may be interrupted indefinitely during the migration session sync operation. This may occur if a backend service restarts during the migration session sync.	Contact your service provider.
MDT-285903	Internal Migration	The migration failed due to a 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-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.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-242178	Internal Migration	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.	Reboot the ESXi host to send an unbind request for the vVol and to allow subsequent migrations of the vVol to the appliance. **Note**: If the cluster contains more than two appliances, the vVol can be migrated to another appliance in the cluster.
MDT-239624	Internal Migration	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: - 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.	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.
MDT-344382	Install or Upgrade	Some appliance models with very large physical capacity may experience a failed upgrade with data unavailability when attempting to upgrade from: 1.0.x.x to 2.0.x.x.	Pre-Upgrade Health Check (PUHC) prevents upgrade to PowerStore 2.0.1.1 when the system capacity is too large. Contact your service representative to continue to upgrade.
MDT-347271	Install or Upgrade	Some appliance models with very large physical capacity may experience a failure to initialize when attempting to boot PowerStore 2.0.x.x for the first time, which will result in data unavailability.	Contact your service representative.
MDT-314978	Install or Upgrade	If an appliance is removed from a PowerStore cluster using the svc_remove_appliance script, the management database must be modified before you can successfully upgrade the PowerStore software on the cluster.	Contact your service provider and refer to Dell knowledge base article 000183374: PowerStore: Cannot NDU after having run Remove Appliance

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-288380	Install or Upgrade	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).	Contact your service provider for assistance.
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-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.
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.

Issue ID	Functional Area	Description	Workaround/Resolution
MDT-282118	Security	Timeout error occurs when trying to log into PowerStore using the LDAP directory service. 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.	Be sure that the Advanced Settings in the PowerStore Manager, Directory Services is configured as follows: 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. Set User Search Settings>Search Path as close as possible to the LDAP user. Set Group Search Settings>Search Path as close as possible to the LDAP groups that have LDAP user as a member. If the above steps still do not help, try to set Group Search Settings> Nested Group Search Level to 1
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
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.

Issue ID	Functional Area	Description	Workaround/Resolution
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.
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.

Issue ID	Functional Area	Description	Workaround/Resolution
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.
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-325206	SupportAssist	When using SupportAssist with a Gateway, SupportAssist connectivity is lost when DNS values are modified.	Disable SupportAssist prior to making a DNS change.
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: 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.
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	Deploying a VM fails, and checking the vvold.log vCenter log file reveals the following messages: 2019-11-15T05:24:35.161Z info vvold[2112740] [Originator@6876 sub=Default] VVolAbandonedObjectScan::ProcessNamespace examining namespace /vmfs/volumes/\<storage-container-identifier>/\<vvol-identifier> and 2019-11-15T05:24:35.318Z warning vvold[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 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.	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."

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	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: Remove the remote system and then add it again. 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. Once all imports have reached the Ready-For-Cutover state, do a manual cutover. 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.
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 *PowerStore Simple Support Matrix* 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 *PowerStore:Supported Host OS for Non-Disruptive migration of Storage resources* at: https://www.dell.com/support/article/en-us/how17129/ 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.

PowerStore T model cluster

To upgrade your PowerStore T model cluster to PowerStore version 2.0.1.1 (Build: 1471924):

1. Download the latest PowerStore T software upgrade package named PowerStoreT-2.0.1.1-1471924-retail.tgz.bin from https://www.dell.com/support.

Upgrade package name	Type	Description	Required
PowerStoreT-2.0.1.1-1471924-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 2.0.1.1 (Build: 1471924):

1. Download the latest PowerStore X model software upgrade .zip file named PowerStoreX-2.0.1.1-1471924-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.x.x.	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.x.x-1345890-retail.tgz.bin	Hotfix	PowerStore node firmware updates	Required when upgrading from PowerStore OS 1.0.4 or earlier.
PowerStoreX-2.0.1.1-1471924-retail.tgz.bin	Software release	PowerStore OS updates	Required
PowerStoreX-2.0.x.x_README.txt	README file	PowerStore software upgrade README file	Not applicable
PowerStoreX-PreUpgrade_Package-2.0.x.x-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.x.x-1345067-retail.tgz.bin	PowerStore X model cluster
2. Node firmware	PowerStoreX-node_firmware-2.0.x.x-1345890-retail.tgz.bin	PowerStore X model cluster
3. ESXi update	ESXi670-202103001.zip	vCenter Server
4. PowerStore OS	PowerStoreX-2.0.1.1-1471924-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.x.x. 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

The Dell Technologies Support site (https://www.dell.com/support) contains important information about products and services including drivers, installation packages, product documentation, knowledge base articles, and advisories.

A valid support contract and account might be required to access all the available information about a specific Dell Technologies product or service.