

Dell EMC PowerStore Release Notes

Release Version: 1.0.1.0.5.002

Release Type: Service (SR)

Contents

| | |
|--|----|
| Revision History | 1 |
| Product Description | 2 |
| Resolved Issues | 3 |
| Known Issues..... | 5 |
| Limitations..... | 14 |
| Environment and System Requirements | 15 |
| Installation and Upgrade Considerations..... | 15 |
| Where to Get Help | 16 |

Revision History

| Document revision | Date | Comments |
|-------------------|-----------|--|
| A01 | July 2020 | Added MDT-163759 to the list of Resolved Issues. |
| A00 | June 2020 | Original release version 1.0.1.0.5.002 |

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 PowerStoreOS
 - Block, File, and vVols
 - NVMe Flash and Storage Class Memory (SCM) media support
 - 25 Gbe iSCSI and 32Gb FC support
 - 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 | 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 (<http://www.dell.com/support>) for detailed user documentation.

Resolved Issues

The following high severity issues have been resolved in this release.

| Issue ID | Functional Area | Description |
|------------|----------------------|---|
| MDT-148696 | Connectivity - Hosts | When there is an active I/O operation occurring between a host and 4-port 32Gb Fibre Channel I/O module on the appliance, and there is a considerable amount of logout and login activity initiated by the hosts, the port may report a BUSY status back to the host even when the port is not oversubscribed from a limits standpoint. |
| MDT-153885 | Internal Migration | If you expand a volume while one of the nodes is booting up, the size of that volume might not change on the node that is booting up. As a result, I/O operations to the newly expanded region of the volume fail if they are sent to that node. |
| MDT-153333 | Data Collection | As the number of targets scale, node CPU stats may include underlying target stats multiple times inflating the value and causing a discrepancy between appliance CPU stats and its underlying node CPU stats. |

| Issue ID | Functional Area | Description |
|------------|-------------------------|---|
| MDT-155853 | Monitoring | If there is a critical capacity-based forecasting alert, a migration recommendation repair flow can be used to migrate storage to another appliance. It is possible that after migration completes, the daily forecast used to trigger the alert will not be updated on the source appliance and the alert will not be cleared until the next daily forecast. After the forecast is triggered again, the capacity based forecast on the source and destination appliance may be inaccurate as a significant amount of storage is migrated |
| MDT-134036 | Monitoring | When there is a contention of resources, the performance policy associated with the volume gets applied. In the current release, you may not notice much difference between the High and Medium performance policies. For volumes with Low performance policy, expect to see a noticeable difference in the amount of I/O operations serviced as compared to volumes with High or Medium performance policies. |
| MDT-153140 | PowerStore Manager(GUI) | The Check for software upgrade option cannot be selected from the ? icon in the top right of PowerStore manager. |
| MDT-144836 | Storage - File | If a snapshot expiration is in progress when the NAS Server has stopped due to issues such as Failover, Failback, or the NAS server is being moved to another node, then a "NAS node [nodename] is down," alert is displayed in PowerStore Manager and NAS management operations become unavailable. |
| MDT-149396 | Support | If an event or issue abruptly interrupts the node while the node, appliance, or cluster is powering off, the node may not complete powering off and a reboot may occur instead. |
| MDT-153919 | SupportAssist | In multi-appliance clusters, data collections taken on a peer appliance may not be automatically uploaded to SupportAssist. |
| MDT-163759 | SupportAssist | Configuring SupportAssist for Gateway Connect with remote access to gateways running SupportAssist 4.0.5 fails on appliances with operating systems earlier than 1.0.1.0.5.002. |

Known Issues

The following high severity issues are reported in this release. For a list of the minor issues, see Knowledge Base article SLN320772 PowerStore Release Notes – Minor issues at: <https://www.dell.com/support/article/SLN320772>.

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|------------------------|--|---|
| MDT-152854 | Connectivity - Hosts | Running an unmap operation might cause a host I/O latency impact of more than 15%. You can disable auto-unmap for applications that have this feature, or manually schedule an unmap operation to reduce the impact on hosts. | Disable auto-unmap for applications that have this feature. Schedule unmap operations manually to reduce the impact on hosts. |
| MDT-154342 | Connectivity - Hosts | When you add a host, in rare cases, you may not see any active sessions or connected paths for the initiators for up to 10 minutes on the Initiators tab. | Do not try adding the host again. Wait for up to 10 minutes for the active sessions to appear. |
| MDT-139488 | Connectivity - Hosts | Xcopy operations can impact host I/O performance. | None |
| MDT-171399 | Connectivity- Networks | The <i>PowerStore Network Configuration for Dell PowerSwitch Series Guide</i> provides examples with the flowcontrol receive set to on. However, it is recommended to set flowcontrol receive off and flowcontrol transmit off for best performance. | Do not follow the flowcontrol settings from the examples in the PowerStore Network Configuration for Dell PowerSwitch Series Guide rather reference the <i>PowerStore Host Configuration Guide</i> for the correct flowcontrol settings. |
| MDT-168926 | Connectivity- Networks | A cluster "Network validation failed with unexpected error (0xE0E70001)" message is returned erroneously while validating initial configuration. The error occurs because an internal entry for DNS was previously configured. | Ignore the warning and continue configuring the cluster. After the cluster is created you can validate that the DNS configuration was successful in the PowerStore Manager. |
| 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-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-148947 | Data Protection | <p>From a protection perspective, a snapshot of a volume group is a single-point-in-time copy of the volume group member volumes. Membership changes to the volume group prevent snapshot recovery operations, such as restore or refresh, unless the volume group is returned to the same membership state at the time the snapshot was created. It is recommended that you do not delete the volumes removed from a volume group until all the snapshots of the volume group either expire or are removed manually.</p> <p>If there are still volume group snapshots, the delete request for the volume returns the following error: "The Volume <volume name> (id: <unique id>) cannot be deleted until all volume group snapshots taken when this volume was a member of a volume group expire or are manually deleted." This error does not indicate the name of the volume group that is needed to delete the volume group snapshots so the volume deletion can succeed.</p> | <p>In PowerStore, first identify the volume group name, if it is unknown, by navigating to Snapshots table of the removed member volume. Volume group snapshots that have been created by a protection policy should have the volume group name as part of the system-created snapshot name for each member volume. If there are only manual snapshots, you can use REST API to identify the volume group name.</p> <p>The volume group snapshots can be deleted by navigating to the protection tab and viewing the snapshots for the specific volume group. Click Show/Hide Table Columns and add the Volume Members column. The column appears and displays a number. When you hover over that number, the volumes within the VG snapshot become visible. You can look for a pattern by sorting this column. A column with more members might indicate that the VG snapshot has the member that is preventing the volume deletion operation, assuming one or more members have been removed and snapshots continued to be created on the remaining VG members.</p> |
| MDT-146780 | Data Protection | <p>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.</p> | <p>Delete the snapshot rule again. The same error appears, but the rule will be deleted.</p> |
| MDT-117061 | Hardware | <p>When beginning cluster creation, the hardware status is incorrectly indicated as not configured for clustering.</p> | <p>Wait a few minutes and try again.</p> |
| MDT-137439 | Import | <p>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.</p> | <p>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.</p> |

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|--------------------------|--|--|
| MDT-164726 | Installation and 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 an error because of a localization issue with the event message. | None |
| MDT-168710 | Internal Migration | In a PowerStore T cluster with more than one appliance, 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-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-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. | Wait for migration to complete to get the correct calculation. |
| MDT-140927 | Internal Migration | After a volume is migrated, there maybe a dip in capacity and performance metrics reporting for that volume at the end of the migration operation. | None |
| 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 or 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. |

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|--------------------------|---|--|
| MDT-148659 | Install or Upgrade | If the peer node is inaccessible for too long, the software upgrade can fail. | Retry the software upgrade. |
| 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-153335 | PowerStore CLI | The PowerStore CLI does not show all the required information for metrics when a roll-up result is returned for long interval values. | It is recommended that you use the REST interface to query metrics. You can also use any REST client (cURL, Insomnia, Postman, etc.) or web browser. |
| 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-86305 | PowerStore Manager(GUI) | 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. | On the Hardware tab, click the Refresh icon (appears on the top right corner of the tab, next to the base enclosure image) or refresh the browser page. |
| MDT-142069 | PowerStore Manager(GUI) | There are situations when the system is Out of Space where the UI is unable to update its statistics. This situation will be evident by a red banner on the top of the screen. You may see objects in the UI that are incorrectly displayed as "0 MB/GB/TB" due to the inability for the UI to query for system statistics. | None |
| 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. | When refreshing a resource, do not select the "Choose One" entry in the Refresh from Volume/Volume Group drop-down. |
| 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-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. |

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|-----------------|--|--|
| 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 ColdStream drive, for example, 7 minutes for 375 G, 15 minutes for 750 G, and 27 minutes for 1.5T. 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 + about 1 minute for other overhead), and then remove the drive. The drive will then become disconnected and automatically removed from the system. |
| MDT-170429 | Storage - File | During manual failover, or failback of the NAS server, the failover, or failback operation may fail. | Wait a few moments and retry the operation. |
| 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-147508 | Storage - File | SMB Server cannot join the Active Directory when the SMB server computer name is more than 15 characters and the first 15 characters of the SMB server computer name are the same as another SMB server's computer name. | Use a unique SMB server Netbios name with 15 characters or less. |

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|-----------------|---|---|
| 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 umount 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-146793 | Storage - File | Uploading configuration files to a NAS server, which were edited in WordPad, does not work as expected. | Use vim or Notepad to edit configuration files before uploading them to the NAS server. |
| 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. |
| 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. |

| Issue ID | Functional Area | Description | Workaround/Resolution |
|------------|-----------------|---|--|
| MDT-151302 | Storage - File | If the network connection is unstable, the DNS may incorrectly be reported as offline. After restarting the NAS server, the health status of DNS is not cleared if the DNS is healthy. | Refer to the Dell knowledge base article SLN320669. Contact your service provider, if you need assistance. |
| MDT-163759 | SupportAssist | The SupportAssis Enterprise 4.00.05.30 gateway is not supported by PowerStore models. | Connect PowerStore models to SupportAssist 3.38 gateway. |
| MDT-148608 | SupportAssist | Attempts to upload Support Materials fail. Support Assist 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 Support Assist 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-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. In case some VMs appear as invalid or orphaned as a result of the ESXi host disconnect, see the VMware Knowledge Base article 1003742 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 SLN320477: "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 Trident in case you have not placed any ISO images into the folder. | Recreate the ISOs folder using the following PowerStore CLI command: <pre>pstcli -service storage_container -id <id where to create the folder> create_directory -size 100G -name ISOs</pre> By default PowerStore creates a 500 GB folder, but you can specify another size if necessary. |

| 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></pre> <p>and</p> <pre>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 SLN320647.</p> |
| MDT-146912 | Virtualization | <p>In PowerStore X model, when both ESXi hosts of the appliance are rebooted simultaneously (for example, due to a power failure), ESXi iSCSI parameters (path selection policy, Round Robin IOPS limit, number of outstanding IOs with competing worlds) of one or both nodes may be reset to the default ESXi values. This may affect IO performance.</p> | <p>Refer to the Dell knowledge base article SLN320780 for more information on resetting the values of the ESXi iSCSI parameters.</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 Management access | <p>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.</p> |
| PowerStore scalability limits | <p>The PowerStore Simple Support Matrix lists the scalability limitations in PowerStore models. The PowerStore Simple Support Matrix is available from the Dell E-Lab Navigator at: https://elabnavigator.emc.com/elN/elNhome</p> |
| Internal migration | <p>Prior to performing an internal migration, please refer to <i>PowerStore:Supported Host OS for Non-Disruptive migration of Storage resources</i> at: https://www.dell.com/support/article/en-us/how17129/ for a list of supported configurations and limitations.</p> |
| SupportAssist | <p>SupportAssist cannot be enabled on PowerStore models configured with IPv6 for the management network. SupportAssist is not supported over IPv6.</p> |

PowerStore X model limitations

The following limitations exist in PowerStore X deployment models.

| Functional Area | Description |
|--|--|
| Clustering | <p>Clustering with multiple PowerStore X appliances is not supported</p> |
| Replication of Virtual Volumes (vVols) | <p>Replication of vVols is not supported</p> |
| Storage – File | <p>NAS services are not supported in PowerStore X models.</p> |

| Functional Area | Description |
|------------------------------|--|
| Support for Virtual Machines | Virtual Machines (VMs) running on PowerStore X internal nodes only supports Virtual Volumes (vVols). Virtual Machine File System (VMFS) is not supported for internally running VMs, but is supported when using PowerStore X with external ESX servers attached via Fibre Channel or iSCSI. |

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

Refer to the *PowerStore Support Matrix* on the support website for compatibility and interoperability information.

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>.

On a PowerStore T model cluster with NAS services enabled, if changes are made to the CHAP settings, such as enabling or disabling CHAP, changing mutual CHAP to single CHAP, or changing single CHAP to mutual CHAP, you must:

- Reboot the cluster nodes one at a time as soon as possible after changing the CHAP settings.
- Perform a nondisruptive upgrade (NDU) only after the nodes have been rebooted.

If the cluster nodes are not rebooted after changes are made to CHAP settings, an NDU will fail.

Where to Get Help

Dell Technologies maintains support pages for all products at 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.

A valid support contract and registration is required to access all information available on the product support sites.